ESSAYS ON THE REAL EFFECTS OF INFORMATIONAL FRICTIONS IN FINANCIAL MARKETS

Michael Daniel Sockin

A DISSERTATION PRESENTED TO THE FACULTY OF

PRINCETON UNIVERSITY IN CANDIDACY FOR THE

DEGREE OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE BY THE

DEPARTMENT OF ECONOMICS

Adviser: Wei Xiong

June 2015

Abstract

This dissertation aims at characterizing the real consequences of informational frictions in financial markets through several applications. The first chapter highlights a dynamic feedback loop in learning that arises between financial markets and real investment that helps explain how financial crises can contribute to slow recoveries. The second chapter investigates feedback effects from learning in commodity markets to understand the role of financial market speculation in the commodity boom and bust cycles of the late 2000's. The third chapter explores the role of informational frictions in housing markets to rationalize the hump-shaped pattern in price volatility observed across supply elasticity in the recent US housing cycle.

Acknowledgements

I am deeply indebted to my adviser, Wei Xiong, for his immeasurable guidance and support in shepherding me through the research leading to this dissertation. I would also like to thank my dissertation committee, Mikhail Golosov and Stephen Morris, for their invaluable advice. I am sincerely grateful to my coauthor, Zhenyu Gao, for his support and encouragement. I would also like to thank Markus Brunnermeier, Maryam Farboodi, Valentin Haddad, Ben Moll, my fellow PhD students, and other seminar participants for their helpful comments and advice. I thank my parents, my brothers, and my sister for their unconditional love.

Table of Contents

Introduction

In the presence of informational frictions, centralized asset markets act as platforms to aggregate the private information of economic agents, and financial prices serve as useful signals about the underlying strength of the real economy. Through this learning channel, noise in financial prices can feed into the real decisions of firms and households, and, by distorting these decisions, feed back into the financial prices that anchor on them. The central research objective of this dissertation is to understand how learning from financial prices impacts real activity. The first chapter examines how learning in financial markets affects broad macroeconomic growth, while the second focuses on feedback effects from learning in commodity markets, and the third on the interaction between learning and supply elasticity in housing markets.

The first chapter is motivated by the lingering effects of the recent recession and financial crisis on the US economy. The model that the chapter develops highlights a dynamic feedback loop in learning that arises between financial markets and real investment that helps to explain how financial crises can contribute to slow recoveries. This occurs because both real and financial signals flatten as a recession deepens after a negative liquidity shock drives down asset prices, making it more difficult for economic agents to act on signs of the recovery. The analysis further highlights that informational frictions give rise to asset return predictability with business-cycle variation and lower welfare by affecting both aggregate growth and cross-sectional inequality. In addition, the welfare analysis illustrates a potential role for policy intervention through several experiments that arises because economic agents do not internalize that their real investment and financial trading decisions impact the provision of public information.

Motivated by the large inflow of financial capital into commodity futures markets in the early 2000's, and the run-up in oil prices in the first half of 2008, the second chapter of the dissertation investigates feedback effects from learning in commodity markets. The informational role of commodity prices gives rise to a channel for futures market speculation to influence spot prices, challenging the conventional wisdom that futures markets are just a shadow of spot markets. This channel also helps clarify how a bubble in commodity markets can persist without a buildup in inventory, since informational frictions can cause real consumers of oil to have a nonnegative price elasticity of demand. A version of this chapter, which is joint work with Wei Xiong, is forthcoming in the Journal of Finance.

Finally, the third chapter develops a model to explore the interaction between informational frictions and supply elasticity in housing markets. Since housing prices are driven entirely by demand shocks in perfectly inelastic markets, and by supply shocks in perfectly elastic markets, distortions to housing prices and demand from learning occur at intermediate supply elasticities. This mechanism helps rationalize the puzzling observations from the recent housing US cycle that counties with intermediate supply elasticities experienced the most dramatic price booms and busts. An extension of the baseline model introduces migrants into the neighborhood to help explain the hump-shaped variability found in the cross-section of investment home purchases.

This dissertation aims to contribute to our growing understanding of the real effects of financial markets in the presence of informational frictions. While many have examined the feedback effects from real investment to financial markets or from financial markets to real investment in the presence of informational frictions, there is relatively little work that examines the dynamic consequences for real activity when there is feedback in both directions. The first chapter helps to fill this gap. The second chapter adds to the ongoing debate about the effects of the financialization of commodities by introducing a mechanism for financial speculation in commodity markets to distort the real decisions of production firm and, through this informational channel, commodity spot prices. The third chapter documents several new empirical features of the recent US housing cycle that cannot be explained by conventional theories of how supply elasticity impacts housing markets, and provides a mechanism based on learning to reconcile these facts. In addition, the dissertation develops new tractable locally linear and log-linear equilibrium frameworks for studying informational frictions with real effects that may be fruitful in future research.

Chapter 1

A Model of Growth and Informational Frictions

1.1. Introduction

In this paper, I introduce a tractable, dynamic framework for studying the feedback loop in learning that occurs between financial markets and firm managers when financial markets aggregate investor private information about the productivity of real investment. Through this informational channel, financial market prices are more important for learning than real activity at the trough of business cycles, and are most informative as signals about investment productivity during downturns and recoveries. My analysis establishes a link between recessions with financial origins and slow recoveries by illustrating how financial crises during downturns can delay recoveries by distorting firm manager expectations, which depresses real investment and feeds back into the incentives for financial market participants to trade on their private information.

Two observations motivate my investigation. The first is that market prices aggregate the private information of investors about macroeconomic and financial conditions, and that firms, in making their real decisions, respond to this useful information.[1] Since the mid-1980's, however, the rapid growth of the market-based financial system (Pozsar et al 2012), especially from 2002-2007 (Philippon (2008)), has increased financial opacity, as intermediaries extended credit and diversified risk through securitization and the OTC derivatives

[1] See, for instance, Luo (2005), Chen, Goldstein, and Jiang (2007) and Bakke and Whited (2010). For evidence that firms learn from their own profit realizations, the other key signal in our model, see, for instance, Moyen and Platikanov (2013).

markets that arose in the wake of LTCM.[2] This heightened opacity has made it difficult for economic agents and policymakers to assess not only the depth of financial distress once a bust occurs, but also its distribution across the financial sector. This was particularly relevant in the recent recession, as regulators scrambled to map out the cross-party linkages of the unregulated financial system in late 2008 (FCIC 2011). As a result, market prices have become noisier signals about the strength of the economy, and economic actors, both real and financial, face more severe informational frictions.

That asset prices contain useful information about the macroeconomy has been well-documented in the literature.[3] Both during and in the aftermath of the financial crisis, many viewed the dramatic fall in asset prices as a signal that the US economy was entering a recession potentially as deep as the Great Depression.[4] When the stock market bottomed out in March 2009, in fact, the Michigan Survey of Consumers "fear of a prolonged depression" question had its lowest score since the 1991 recession.

The second observation is that recessions with financial origins appear to be deeper and have slower recoveries. A salient feature of the recent US experience, for instance, is the anemic economic recovery compared to previous cycles, especially in GDP, lending, and productivity (Haltmaier (2012), Reifschneider et al (2013)). As highlighted in a speech by former Federal Reserve Chairman, Ben Bernanke, this weak growth in productivity following the 2007 to 2009 recession represents "a puzzle whose resolution is important for shaping expectations about longer-term growth" (Bernanke (2014)). While there is growing evidence that financial crises lead to deeper recessions, however, it is less clear if, and how, they

[2] Former FRBNY President and Treasury Secretary Timothy Geitner, in fact, made it part of his agenda before the financial crisis to move the OTC derivatives market onto exchanges to increase transparency.

[3] For stock prices, for instance, see Fama (1981), Barro (1990), and Beaudry and Portier (2006), while for credit spreads, see Gertler and Lown (1999), Gilchrist, Yankov, and Zakrasjek (2009), Gilchrist and Zakrajsek (2012), and Ng and Wright (2013), and for a wide cross-section of asset classes, see Stock and Watson (2003) and Andreu, Ghysels, and Kourtellos (2013).

[4] For evidence regarding the fall in the stock market, see, for instance, Robert Barro's March 2009 WSJ Article "What Are the Odds of a Depression?" that accompanies Barro and Ursúa (2009), and Gerald Dwyer's September 2009 article, "Stock Prices in the Financial Crisis" from FRB Atlanta's Notes from the Vault.

also slow recoveries.[5] My model provides a framework for addressing conceptual questions about business cycles and uncertainty that explicitly incorporates a financial sector, and can also help explain why financial shocks have asymmetric impacts over the business cycle (Aizenman et al (2012)).[6] The uncertainty I consider here that distorts investment arises from learning, and is therefore different from that in Bloom (2009), which focuses on shocks to firm fundamental volatility. It is also different from the policy uncertainty featured in Fernández-Villaverde et al (2013), which is over future corporate tax policies, and in Baker, Bloom, and Davis (2013).

Informational frictions can lead firms to voluntarily withdraw from investment because of weak expectations about the state of the economy, rather than from uncertainty itself, a phenomenon which can help explain several stylized facts. First, the FRB Senior Loan Officer Survey cites weak credit demand as a reason for the low level of C&I loans until the end of 2010. Second, since the recession, firms have been increasing the cash on their balance sheets and saving their income as retained earnings rather than investing (Baily and Bosworth (2013), Sanchez and Yurgadul (2013), Kliesen (2013)).[7] Third, firms appear reluctant to fill vacancies, as studies such as Daly et al (2012) and Leduc and Liu (2013) find a potential shift in the Beveridge Curve after the recent recession, which reflects a higher vacancy rate compared to the unemployment rate, while Davis, Faberman, and Haltiwanger (2013) document a fall in recruiting intensity. Though, for simplicity, my model will only involve capital, the same forces depressing real investment would also depress labor market demand in a more general framework. This evidence suggests that the slow recovery may, at

[5] While studies like Reinhart and Rogoff (2009a,b, 2011), Ng and Wright (2013), and Jorda, Schularick and Taylor (2013), for instance, argue that financial crises result in slower recoveries, others such as Haltmaier (2012) and Stock and Watson (2012) find little difference, and those such as Bloom (2009), Muir (2014), and Bordo and Haubrich (2012) predict faster upswings following financial crises.

[6] For instance, while the S&L crisis and the bursting of the housing bubble accompanied recessions that had slow recoveries, the collapse of Long-Term Capital Management (LTCM) in 1998, arguably an event that almost led to the meltdown of the whole financial system, had no significant impact on the real economy.

[7] Pinkowitz, Stulz, and Williamson (2013) provide evidence that this increase in cash holdings is driven by perceived low investment opportunities by firms, since it is concentrated among the highly profitable firms in their sample.

least in part, be driven by firms choosing to delay investment because of a persistent poor economic outlook.

To study the implications of learning in the presence of informational frictions for financial market trading and real activity, I integrate the classic information aggregation framework of Grossman and Stiglitz (1980) and Hellwig (1980) into a standard, general equilibrium macroeconomic model in continuous-time. This setting allows me not only to examine the dynamic, real consequences of informational frictions when there is a feedback loop between real activity and financial markets, but also to depart from the CARA-normal and risk-neutral-normal frameworks, which are less desirable for addressing macroeconomic questions, and to study agents with log utility without the need for approximation. Both tasks have posed a well-known and substantial challenge in the information aggregation literature, and separate strands have developed to examine feedback in each direction. A finance literature, including Albagi (2010), Goldstein, Ozdenoren, and Yuan (2013), and Subrahmanyam and Titman (2013), examines how asset prices impact real activity through the learning channel, while a macroeconomic literature, including Angeletos, Lorenzoni, and Pavan (2012), investigates how real investment decisions are distorted by the ability to manipulate asset prices in the presence of informational frictions. I am able to make progress by appealing to the local linearity inherent in working in continuous-time, as well as to a standard assumption about the information structure of households and a convenient functional form for firm real investment.

The model presented herein features a continuum of non-overlapping generations of households that trade riskless debt and claims to the assets of firms in centralized financial markets. Households here represent the hedge funds, financial analysts, intermediaries, and other investors that participate in financial markets. Households each possess a private signal regarding the underlying strength of the economy when they trade, and are subject to preference shocks that reflect their private liquidity needs. Asset prices in my economy aggregate the private information of agents, and liquidity shocks represent a source of noise that prevents them from being fully revealing to both households and firms. To avoid both the infinite regress problem of Townsend (1983) and a time-varying correlation between the wealth of households and the persistence of their beliefs, I follow Allen, Morris, and Shin (2006), Bacchetta and van Wincoop (2008), and Straub and Ulbricht (2013) and assume that, though households in each generation pass along their wealth to their children, they do not pass along their private information. This assumption of investor myopia is necessary to maintain tractability in learning by helping me avoid these two issues.

Perfectly competitive, identical firms in my economy produce output and are run by managers who use financial prices, which aggregate private information dispersed among households, and real signals from production to form their expectations about the underlying state of the economy when making investment decisions. This introduces a channel for liquidity shocks from financial markets to feed into real activity by distorting the expectations of firm managers, since the impact of financial shocks on prices cannot be fully disentangled from fundamental trading. By affecting the returns on their securities and the informativeness of real economic signals through their investment choices, firms, in turn, impact the incentives of investors to trade on their private information to take advantage of the uncertain economic environment. This can lead to an adverse feedback loop that exacerbates real shocks to the economy during downturns that can deepen and lengthen recessions.

With these ingredients, I derive a tractable, linear noisy rational expectations equilibrium that offers several insights about learning from real and financial signals over the business cycle when there is this feedback loop. First, time-varying second moments are important for macroeconomic dynamics even without the real-options "wait-and-see" channel of Bernanke (1983) and Bloom (2009) for investment. In most environments with learning and asymmetric information, the conditional variance of beliefs is either constant or deterministically converging toward a (possibly trivial) limit. In my setting, this conditional variance varies stochastically with the level of investment, and this gives rise to countercyclical uncertainty in the economy. The second insight is that, while real signals about the macroeconomy are procyclical in their informativeness in learning, similar to the mechanism in Van Nieuwerburgh and Veldkamp (2006), financial signals are strongest during downturns and recoveries. This feature arises because households have dispersed information and trade more aggressively against each other when there is uncertainty about the state of the economy, and this increase in trading leads more of their private information to be incorporated into prices. The strength of the financial signal trades off the return to investment with the level of uncertainty in the economy, and these two quantities are negatively correlated over the business cycle. Finally, nonlinearity in investment slows recoveries since the informativeness of real and financial signals is tied to real investment. As investment falls, both real and financial signals weaken, which leads uncertainty to remain high and persistent until investment recovers. Real signals flatten because firms are less active, and financial signals flatten because the value of household private information anchors on the return to real investment.

I next offer an explanation of the slow US recovery in the context of my mechanism as stemming from confusion in financial price signals brought about by the financial crisis. This confusion led real investment to fall further during the recession and real and financial signals to flatten, which made it more difficult for agents to act on the recovery. I characterize welfare in the economy and identify a role for policy in improving the provision of public information about current economic conditions, since investors and firms do not fully internalize the benefit of the information that their activities produce.

Lastly, I turn to some of the empirical implications of my framework. I illustrate how informational frictions give rise to an informational component in risk premia. This component has predictive power for future returns and real activity, which varies with the level of uncertainty and investment in the economy. It also gives rise to business cycle varation in asset turnover based on informational trading. I then conclude by discussing how taking advantage of the business cycle behavior of financial market signals can help macroeconomic forecasting, as well as conceptual issues that informational frictions raise for identifying structural shocks originating from financial markets.

1.2. Related Literature

I view my amplification mechanism from feedback in learning as playing a contributing role in transmitting financial shocks to the real US economy to bring about deeper recessions and anemic recoveries, and frame it as being complementary to other channels highlighted in the macroeconomics literature linking recessions and financial crises. My paper is also part of several literatures on asymmetric information and the real consequences of asset prices. I discuss my relation to each of these literatures in turn.

Most such studies focus on the balance sheet and/or collateral channels for financial crises to amplify real shocks and depress real activity. He and Krishnamurthy (2012), for instance, explores the quantitative impact of the balance sheet channel for constrained intermediaries, while Mian and Sufi (2012) examines empirically how the deleveraging of household balance sheets can prolong recessions through debt overhang. A slow recovery explained purely by intermediary balance sheet impairment, for instance, is difficult to reconcile with the quick recapitalization of banks by early 2009 because of the TARP and SCAP programs. An explanation based purely on credit constraints confronts the empirical challenges that C&I loan terms had, on average, loosened to around 2005 levels by mid-2011, according to the FRB Senior Loan Officer Survey, and that corporate bond markets continued to function both during and after the recession.[8]

The channel I highlight is also distinct from those in other models of financial opacity, such as Gorton and Ordoñez (2012), Dang, Gorton, and Holmström (2013), and Hanson and Sunderam (2013). These studies tend to focus on the time-inconsistency in the design of

[8] According to sifma statistics, for example, US Corporate Bond and ABS issuance, for instance, actually climbed in 2009.

informationally-insensitive securities that are deployed as collateral in lending agreements. Through a similar mechanism, Moreira and Savov (2013) attempt to explain the slow US recovery in the context of neglected risk and the fragility of the shadow banking system. A similar literature, which includes Kobayashi and Nutahara (2007), Kobayashi, Nakajima, and Inaba (2012), and Gunn and Johri (2013), explores the impact of news shocks on business cycles in the presence of financial market imperfections, such as collateral constraints or costly state verification.

My work is related to the literature on dynamic models of asymmetric information, such as Foster and Viswanathan (1996), He and Wang (1995), and Allen, Morris and Shin (2006), which do not have real sectors and feature static economic environments where the asset's fundamental is fixed. Foster and Viswanathan (1996) models strategic, dynamic trading between investors with private information and a market maker in a static informational environment, while He and Wang (1995) examines the impact on trading volume when investors trade on public signals and dynamic private information in the presence of persistent noise supply shocks. Allen, Morris, and Shin (2006) and Bacchetta and van Wincoop (2006, 2008) investigate the role of higher-order expectations introduced by dispersed information in the determination of asset prices, and Nimark (2012) extends these implications to the term structure of interest rates. Albagi, Hellwig, and Tsyvinski (2013) rationalizes the credit spread puzzle with dynamic dispersed information and the nonlinear payoff profile of debt, and neither has a real sector nor long-lived incomplete information about the firm's fundamentals. My study focuses on the impact on asset prices and real activity when agents learn not only from endogenous information in prices generated by dispersed information, but also from the endogenous information in the return process governing the asset's time-varying fundamentals. To my knowledge, my work is also one of the first studies to study the long-run implications of a dynamic model of asymmetric information.

While my work exploits the local linearity of continuous-time and a non-overlapping generational informational structure for investors to help maintain tractability, the literature has developed other settings of information aggregation that deliver tractable equilibria outside of the CARA-Normal paradigm. Albagi, Hellwig, and Tsyvinski (2012), for instance, construct an equilibrium with log-concavity and an unboundedness assumption on the distribution of private signals that delivers a sufficient statistic for the market price as the private signal of the marginal trader. Goldstein, Ozdenoren, and Yuan (2013) and Albagi, Hellwig, and Tsyvinski (2012, 2014) employ risk-neutral agents with normally-distributed asset fundamentals and position limits to deliver tractable nonlinear equilibria in a static setting. Other papers like Sockin and Xiong (2014a,b) develop analytic log-linear equilibria in a static setting by exploiting Cobb-Douglas utility with fundamentals that have log-normal distributions. Straub and Ulbricht (2013) makes use of a conjugate prior framework with one period-lived, risk-neutral agents to maintain tractability in learning in a dynamic setting.

My work also contributes to the literature on informational frictions and the macroeconomy, which include Greenwood and Jovanovic (1990), Woodford (2003), Van Nieuwerburgh and Veldkamp (2006), Lorenzoni (2009), Kurlat (2013), Angeletos and La'O (2013), Blanchard, L'Huillier, and Lorenzoni (2013), Straub and Ulbricht (2013), Hassan and Mertens (2014a,b), Fajgelbaum, Schaal, and Taschereau-Dumouchel (2014), and David, Hopenhayn, and Venkateswaran (2014).[9] Only Straub and Ulbricht (2013), Hassan and Mertens (2014a,b), and David, Hopenhayn, and Venkateswaran (2014) consider the real consequences of informational frictions with centralized asset market trading to aggregate information. Informational frictions are, however, static in Hassan and Mertens (2014a,b), because of the assumption of perfect consumption insurance across agents, and in David, Hopenhayn, and Venkateswaran (2014), who focus on resource misallocation across firms from imperfect information, because firms observe their fundamentals after revenue is realized each period.[10] Straub and Ulbricht

[9] There is also a large literature on quantifying the impact of news shocks, which stresses the informational asymmetry between private agents and the econometrician, as well as situations in which agents have incomplete information. For a survey of this literature, see Beaudry and Portier (2013).

[10] In my setting, firms face more severe information frictions than in Hassan and Mertens (2014) and David,

(2013) explore the feedback loop between learning and the collateral channel, which destroys information during busts when agents become financially constrained because of a decline in the value of collateral with an exogenous, but hidden fundamental.[11] My focus instead is on the adverse feedback between asset prices and real investment that arises through the persistent distortion of the beliefs that govern real investment. In contrast to models like Albagi (2010), Kurlat (2013), and Straub and Ulbricht (2013), my learning mechanism does not arise because of financial frictions, but only informational frictions, which implies, for instance, that relieving credit conditions for firms will do little in my setting to improve economic conditions.

Finally, my paper also relates to the growing literature on the real effects of asset prices, which includes Bray (1981), Subrahmanyam and Titman (2001), Albagi (2010), Tinn (2010), Goldstein, Ozdenoren, and Yuan (2011, 2013), Angeletos, Lorenzoni, and Pavan (2012), Ordoñez (2012), Albagi, Hellwig, and Tsyvinski (2014), Sockin and Xiong (2014), and Gao, Sockin, and Xiong (2014).[12] Goldstein, Ozdenoren, and Yuan (2013) explores the coordination motive among financial investors when stock prices inform real investment decisions, while Albagi (2010) examines the distortion to real investment that occurs when financial market participants face funding constraints. Angeletos, Lorenzoni, and Pavan (2012) investigates the distortion to real investment and financial prices in a sequential game when entrepreneurs make investment decisions before claims are sold to the market to rationalize the dot com bubble. Albagi, Hellwig, and Tsyvinski (2014) highlights the inefficiency that asymmetric information introduces into real investment when existing shareholders extract informational rent by making investment decisions before selling shares to imperfectly-informed capital markets. Tinn (2010) features a similar setup to Angeletos, Lorenzoni,

Hopenhayn, and Venkateswaran (2014) because they neither observe private signals nor the past history of the realized fundamental. As a result, learning occurs more slowly and uncertainty about the fundamental fluctuates endogenously over time.

[11] In a similar spirit, a working paper version of Kurlat (2013) illustrates how adverse selection in asset markets can lead to countercylical uncertainty when there is incomplete information.

[12] See Bond, Edmans, and Goldstein (2012) for a survey of this literature.

and Pavan (2012) of perfectly informed entrepreneurs selling to investors who observe a noisy public signal, where uncertainty is short-lived and again entrepreneurs have superior information to market participants. My dynamic model features feedback both from real investment to the beliefs and trading incentives of financial market participants, as in Tinn (2010), Angeletos, Lorenzoni, and Pavan (2012), Ordoñez (2012), and Albagi, Hellwig, and Tsyvinski (2014), and from financial markets back to real investment, as in Albagi (2010), Goldstein, Ozdenoren, and Yuan (2013), and Sockin and Xiong (2014) for firms, and Gao, Sockin and Xiong (2014) for home buyers. In contrast to these studies, my focus is on the dynamic consequences for real activity of learning from endogenous real and financial signals.

1.3. A Model of Informational Frictions

1.3.1. The Environment

I consider an infinite-horizon production economy in continuous-time on a probability triple $(\Omega, \mathbf{F}, \mathcal{P})$ equipped with a filtration $\mathcal{F}_t$. There are three fundamental shocks in the economy $\left\{Z_t^\theta, Z_t^\xi, Z_t^k\right\}$ which are standard independent Weiner processes. To focus on the impact of informational frictions in financial markets on real activity, I turn off the conventional channels for financial markets to feed back to real activity through financial frictions in borrowing and lending.

There are perfectly competitive, identical firms in the economy that manage capital K_t for households with which they produce output Y_t according to

$$Y_t = aK_t,$$

for $a > 0$. Firm managers are able to grow capital according to

$$\frac{dK_t}{K_t} = \left(I_t\theta_t - \delta\right) dt + \sigma_k dZ_t^k, \tag{1}$$

where I_t is investment per unit of assets, θ_t is the productivity of real investment in installing new capital, similar to the investment-specific technology shock of Greenwood, Hercowitz, and Krusell (1997, 2000), δ is depreciation, and Z_t^k is a Total Factor Productivity (TFP) shock to existing capital. Importantly, the productivity of real investment θ_t is unobservable to firm managers and all other economic agents in the economy.[13] It evolves according to an Ornstein-Uhlenbeck process

$$d\theta_t = \lambda\left(\bar{\theta} - \theta_t\right)dt + \sigma_\theta dZ_t^\theta, \tag{2}$$

which has the known solution, found by applying Itô's Lemma to $e^{\lambda t}\theta_t$ and integrating from 0 to t,

$$\theta_t = \theta_0 e^{-\lambda t} + \bar{\theta}\left(1 - e^{-\lambda t}\right) + \int_0^t \sigma_\theta e^{\lambda(s-t)} dZ_s^\theta. \tag{3}$$

The OU process is the continuous-time analogue of an AR(1) process in discrete-time and has a mean-reverting drift and iid shocks.[14]

Households consume the output from firms and invest in two assets in the economy: claims to the cash flows of the assets of firms which have price q_t and in (locally) riskless debt, which is an inside asset, with instantaneous interest rate r_t. Importantly, both assets are traded in centralized asset markets, so that prices are observable to both households and firm managers when forming their expectations about θ_t.

1.3.2. Households

There is a continuum $\mathcal{I} = [0, 1]$ of risk-averse households that are part of a non-overlapping generational structure with wealth $w_t(i)$ that invest in firm claims and riskless debt. Each

[13] Kogan and Papanikolaou (2013) consider a setting where agents are trying to learn about the growth opportunities of firms and know the investment-specific technology shock.

[14] Theoretically, it is possible for θ_t to take negative values, similar to dividends in Wang (1993) and Campbell and Kyle (1993), though one can choose parameter values so that this occurs with negligible probability. Since beliefs over θ_t must be absolutely continuous with respect to the true distribution, such restrictions would apply to the posterior for θ_t as well.

That θ_t can potentially be negative may reflect that the scale of a firm can be suboptimally large during economic contractions, and that firms would strongly benefit from consolidating their businesses and shedding assets.

household invests a fraction $x_t(i)$ of its wealth $w_t(i)$ in firm claims, which are perfectly divisible, and $1 - x_t(i)$ in riskless debt. I index time for households as $t, t+\Delta t, t+2\Delta t$ and consider the continuous-time limit when Δt is of the order dt. Households have log utility over flow consumption $\log c_t(i)$ and subjective discount rate ρ over the bequest utility $v_{t+\Delta t}(i)$ they leave to future generations. I work with bequest utility instead of a preference over final wealth, as in He and Krishnamurthy (2012), to derive several asset pricing relationships relevant to the problem of firms. All prices, however, are ultimately pinned down by market clearing and not these relationships. Since households have log utility, and are therefore myopic, their optimal policies for consumption and investment, as well as the pricing kernel implied by their marginal utilities, will be the same regardless of whether they are part of a non-overlapping generational structure or long-lived.[15]

Households are subject to a random, private preference shock at each instant, which represents a liquidity shock and is the outcome of a Poisson random variable $N_t(i)$ with intensity $\pi \in (0,1)$, where $l_t(i) = \Delta N_t(i)$ is an indicator variable that the household has been hit. If hit by the preference shock, a household must take a fixed position in asset markets by divesting a fraction ξ_t of its wealth invested in firm claims and moving it into riskless bonds. Only those households hit by the shock observe its size ξ_t. The size of the shock may be correlated with investment productivity θ_t, and follows the law of motion

$$d\xi_t = \alpha\sigma_\xi dZ_t^\theta + \sqrt{1-\alpha^2}\sigma_\xi dZ_t^\xi,$$

where $\alpha \in (-1,1)$ represents this correlation. The innovation Z_t^ξ represents the pure liquidity shock to ξ_t. Later, when I consider the impact of financial crises in my economy, a financial crisis will be a large positive realization of this common liquidity shock. This allows me to focus on the informational effect of one feature of financial crises: asset firesales that

[15] From Gennotte (1986), general homothetic preferences with incomplete information introduce a negative dynamic hedging term in addition to agents' myopic demand. Brown and Jennings (1989) provides a numerical analysis of the impact on investor trading that this additional hedging term introduces with dispersed information.

depress financial prices. Other important features of financial crises, such as credit rationing and balance sheet impairment, would exacerbate the impact of financial crises through my channel.

Households are part of a continuum, and therefore exactly a fraction π will receive the liquidity shock at time t. Since those hit by the shock take a fixed position in asset markets, they do not trade on their superior information about its magnitude. Furthermore, because households are atomistic and, as such, do not view their preference shock as having any impact on the aggregate dynamics of market prices, those hit by the shock do not have an incentive to sell the private information of its magnitude to other households.

An unrealistic feature of the liquidity shock ξ_t is that it is not bounded between zero and one, and can also be negative. This implies that a household hit by the preference shock may be induced to take a positive position in the risky asset or a levered short position. Since ξ_t represents the noise in financial market prices that prevents them from being fully revealing about investment productivity θ_t, it is necessary that ξ_t have Gaussian innovations for tractability in learning, and therefore it cannot be restricted to the interval $[0, 1]$. Given that the prices and investment will not depend on the wealth distribution of households in equilibrium, the redistributional consequences of the liquidity shocks are not significant for my results. In all discussions of welfare, I focus on the redistributional consequences of informational frictions by comparing welfare in my economy to one in which households and firms have perfect information.

Households in my economy have private information about its unobserved strength θ_t. At each date t, household i receives news about θ_t through a private signal $s_t(i)$

$$s_t(i) = \theta_t + \sigma_s Z_t^s(i),$$

where $Z_t^s(i)$ is a standard $N(0, 1)$ random variable that represents household $i's$ idiosyncratic signal noise that is independent across (i, t) and independent from Z_t^θ and Z_t^ξ

$\forall \ (t,i)$.[16] Households are part of a continuum and, as such, there is no aggregate risk from their idiosyncratic signal noise in the sense that the sum of the noise converges to zero in the $\mathcal{L}^2 - norm$.[17] Households at $t=0$ have a common Gaussian prior $\theta_0 \sim N\left(\hat{\theta}_0, \Sigma_0\right)$.

To simplify my analysis, and to focus on the feedback between the real sector and financial markets from learning, I assume that, while parents in a generation pass along their wealth to their children within a household, they do not pass along their private information, which includes their own private signal, the size of the liquidity shock if they were hit by it, and their initial wealth. As discussed in the introduction, models of information aggregation even in static settings are very difficult to solve, and I make this common, simplifying assumption so that learning by households and firm managers remains tractable. This lets me avoid both the infinite regress problem of Townsend (1983), where market prices partially reveal a moving-average representation of the investment productivity θ_t, and a time-varying correlation between the persistence of wealth of households and the persistence of their private beliefs.[18] In addition to making learning intractable, it would also render the equilibrium no longer Markovian.

This assumption about the information structure, however, is not material for the main qualitative insights of my analysis. Relaxing it would introduce an additional component to the riskless rate that reflects that optimistic households tend to be wealthier during booms and poorer during recessions, similar to Detemple and Murphy (1994), Xiong and Yan (2010), and Cao (2011) for heterogeneous beliefs. This effect, however, is not likely to be significant given the nature of the equilibrium. The low uncertainty at business cycle peaks mitigates wealth inequality at peaks and during busts because households hold similar beliefs about investment productivity. This dampens the increased interest rate volatility

[16] One can model this Gaussian process, for instance, as a time-change Wiener process.

[17] Since convergence of stochastic objects in continuous-time is in the $\mathcal{L}^2 - norm$, there is little reason to think about convegence in an $a.s.$ sense. There do, however, exist Fubini extensions of the Lebesgue measure for the index of agents such that the convergence is $a.s.$ See, for instance, Sun and Zhang (2009).

[18] Nimark (2012) instead takes the approach of having traders with long-lived private information but static wealth to break the time-varying correlation between trader wealth and private beliefs.

that the interaction between wealth and beliefs would normally introduce.

In addition to their private signal $s_t(i)$, all households in a generation observe the history of firm asset growth in the economy $\log K_t$, investment I_t, the price of firm claims q_t, and the riskless rate r_t.[19] While private information is known by an individual, and would have to be remembered and passed along to progeny, historical public information is kept in public records and is readily available. Let the common knowledge, or public, filtration $\mathcal{F}_t^c$ be the minimal sigma-algebra generated by these public signals.

Households form rational expectations about the underlying state θ_t by Bayes' Rule given their information set $\mathcal{F}_t^i = \mathcal{F}_t^c \vee \{w_t(i), s_t(i)\}$, which is the public filtration $\mathcal{F}_t^c$ augmented with the household's current private wealth and signal. One can interpret the information structure of my economy as all households entering the current period with a common, time-varying prior based on the full history of public information $\mathcal{F}_t^c$, and then each updates its prior based on its private signal $s_t(i)$. Define $\hat{\theta}_t(i) = E[\theta_t \mid \mathcal{F}_t^i]$ to be the conditional expectation of θ_t of household i, where $E[\cdot \mid \mathcal{F}_t^i]$ is the conditional expectations operator with respect to the information set $\mathcal{F}_t^i$.

Households in each generation choose their consumption and investment to maximize their utility and their utility bequest to future generations $v_{t+\Delta t}(i)$, according to

$$0 = \sup_{c_t(i), x_t(i)} \left\{ \rho \Delta t \log c_t(i) + (1 - \rho \Delta t) E\left[v_{t+\Delta t}(i) \mid \mathcal{F}_t^i \right] - v_t(i) \right\}, \tag{4}$$

subject to the law of motion of their wealth $w_t(i)$ derived below. All households have the same initial wealth w_0. The optimization problem is solved under household $i's$ filtration $\mathcal{F}_t^i$ which incorporates household $i's$ private beliefs about investment productivity θ_t.

1.3.3. Firms

I keep the model of firms as simple as possible. There is a continuum of perfectly

[19] Since output is related to asset growth by $y_t = aK_t$, observing asset growth is the same as observing output.

competitive, identical firms in the economy who issue claims to households. Firms issue equity claims to households and are run by managers who have two responsibilities: to oversee the firm's operations and to invest $I_t K_t$ to grow the firm's assets K_t according to equation (1). Firms must maintain a minimal level of investment $\underline{I}$ such that $I_t \geq \underline{I}$. This prevents the signals about investment productivity θ_t in the economy from fully flattening, since, if $I = 0$, then neither firms nor households care about the productivity of investment.

The choice of functional form for the capital accumulation equation (1) makes transparent the impact of firm beliefs on real investment and uncertainty in the economy, as well as shuts down any variation in the second moment of firm capital accumulation because of investment to turn off the real-options "wait-and-see" channel of Bernanke (1983) featured in Bloom (2009). While this law of motion will mechanically give rise to a stark relationship between asset growth and the signal strength of real investment, similar to the choice of the production function of firms in Van Nieuwerburgh and Veldkamp (2006), as well as between investment and the Sharpe ratio of the return on firm claims, the interaction between investment and the level of uncertainty in determining the behavior of the market price, which is the focus of my analysis, will be an equilibrium outcome. The insights about the relationships explored here will hold more generally as long as firms care about the current, hidden state of the economy when they invest, and that there is more information from real signals when real activity is high.[20]

Firm managers invest I_t to maximize the value to shareholders of its claims. Firms face frictions in adjusting their level of investment, and can only imperfectly control it by choosing effort g_t so that I_t evolves according to

$$dI_t = g_t I_t dt,$$

[20] One may notice that learning from capital accumulation would be strong during recessions as well as expansions if real investment became largely negative, and firms, on aggregate, rapidly disinvested. Since aggregate US private nonresidential fixed investment historically has been nonnegative, I abstract from this artifact of the specification of the capital accumulation process.

with $g_t \geq 0$ if $I_t = \underline{I}$. Thus $\underline{I}$ is a reflecting boundary for I_t. Managers incur a linear cost $\frac{1}{\rho} g_t$ for this adjustment per unit of current investment $I_t K_t$, which is rebated back to the firm as a subsidy τ_t. The cost is meant to slow the adjustment of real investment and captures that real investment, in practice, is sluggish. As will be shown, if firms could choose the level of investment I_t directly, then I_t would have a well-defined solution between $\underline{I}$ and a because the value of their claims q_t is pinned down by household risk aversion, and is decreasing in I_t. With this slow adjustment, firms will have this same optimal I_t that they are slowly trying to adjust to by varying g_t, and therefore the policies with and without the technical restrictions are qualitatively similar. Since my comparisons for the dynamics of the economy will be relative to a perfect-information benchmark, relative business cycle asymmetries will not be driven by this assumption.

Households that hold firm claims receive a payment D_t of the residual cash flow from operations and investment

$$D_t = \left(a - I_t - \frac{1}{\rho} g_t I_t + \tau_t \right) K_t,$$

Firms finance their investment $I_t K_t$ from their cash flow from operations, the shortfall of which is made up by households through the sale of additional claims. Since financial markets are frictionless, they do not need to hold cash reserves.

For simplicity, managers do not have access to the private information of households and choose investment using only public information. While, in reality, firms are likely to have private information about the idiosyncratic component of their businesses or industries, they still have imperfect knowledge of general macroeconomic trends.[21] For managers to have access only to public information, they cannot observe the pricing kernels of their investors or their investors' ownership stakes in the firm. If they did, then managers would

[21] My mechanism is robust to managers having private information as long as they do not have superior information to households, in which case they would not need to learn from prices. See, for instance, David, Hopenhayn, and Venkateswaran (2014) for a setting in which firms also observe noisy private signals about their fundamentals.

know the identity of the marginal buyer of its firm's claims, which would allow it to infer information about investment productivity θ_t. Given that managers make use of only public information, their investment strategies must be measurable with respect to the common knowledge filtration $\mathcal{F}_t^c$.

I assume that firm managers attempt to maximize shareholder value for their investors who are not hit by the preference shock ξ_t. The logic behind this choice is that households who trade because of the preference shock are trading for reasons unrelated to the return on firm claims, reasons for which they are happy to take whatever position the shock demands regardless of managers' investment policies, and therefore it is unclear that maximizing shareholder value is the appropriate objective for them. Though managers must choose their investment policies from "behind the veil", since they do not know the composition of their shareholders, their policies in equilibrium will be robust to this uncertainty.

Let Λ_t be the pricing kernel of its shareholders not hit by the preference shock and E_t the value of firm claims. Firm managers then solve the optimization problem

$$E_0 = \sup_{\{g_t\}_{s \geq 0}} E\left[\int_0^\infty \frac{\Lambda_s}{\Lambda_0} D_s ds \mid \mathcal{F}_0^c\right], \tag{5}$$

subject to the transversality condition

$$\lim_{T \to \infty} E\left[\Lambda_T E_T \mid \mathcal{F}_0^c\right] = 0.$$

Since firms are perfectly competitive and atomistic, they take the pricing kernel of their shareholders as given. Though I restrict my attention to firm equity claims, it is worth mentioning that, since households have superior information compared to firms about general macroeconomic trends, firms could find it optimal to issue additional securities in this economy in order to have more signals from which to learn about the underlying state θ_t. Such a richer setting would introduce additional complexity, since instruments like risky debt are likely to have nonlinear payoffs, without adding much additional insight.

1.3.4. Market Clearing

Household i takes the net position firm claims $x_t(i)\, w_t(i) - q_t k_t(i)$, where $q_t k_t(i)$ is its initial holdings. Aggregating over all households then imposes the market clearing condition for the market for firm claims

$$\int_0^1 \left(x_t(i)\, w_t(i) - q_t k_t(i)\right) di = \int_0^1 x_t(i)\, w_t(i)\, di - q_t K_t = 0,$$

where $K_t = \int_0^1 k_t(i)\, di$ is the total assets of the firm at time t. Market clearing in the market for riskless debt additionally imposes that

$$\int_0^1 \left(1 - x_t(i)\right) w_t(i)\, di = 0.$$

Figure 1 in the Appendix illustrates the structure of the model. I search for a recursive competitive noisy rational expectations equilibrium.

1.3.5. Recursive Competitive Noisy Rational Expectations Equilibrium

Let ω be a state vector of publicly observable objects. A recursive competitive equilibrium for the economy is a list of policy functions $c\left(w(i), \hat{\theta}(i), \omega\right)$, $x\left(w(i), \hat{\theta}(i), \omega\right)$, $y(j, \omega)$, and $i(\omega)$, value functions $v\left(w(i), \hat{\theta}(i), \omega\right)$ and $E(\omega)$, and a list of prices $\{q(\omega), r(\omega)\}$ with $q(\omega) \geq 0$ such that

- Household Optimization: For every ω and i, given prices $\{q(\omega), r(\omega)\}$, $c\left(w(i), \hat{\theta}(i), l(i), \omega\right)$, and $x\left(w(i), \hat{\theta}(i), l(i), \omega\right)$ solve each household's problem (4) and deliver value $v\left(w(i), \hat{\theta}(i), \omega\right)$

- Firm Manager Optimization: For every ω, given prices $\{q(\omega), r(\omega)\}$, $g(\omega)$ solves the firm manager's problem (5) and delivers value $E(\omega)$

- Market Clearing: The markets for output, firm claims, and riskless debt clear

$$
\begin{aligned}
&: \int_0^1 c\left(w(i), \hat{\theta}(i), l(i), \omega\right) di + I(\omega) K = aK && (output) \quad (6)\\
&: \int_0^1 x\left(w(i), \hat{\theta}(i), l(i), \omega\right) w(i)\, di = qK && (firm\ claims\ market) \quad (7)\\
&: \int_0^1 \left(1 - x\left(w(i), \hat{\theta}(i), l(i), \omega\right)\right) w(i)\, di = 0 && (riskless\ debt\ market), \quad (8)
\end{aligned}
$$

- Consistency: $w(i)$ follows its law of motion $\forall\, i \in [0,1]$, household i forms its expectation about θ based on its information set $\mathcal{F}^i$ and firm managers form their expectation about θ based on their information set $\mathcal{F}^c$ according to Bayes' Rule

and the transversality conditions are satisfied.

1.4. The Equilibrium

I first state the main proposition of the section and then build up to this proposition in a sequence of key steps.

PROPOSITION 1.1: There exists a (locally) linear noisy rational expectations equilibrium in which the riskless return r is given by

$$
r = \frac{a}{a-I}\rho - \delta - \frac{\sigma_k^2}{1-\pi} + I\frac{\Sigma}{\Sigma + \sigma_s^2}\left(\theta - \hat{\theta}^c\right) - \frac{\pi\sigma_k^2}{1-\pi}\xi,
$$

when $I > \underline{I}$, and each household's investment in firm equity $x(i)$ can be decomposed into

$$
x(i) = x_c + x_i\left(\hat{\theta}(i) - \hat{\theta}^c\right),
$$

where

$$
\begin{aligned}
x_c &= \frac{\frac{a}{a-I}\rho - r - \delta}{\sigma_k^2},\\
x_i &= \frac{I}{\sigma_k^2}.
\end{aligned}
$$

When $I = \underline{I}$ and $g = 0$, then r is instead given by

$$r = \rho - \delta - \frac{\sigma_k^2}{1-\pi} + I\hat{\theta}^c + I\frac{\Sigma}{\Sigma + \sigma_s^2}\left(\theta - \hat{\theta}^c\right) - \frac{\pi\sigma_k^2}{1-\pi}\xi,$$

and x_c is given by

$$x_c = \frac{\rho + I\hat{\theta}^c - r - \delta}{\sigma_k^2}.$$

Similar to He and Wang (1995), individual households take a position in firm claims that can be decomposed into a component common to all households $x_c(\omega)$ and a term that reflects their informational advantage based on their private information $x_i(\omega)\left(\hat{\theta}(i) - \hat{\theta}^c\right)$. This informational advantage term reflects disagreement among households about the Sharpe Ratio of investing in firm claims. In contrast to He and Wang (1995), and other models of dispersed information like Foster and Viswanathan (1996) and Allen, Morris, and Shin (2006), the intensity with which households trade on their private information is influenced by real factors in the economy. Though private information is static, since the private information of households is short-lived because of the generational structure and because the signal-to-noise ratio of the private signals $s_t(i)$ is constant, the intensity with which households trade on their private information is now dynamic because the environment in which they trade is time-varying.

As is common in general equilibrium models of production, such as Cox, Ingersoll, and Ross (1985), interest rates adjust until all wealth is invested in firm assets. Focusing on the interaction between financial markets and real investment necessitates the adoption of such a setting that has this feature. In models of heterogeneous beliefs, such as Detemple and Murphy (1994) and Xiong and Yan (2010), the riskless rate r, which is the price at which relative pessimists are willing to offer leverage to relative optimists to hold all firm claims in equilibrium, reflects the disagreement among households about investment productivity θ_t. In my setting, it serves to aggregate their private information. This riskless rate falls during recessions to raise the expected excess return to firm claims, and shift down the level of optimism of the marginal buyer so that enough households purchase claims for asset markets to clear. Similarly, it rises during booms to shift up the level of optimism of the marginal buyer to curb the high demand of households for claims because of limited supply.

The market clearing condition for riskless debt effectively pins down the risk premium on firm claims required for asset markets to clear. As such, one can view market risk premium, whether it be the equity premium or a credit spread, as being the relevant market rate that aggregates information. Alternatively, one could interpret the interest rate in my stylized setting as being a composite market rate that arises from the trading of a well-diversified portfolio of securities. In the empirical discussion, I focus on the excess return to firm claims, or the spread between the return to firm claims and this riskless interest rate, to try to avoid taking a stance on which market rates have informative content.

The first step toward solving the equilibrium is to solve for the consumption and portfolio choice of household i given its information set $\mathcal{F}^i$. In what follows, I anticipate that the price of firm claims q will be a continuous, nonnegative function of finite total variation with respect to the level of investment I. Since q will have zero continuous quadratic variation, one has by a trivial application of Itô's Lemma that $\frac{dq}{q} = \frac{\partial_I q}{q} dI$.

I now derive the law of motion of the wealth of household i $w(i)$. Applying Itô's Lemma to K, the wealth of household i $w(i)$ then evolves according to

$$dw(i) = (rw(i) - c(i))\,dt + x(i)\,w(i)\left(\frac{(a-I)\,Kdt + Kdq + qdK}{qK} - rdt\right),$$

which can be expanded to yield

$$dw(i) = (rw(i) - c(i))\,dt + x(i)\,w(i)\left(\frac{a-I}{q} - r\right)dt + x(i)\,w(i)\left(\frac{dq}{q} + \frac{dK}{K}\right), \tag{9}$$

and is irrespective of the measure. The variance term for $\frac{dK}{K}$ is irrespective of the measure because of diffusion invariance. Intuitively, it is easier to estimate variances than means of processes, so that even if two households disagreed on the drift of a process, they cannot disagree on its variance. The dividend $a - I$ reflects the dividend after the rebate for the adjustment cost.

To make progress in solving household $i's$ problem, I analyze each household's problem (4) in the limit as $\Delta t \searrow dt$. Since uncertainty over θ_t represents a compound lottery for households over the uncertainty in the change in θ_t, I can separate their filtering from their optimization problem and treat $\hat{\theta}_t(i) = E\left[\theta_t \mid \mathcal{F}_t^i\right]$ with variance $\Sigma_t(i) = E\left[\left(\theta_t - \hat{\theta}_t(i)\right) \mid \mathcal{F}_t^i\right]$ as θ_t in their optimization problem.

Given that households have log preferences over consumption, and that liquidity shocks are proportional to wealth, households will optimally consume a fixed fraction of their wealth at each date t. Furthermore, when they are unconstrained in investment, they will also choose a myopic portfolio in the sense that it maximizes the Sharpe Ratio of their investment and ignores market incompleteness. This is summarized in the following proposition.

PROPOSITION 1.2: The household's value function takes the form $v\left(w(i), \hat{\theta}(i), l(i), h\right) = \frac{1}{\rho}\log w(i) + f\left(\hat{\theta}(i), l(i), h\right)$, where h_t is a vector of general equilibrium objects. Furthermore, the household's optimal consumption and portfolio choice take the form

$$
\begin{aligned}
c(i) &= \rho w(i), \\
x(i) &= \begin{cases} \dfrac{\frac{a-I}{q} + \frac{\partial_I q}{q} Ig + I\hat{\theta}(i) - r - \delta}{\sigma_k^2} & l(i) = 0 \\ -\xi & l(i) = 1 \end{cases}.
\end{aligned}
$$

Furthermore, define $\Lambda_t(i) = e^{-\rho t}\frac{1}{w_t(i)}$ to be the pricing kernel of household i that is not hit by a liquidity shock. Then the riskless rate and risky firm claims satisfy $\forall\ i$

$$
\begin{aligned}
r &= -\frac{1}{dt}E\left[\frac{d\Lambda(i)}{\Lambda(i)} \mid \mathcal{F}^i\right], \\
0 &= \frac{a-I}{q}dt + E\left[\frac{d\left(\Lambda(i)\, qK\right)}{\Lambda(i)\, qK} \mid \mathcal{F}^i\right].
\end{aligned}
$$

An immediate observation is that, similar to Detemple (1986), a separation principle applies in my noisy rational expectations equilibrium: the optimal consumption and investment policies are chosen independent of the learning process. Intuitively, since households are fully rational and update their beliefs with Bayesian learning, I can separate the filtering problem faced by households from their consumption choices and portfolio optimization.

Given the optimal choice of consumption $c(i) = \rho w(i)$ from the proposition, it follows that the law of motion of $w(i)$ can be written as

$$\frac{dw(i)}{w(i)} = (r - \rho)\, dt + x(i) \left(\frac{a - I}{q} dt + \frac{\partial_I q}{q} I g dt + \frac{dK}{K} - r dt \right), \tag{10}$$

which is also irrespective of the measure because of diffusion invariance.

From the market clearing conditions for the market for firm equity and riskless inside debt (7) and (8), the price of firm securities is given by

$$W = qK. \tag{11}$$

Equation (11) states that, in equilibrium, the total wealth in the economy W is equal to the total value of firm assets qK. Substituting $c(i) = \rho w(i)$ and equation (11) into the market clearing condition for output (6), it follows that

$$q = \frac{a - I}{\rho}, \tag{12}$$

from which follows that $\frac{a-I}{q} = \rho$, and the household, in equilibrium, receives a constant dividend yield from firm claims.

I now derive the conditional beliefs of households and firms about θ_t with respect to the common knowledge filtration $\mathcal{F}^c$ and their private information sets $\mathcal{F}^i$. The public signals that households have available for forming their expectations are $\log K$, q, I, and r. Since firm managers only have access to public information, it must be the case that firm investment $I \in \mathcal{F}^c$. Consequently, there is no additional information contained in I, or q given equation (12), once households have formed their beliefs. I can then generate the public filtration $\mathcal{F}^c$ with these two public signals $\mathcal{F}^c = \sigma\left(\{\log K_u, r_u\}_{u \leq t}\right)$.

Given the results of the main proposition, Proposition 1.1, let me now conjecture that the riskless rate r takes the form

$$r = r_0 + r_\theta(I, \Sigma)\left(\theta - \hat{\theta}^c\right) + r_\xi \xi, \tag{13}$$

where $r_\theta(I,\Sigma) \in \mathcal{F}^c$ since $(I,\Sigma) \in \mathcal{F}^c$. I assume that $|r_\xi|^{-1} > 0$ and that $r_\theta(I,\Sigma)$ is uniformly bounded and nonvanishing $a.s.$ Given equation (13), one can construct the public signal S

$$S = \frac{r - r_0 + r_\theta(I,\Sigma)\hat{\theta}^c}{r_\xi} = R_\theta(I,\Sigma)\theta + \xi. \tag{14}$$

Comparing equation (14) with the expression for the riskless rate r in Proposition 1.1, it follows that $R_\theta = -\frac{1-\pi}{\pi}\frac{I}{\sigma_k^2}\frac{\Sigma}{\Sigma+\sigma_s^2}$. Assuming that R_θ is a process of finite total variation, applying Itô's Lemma to S, S follows the law of motion

$$dS = \left(\partial_\Sigma R_\theta \frac{d\Sigma}{dt} + \partial_I R_\theta I g\right)\theta dt + R_\theta\lambda\left(\bar{\theta} - \theta\right)dt + \left(R_\theta\sigma_\theta + \alpha\sigma_\xi\right)dZ^\theta + \sqrt{1-\alpha^2}\sigma_\xi dZ^\xi.$$

Given these arguments, I can construct the vector of public signals $\zeta = \left[\begin{array}{cc}\log K & S\end{array}\right]'$ whose history, along with initial household wealth w_0 and firm assets K_0, generate the information set $\mathcal{F}^c$. Assuming that households using only the history of the public signals have a normal prior about θ_t, then after observing the two conditionally normal signals ζ_t their optimal updating rule for their beliefs about θ_t is linear, and their posterior belief about θ_t will also be conditionally normal. In continuous-time, these updating rules characterize the laws of motion for the conditional expectation and variance of these beliefs, $\hat{\theta}^c = E\left[\theta \mid \mathcal{F}^c\right]$ and $\Sigma = E\left[\left(\theta - \hat{\theta}^c\right)^2 \mid \mathcal{F}^c\right]$, respectively. In addition, ζ contains $\hat{\theta}^c$, Σ, and the level of investment I, which are all publicly observable, though we supress these arguments from the vector for simplicity since they do not contain new information about θ_t. Households then update these public estimates with their normally distributed private signals following another linear updating rule, and I have the following result.

PROPOSITION 1.3: The conditional belief of households using only public information is Gaussian with conditional expectation $\hat{\theta}^c = E\left[\theta \mid \mathcal{F}^c\right]$ and conditional variance $\Sigma = E\left[\left(\theta - \hat{\theta}^c\right)^2 \mid \mathcal{F}^c\right] \in \left[0, \frac{\sigma_\theta^2}{2\lambda}\right]$ that follow the laws of motion

$$d\hat{\theta}^c = \lambda\left(\bar{\theta} - \hat{\theta}^c\right)dt + \sigma_{\hat{\theta}k}(I,\Sigma)\,d\tilde{Z}^k + \sigma_{\hat{\theta}r}\left(I,\hat{\theta}^c,\Sigma\right)d\tilde{Z}^r,$$

where

$$\sigma_{\hat{\theta}k}(I,\Sigma) = I\frac{\Sigma}{\sigma_k},$$

$$\sigma_{\hat{\theta}r}\left(I,\hat{\theta}^c,\Sigma\right) = \frac{R_\theta\sigma_\theta^2 + \alpha\sigma_\xi\sigma_\theta + R_\theta\left(\frac{\sigma_s^2}{\Sigma+\sigma_s^2}\frac{d\Sigma}{dt} + g\Sigma - \lambda\Sigma\right)}{\sqrt{\left(R_\theta\sigma_\theta + \alpha\sigma_\xi\right)^2 + \left(1-\alpha^2\right)\sigma_\xi^2}},$$

and

$$\frac{d\Sigma}{dt} = -\frac{B}{2A} \pm \frac{1}{2A}\sqrt{2B + 4A\left(\sigma_\theta^2 - 2\lambda\Sigma - I^2\frac{\Sigma^2}{\sigma_k^2}\right) - 1}$$

with

$$A = \frac{\left(R_\theta\frac{\sigma_s^2}{\Sigma+\sigma_s^2}\right)^2}{\left(R_\theta\sigma_\theta + \alpha\sigma_\xi\right)^2 + \left(1-\alpha^2\right)\sigma_\xi^2},$$

$$B = 1 + 2R_\theta\frac{\sigma_s^2}{\Sigma+\sigma_s^2}\frac{R_\theta\sigma_\theta^2 + \alpha\sigma_\xi\sigma_\theta + R_\theta\left(g-\lambda\right)\Sigma}{\left(R_\theta\sigma_\theta + \alpha\sigma_\xi\right)^2 + \left(1-\alpha^2\right)\sigma_\xi^2},$$

and

$$d\tilde{Z}^k = \frac{1}{\sigma_k}\left(d\log K + \left(\frac{1}{2}\sigma_k^2 + \delta - I\hat{\theta}^c\right)dt\right),$$

$$d\tilde{Z}^r = \frac{1}{\sqrt{\left(R_\theta\sigma_\theta + \alpha\sigma_\xi\right)^2 + \left(1-\alpha^2\right)\sigma_\xi^2}}\left(dS - R_\theta\left(\frac{\sigma_s^2}{\Sigma+\sigma_s^2}\frac{1}{\Sigma}\frac{d\Sigma}{dt} + g\right)\hat{\theta}^c dt - R_\theta\lambda\left(\bar{\theta} - \hat{\theta}^c\right)dt\right),$$

is a vector of standard Wiener processes with respect to $\mathcal{F}^c$.

The conditional expectation of θ_t of household i of generation t $\hat{\theta}(i) = E\left[\theta \mid \mathcal{F}^i\right]$ and the conditional variance $\Sigma(i) = E\left[\left(\theta - \hat{\theta}(i)\right)^2 \mid \mathcal{F}^i\right]$ are related to the average household estimates $\hat{\theta}^c$ and Σ by

$$\hat{\theta}(i) = \hat{\theta}^c + \frac{\Sigma}{\Sigma+\sigma_s^2}\left(s(i) - \hat{\theta}^c\right),$$

$$\Sigma(i) = \frac{\sigma_s^2}{\Sigma+\sigma_s^2}\Sigma.$$

The public or common knowledge belief $\hat{\theta}^c$ is derived from the endogenous public signals $\log K$ and r, while each household's private belief $\hat{\theta}(i)$ is a linear combination of this public belief and their private signal. This public belief $\hat{\theta}^c$ is an important state variable because it survives the aggregation of the beliefs of households, and because it is the forecast of firm managers. Similar to the Kalman Filter in discrete-time, the loadings on the normalized innovations $d\tilde{Z}^k$ and $d\tilde{Z}^r$ formed from the real investment and market signals, $\sigma_{\hat{\theta}k}$ and $\sigma_{\hat{\theta}r}$, respectively, represent the Kalman Gains of the public signals. Changes in the first moment of public beliefs $\hat{\theta}^c$ are a linear combination of a term capturing the deterministic mean-reversion of investment productivity, $\lambda\left(\bar{\theta}-\hat{\theta}^c\right)dt$, and a stochastic component related to the news from the innovations to the public signals, $\sigma_{\hat{\theta}k}\left(I,\Sigma\right)d\tilde{Z}^k+\sigma_{\hat{\theta}r}\left(I,\hat{\theta}^c,\Sigma\right)d\tilde{Z}^r$. The law of motion of the second moment of public beliefs Σ, in contrast, is (locally) deterministic and is the continuous-time analogue of the Ricatti equation for the Kalman filter, yet it is stochastic unconditionally.

An important feature of the optimal filter is that the conditional variance of public beliefs Σ is time-varying over the business cycle, and fluctuates endogenously according to its law of motion given in Proposition 1.3, which depends on its current value, the perceived investment productivity $\hat{\theta}^c$, and the level of investment by firms I. The stochastic time-variation in Σ is in contrast to dynamic models of asymmetric information like Wang (1993) that focus on the steady-state solution for the conditional variance of beliefs to which the economy tends deterministically. In this setting, Σ influences the quantity of private information households have, and how they trade on it in financial markets. As a result of shutting down the "wait and see" channel of Bloom (2009) for uncertainty to feed into firm investment behavior, firm investment decisions are indirectly influenced by Σ purely through how it affects the informativeness of the financial signal. Since Σ is time-varying, it is part of the state vector, along with I and $\hat{\theta}^c$, that summarizes the current state of the economy.

Learning from the endogenous market signal r that aggregates households' private information leads to either zero or two solutions for the (locally) deterministic change in the conditional variance $\frac{d\Sigma}{dt}$, which can result in nonexistence and multiple solutions.[22] With

[22] Nonexistence can occur because learning from market prices leads to the simultaneous determination of the change in the conditional variance $\frac{d\Sigma}{dt}$ and the strength of the market signal $\sigma_{\hat{\theta}r}$. There are situations

two solutions, households and firms can coordinate around either solution for the change in Σ, one which leads them to learn about investment productivity θ_t faster, and one in which they learn more slowly. In all the numerical applications, I follow the convention of selecting the larger root when two real solutions to $\frac{d\Sigma}{dt}$ exist, since the smaller, more negative root tends to lead households and firms to learn about θ_t extremely quickly.

In addition to their private signals, households learn about the underlying strength of the economy θ from the growth of firm assets $\log K$, whose informativeness (signal-to-noise ratio) is increasing in the level of firm investment I, and from the riskless rate, whose informativeness $R_\theta(I, \Sigma)$ is also influenced by I. This link from the investment choices of firms to the learning process of households represents one part of the feedback loop between real activity and asset markets that I wish to highlight. The ability of real investment decisions to distort investor expectations is similar to the channel explored in Angeletos, Lorenzoni, and Pavan (2012) to rationalize the tech bubble of the early 2000's.

I now turn to the problem faced by firm managers. Given that firm managers only have access to public information, their conditional expectation of θ when making their investment decision g is $\hat{\theta}^c$. Furthermore, since the price of firm claims is pinned down by market clearing $q = \frac{a-I}{\rho}$, it must be the case that the optimal choice of g under the pricing kernel of investors confirms this price.

PROPOSITION 1.4: The value of firm claims is given by $E = qK$, and the optimal level of investment is given by

$$g = \rho\left(q\hat{\theta}^c - 1\right)\mathbf{1}\left\{I > \underline{I} \,\cup\, \hat{\theta}^c \geq \frac{\rho}{a - \underline{I}}\right\}. \tag{15}$$

From the functional form of the optimal investment policy, it is apparent that $I = \underline{I}$ and

when the real signal and the natural mean-reversion of θ_t are so strong that the conditional variance Σ falls too precipitously, as measured by $\frac{d\Sigma}{dt}$, for $\sigma_{\hat{\theta}r}$, which depends on $\frac{d\Sigma}{dt}$, to be sufficient to justify the fall in Σ. This result is reminiscent of the finding of Futia (1981) that price formation in a linear rational expectations framework can exhibit nonexistence pathologies.

$I = a$ are reflecting boundaries, since when $I = a$, then $q = 0$ and $g < 0$. As a result, the price of firm claims can never be negative. Similarly, when $I = \underline{I}$ and $\hat{\theta}^c \leq \frac{\rho}{a - \underline{I}}$, then $dI = 0$ and investment stays at $\underline{I}$ until g becomes positive. Since I has finite variation, its sample paths are continuous in time, and I will approach its two boundaries continuously.

To see how investment in my setting compares to one in which I allow firms to freely choose I, it is easy to see that the FOCs for the firm's problem would then be

$$-1 + q\hat{\theta}^c \leq 0,$$

with equality when $\hat{\theta}^c \geq \frac{\rho}{a - \underline{I}}$ since $q = \frac{a - \underline{I}}{\rho}$, from which it follows that $I^{opt} = \underline{I} + \left(a - \frac{\rho}{\hat{\theta}^c} - \underline{I}\right)$ $\mathbf{1}\left\{\hat{\theta}^c \geq \frac{\rho}{a - \underline{I}}\right\}$. Since firms choose bang-bang policies, the price of capital q adjusts to make them indifferent to the optimal level of investment I^{opt}. Notice that when $I = a - \frac{\rho}{\hat{\theta}^c}$ when I can only be slowly adjusted, then $I = I^{opt}$ and $g = 0$. If I were above its optimal value $I > I^{opt}$, then $g > 0$, and similarly $g < 0$ when I is below its optimal value $I < I^{opt}$. Thus g tries to adjust I toward the optimal level the firm would choose if I could be chosen freely. This is the sense in which investment is sluggish.

Given the solution to the optimal investment strategy of firms, q has the interpretation of being Tobin's q. Investment by firms aims to equate the perceived productivity of real investment $\hat{\theta}^c$ to $1/q$, the book-to-market value of its assets. Thus informational frictions distort real investment by creating a misperception about the value of its assets. This highlights a key difference between my channel for firm beliefs to distort real activity and that of Straub and Ulbricht (2013). In their setting, entrepreneurs are never confused about the optimal level of production, but rather about the value of the collateral they must pledge to workers because of financial frictions. In my setting, firms optimally choose a level of production that is distorted because of their beliefs about investment productivity. Also, in contrast to models of uncertainty like Bloom (2009), investment in my economy declines because of shocks to the first moment of productivity rather than from shocks to the second moment through a "real-options" channel.

Learning by firm managers introduces a channel through which the first moment of beliefs about investment productivity $\hat{\theta}^c$ influences the second moment Σ. From Proposition 1.3, the change in uncertainty Σ $\frac{d\Sigma}{dt}$ depends on the change in investment g, which is a function of firm manager beliefs $\hat{\theta}^c$. Thus the filtering equations for $\hat{\theta}^c$ and Σ are coupled because there is feedback from second moments to first moments, which is a natural feature of the optimal nonlinear filter, and from first moments to second moments, because learning by firm managers determines their investment decisions, which influences the informativeness of the two public signals.

Since household trading behavior impacts the riskless rate r, from which both households and firms learn, the riskless rate acts as a channel for liquidity shocks in financial markets to feed into real investment decisions by influencing manager expectations. This mechanism for asset prices to distort firm investment is similar to Goldstein, Ordozen, and Yuan (2013). Along with the impact of investment decisions on household learning discussed above, these two forces characterize the feedback loop in learning between financial markets and real activity.

To derive the functional form for the riskless rate r, I must aggregate the wealth-weighted private expectations of all households, which will reveal the current true θ_t and the signal noise of households. Given that the private beliefs of each household are uncorrelated with their wealth share because households do not pass along their private information to later generations, the Law of Large Numbers will cause the aggregation of idiosyncratic signal noise to vanish. Let $\mathcal{D}_t$ be the set of households hit by the liquidity shock at time t. Let $W = \int_0^1 w(i)\, di$ be the total wealth of all households. Then I obtain the following result.

PROPOSITION 1.5: Aggregating the wealth-weighted deviation in the conditional expectation θ of household i $\hat{\theta}(i)$ from the common knowledge expectation $\hat{\theta}^c$ yields $a.s.$

$$\int_{\mathcal{D}^c} \frac{w(i)}{W} \left(\hat{\theta}(i) - \hat{\theta}^c\right) di - \int_{\mathcal{D}} \frac{w(i)}{W} \xi di = (1-\pi) \frac{\Sigma}{\Sigma + \sigma_s^2} \left(\theta - \hat{\theta}^c\right) - \pi \xi,$$

and the convergence $\forall\ t$ is in the $\mathcal{L}^2 - norm.$

By aggregating the beliefs of individual households, the riskless rate r will depend on $\hat{\theta}^c$ and θ through the productivity of investment θ revealed by the households' private signals. An important caveat to this result is that it relies on households being symmetrically informed. If, instead, households had different signal precisions $\sigma_s(i)$, then the wealth distribution of households would matter for prices.[23] Given this aggregation result, the noisy rational expectations equilibrium and the riskless rate r then satisfy the main theorem of the section. Thus it follows that the state vector ω for the economy is $\omega = \left[\theta, \xi, \hat{\theta}^c, I, \Sigma\right]$.

While the intensity with which households trade on their private information is procyclical, since $x_i(\omega)$ is monotonically increasing in the investment by firms I, the information content in the market price is monotonically increasing in uncertainty about θ, measured by Σ, because market prices aggregate the private information of households to partially reveal θ. These two forces interact so that asset prices will be strongest during downturns and recoveries, in the sense that the variation in $\hat{\theta}^c$ driven by the market signal is largest when I and Σ are in an intermediate range. To see this, Figure 2 in the Appendix plots, as a numerical example, the loading of the market signal $\sigma_{\hat{\theta}r}\left(I, \hat{\theta}^c, \Sigma\right)$ on beliefs for a fixed level of perceived investment productivity $\hat{\theta}^c$ for a set of parameters listed in the Appendix. The figure reveals that the variation from the market signal is increasing in the level of investment by firms, and increasing in uncertainty about investment productivity Σ, though for other parameter values it can be non-monotonic. Furthermore, since a decline in the perceived investment productivity $\hat{\theta}^c$ lowers investment, and also leads to greater uncertainty, it follows that $\sigma_{\hat{\theta}r}\left(I, \hat{\theta}^c, \Sigma\right)$ can be increasing or decreasing in $\hat{\theta}^c$ depending on I and Σ. These observations illustrate that more of the variation in the beliefs in households and firm managers is driven by the market signal when I and Σ are in an intermediate range. As

[23] Asymptotically, however, one would expect households with superior information to eventually drive out the less well-informed households. This would lead to a degenerate wealth distribution in which wealth once again does not matter.

$\Sigma \to 0$, the market price contains little information about θ at peaks, since $R_\theta \to 0$, and households do not react strongly to it.

This last point merits some emphasis. While it is well-appreciated that risk premia in financial markets are countercyclical, it is less appreciated that the strength of asset prices as a signal of economic strength also exhibits business cycle asymmetries. This asymmetry arises because the incentives for investors to trade on their private information anchors on both the level of real investment and uncertainty in the economy.

1.5. The Impact of Feedback in Learning

To assess the impact of feedback in learning, I first derive the equilibrium in two benchmark economies, one with perfect information and one in which only households have perfect information, as helpful anchors for my analysis. The first benchmark gives us insight into how the economy behaves in the absence of any informational frictions, while the second will help to clarify the role that dispersed information among households plays in influencing the business cycle behavior of the market signal. I then explain the slow US recovery in the context of this feedback loop.

1.5.1. Two Benchmarks

Suppose that $\theta(t)$ is observable to all households and firm managers. Then all households will allocate identical fractions of their portfolios to risky projects and the riskless asset. In this benchmark setting, it is sufficient to solve the equilibrium for the aggregate state variables, since the wealth of households will only differ in their history of preference shocks. The following proposition summarizes the recursive competitive equilibrium that the recursive noisy rational expectations equilibrium tends to, in the aggregate, as informational frictions vanish for all agents.

PROPOSITION 1.6: When θ is observable to all households and firm managers, a) the price of firm equity is given by

$$q = \frac{a - I}{\rho},$$

b) the riskless return r satisfies

$$r = \frac{a}{a - I}\rho - \delta - \frac{\sigma_k^2}{1 - \pi} - \frac{\pi\sigma_k^2}{1 - \pi}\xi,$$

when $I > \underline{I}$, c) optimal consumption and investment in firm equity by households who are not hit by the liquidity shock satisfy

$$\begin{aligned} c(i) &= \rho w(i), \\ x(i) &= \frac{\frac{a-I}{q} - \frac{I}{a-I}g + I\theta - r - \delta}{\sigma_k^2}, \end{aligned}$$

and d) optimal investment by managers is given by

$$g = \rho(q\theta - 1)\,\mathbf{1}\left\{I > \underline{I} \,\cup\, \theta \geq \frac{\rho}{a - \underline{I}}\right\}.$$

The equilibrium with perfect information appears similar to the one with informational frictions, except that the riskless rate no longer reflects the wedge between the beliefs of agents and the true underlying strength of the economy θ because households and firm managers are now perfectly informed. The economy is isomorphic to one with a representative agent household who owns and manages all assets in the economy, and chooses the riskless rate so that it invests all its resources in assets given its preference shock. In this setting, there is no role for noise from preference shocks ξ to transmit to real investment decisions because manager do not learn from prices. Financial market activity has no consequence for the business cycle at all.

The second benchmark provides an intermediate case between the informational frictions economy of the previous section and the perfect-information benchmark. Though households behave identically when they have perfect information, there is still feedback from financial market noise ξ to real investment decisions because managers still must learn about θ from market prices. The behavior of this economy is summarized in the next proposition.

PROPOSITION 1.7: When θ is observable to all households, a) the price of firm equity is given by

$$q = \frac{a - I}{\rho},$$

b) the riskless return r satisfies

$$r = \frac{a}{a - I}\rho - \delta - \frac{\sigma_k^2}{1 - \pi} + I\left(\theta - \hat{\theta}^c\right) - \frac{\pi\sigma_k^2}{1 - \pi}\xi,$$

c) optimal consumption and investment in firm equity by households who are not hit by the liquidity shock satisfy

$$\begin{aligned} c(i) &= \rho w(i), \\ x(i) &= \frac{\frac{a-I}{q} - \frac{I}{a-I}g + I\theta - r - \delta}{\sigma_k^2}, \end{aligned}$$

and d) optimal investment by managers is given by

$$g = \rho\left(q\hat{\theta}^c - 1\right)\mathbf{1}\left\{I > \underline{I} \;\cup\; \hat{\theta}^c \geq \frac{\rho}{a - \underline{I}}\right\}.$$

Furthermore, beliefs, prices, and optimal policies in the economy with informational frictions approach their representative agent benchmark values as $\sigma_s \searrow 0$.

In this intermediate case, firm managers must still learn from both the growth of firm assets and market prices. Noise from market prices from preference shocks ξ can potentially feed back into firm manager learning, and therefore their investment decisions, yet there is an important distinction from the NREE equilibrium. Since households have perfect information, the level of uncertainty in the economy Σ does not affect their trading behavior, and consequently it has a smaller role in determining the influence and strength of the market signal. This can be seen from the difference in the loadings on the tracking error $\theta - \hat{\theta}^c$ in the expressions for r in Propositions 1.6 and 1.7. In the NREE economy, the signal-to-noise ratio $R_\theta = -\frac{1-\pi}{\pi}\frac{I}{\sigma_k^2}\frac{\Sigma}{\Sigma+\sigma_s^2}$, while in this representative agent setting $R_\theta = -\frac{1-\pi}{\pi}\frac{I_t}{\sigma_k^2}$. This implies that the market signal in the representative agent setting mimics much of the cyclical behavior of the real investment signal (though it is not redundant because the noise in the two signals are conditionally independent of each other). The market signal $S = \frac{1-\pi}{\pi\sigma_k^2}\left(r - \frac{a}{a-I}\rho + \frac{\sigma_k^2}{1-\pi} + I\hat{\theta}^c\right)$ for households and firm managers then has the law of motion

$$dS = R_\theta\left(g-\lambda\right)\theta dt + \left(R_\theta\sigma_\theta + \alpha\sigma_\xi\right)dZ^\theta + \sqrt{1-\alpha^2}\sigma_\xi dZ^\xi,$$

where R_θ is increasing in I and unrelated to uncertainty Σ. The market signal is, consequently, strongest during booms when uncertainty $\Sigma = E\left[\left(\theta - \hat{\theta}^c\right)^2 \mid \mathcal{F}^c\right]$ is low.

This setting consequently highlights the importance of dispersed information for the mechanism of the NREE economy: aggregation of dispersed information gives the market signal much of its countercyclical behavior because the quantity of private information Σ matters for how households trade on their private information. There is a dramatic difference, then, in the predictions of how an economy with a representative household behaves compared to an economy with households with heterogeneous information.

One could also consider a benchmark with a representative household that receives a noisy private signal instead of having perfect information. In this benchmark, the conditional variance of public beliefs Σ would be important for the information content of the market signal, and the market signal would exhibit more countercyclical behavior. Since the noise in the household's private signal would not vanish from market prices, however, it is less clear how the informativeness of the market signal would change over the cycle, since the noise in the price from the household's private signal would also increase as Σ increased.

1.5.2. Explaining the Slow US Recovery

My analysis highlights a potential channel by which recessions with financial origins can have deeper recessions and slower recoveries, and can help explain how the financial crisis of late 2008 may have contributed to the anemic US recovery. Economic agents rely more on price signals for helpful guidance about the state of the economy as the economy enters a downturn. Financial crises during downturns distort these price signals and, as a result of severe informational frictions, investors and firms interpret part of the collapse in asset prices as a signal of severe economic weakness. This further depresses real activity, causing both real and financial signals to flatten, which increases uncertainty and causes it to remain elevated. This makes it harder for private agents to act on signs of a recovery. Despite evidence of economic improvement, and a rebounding of financial markets, the heightened level of uncertainty makes it difficult for a recovery to gain traction and stifles growth.

To illustrate this story, Figure 3 depicts the impact of a unit dZ^{ξ} shock (a one standard deviation negative liquidity shock) to financial prices in the economy during a boom $\left(\hat{\theta}^{c}, I, \Sigma\right) = (.3, .1, 10^{-6})$ and during a bust $\left(\hat{\theta}^{c}, I, \Sigma\right) = (.2, .04, 10^{-5})$.[24] As a result of informational frictions, the recession is deeper in this numerical experiment compared to the perfect-information benchmark, and the recovery is also more gradual. In contrast, a one standard deviation negative financial shock during a boom has a much more attenuated impact on growth, which can help explain why financial events like the LTCM crisis had little effect on the real economy. Key to this result is that uncertainty is time-varying, with a law of motion given in Proposition 1.3, and countercyclical. When uncertainty is higher, noise in financial prices that is interpreted as bad news perpetuates low investment. This, in turn, perpetuates high uncertainty and allows the distortion to beliefs from the noise in financial prices to persistent.

My analysis consequently identifies a potential benefit of unconventional monetary policy in the presence of informational frictions. By buying treasury and mortgage-backed securities through Quantitative Easing (QE), the US government provided financing for investors to purchase assets from riskier asset classes, such as equities and speculative-grade debt. This

[24] Since time is continuous, we feed the quarterly negative shock to the model as one large innovation at time 0 equal to one fourth the annual variance of the financial shock.

injection of capital may have lessened the noise that constrained investors introduced into financial prices during the financial crisis that distorted the expectations of private agents about the strength of the US economy. In continuing QE in its various forms of QE1-QE3 until late 2014, however, the buoying of financial markets may have later added noise to financial prices that confused agents about the strength of the US recovery. The April 2011 WSJ article "Is the Market Overvalued?", for instance, discusses how market participants and economists like Robert Shiller could not disentangle signs of strong corporate profitability from the effects of QE behind the high valuations in the stock market.

1.6. Welfare

I now turn to the welfare implications of my analysis. The economy with informational frictions may be constrained inefficient because households and firms do not fully internalize the benefit of the public information they produce by trading in asset markets and engaging in real investment. As emphasized in Greenwald and Stiglitz (1986), economies with incomplete markets and incomplete information are generically not constrained Pareto efficient, and there is a role for welfare-improving policies. In this spirit, I consider several thought experiments that augment the provision of public information in the economy to highlight this potential externality.

I begin this section by characterizing ex-ante welfare in the economy. I adopt a utilitarian weighting scheme to aggregate utility across the heterogeneous households, normalizing welfare to initial household consumption to remove the level effect of initial conditions. This helps me construct a measure of welfare in the economy that has a stationary distribution conducive to conducting thought experiments. Since the noise in financial prices stems from the preference shocks of households, the analysis avoids the issue of characterizing welfare in the presence of exogenous "noise traders" discussed in Wang (1994). Informational frictions impact welfare through two channels: a distortion to real investment and household trading, and a cost that comes from the inequality in household wealth that arises because of the dispersion of private beliefs. This is summarized in the following proposition.

PROPOSITION 1.8: Ex-ante utilitarian welfare in the economy with informational frictions is given by

$$
\begin{aligned}
U \;=\; & \underbrace{\frac{1}{\rho}E\left[\int_0^\infty e^{-\rho t}\left(\frac{\rho}{a-I_t}+\theta_t-\hat{\theta}_t^c\right)I_t dt \mid \mathcal{F}_0\right]}_{Efficiency\ of\ Real\ Investment} \\
& \underbrace{-\frac{1-\pi}{2\rho}\left(\frac{\sigma_s}{\sigma_k}\right)^2 E\left[\int_0^\infty e^{-\rho t}\left(\frac{I_t\Sigma_t}{\Sigma_t+\sigma_s^2}\right)^2 dt \mid \mathcal{F}_0\right]}_{Cross-Sectional\ Inequality} \\
& -\frac{\delta}{\rho^2}-\frac{1}{2\rho^2}\left(\sigma_k^2+\frac{\pi\sigma_k^2}{1-\pi}\left(1+\xi_0\right)^2\right)-\frac{1}{2\rho^3}\frac{\pi\sigma_k^2}{1-\pi}\sigma_\xi^2.
\end{aligned}
$$

Under this welfare criterion, there exists a representative household in the economy who holds all claims to firm assets and whose wealth w evolves according to

$$
\frac{dw}{w}=\left(\frac{\rho I}{a-I}-\delta+I\left(\theta-\hat{\theta}^c\right)-\frac{1}{2}\left(\frac{\pi\sigma_k^2}{1-\pi}\left(1+\xi\right)^2+\left(1-\pi\right)\left(\frac{I}{\sigma_k}\frac{\Sigma}{\Sigma+\sigma_s^2}\sigma_s\right)^2\right)\right)dt+\sigma_k dZ^k.
$$

From Proposition 1.7, the representative household under this welfare criterion is different from a representative household who holds all firm claims since the criterion reflects the inequality in wealth that arises because of informational frictions and liquidity shocks. This distinction is absent from representative agent models and comes from the aggregation of flow utility $\log c(i)$ rather than consumption $c(i)$ in the utilitarian welfare function. The effects of the distortion show up as a tax on the representative household, and consequently one can think of the transfer of wealth from liquidity shocks and the presence of informational frictions as imposing a tax on the economy. This tax vanishes when households have identical beliefs, which occurs in the limiting cases when $\sigma_s \searrow 0$, $\sigma_s \nearrow \infty$, or $\Sigma \equiv 0$.

Having derived ex-ante utilitarian welfare to understand the forces that impinge on household utility, I construct a measure of expected welfare using only public information once the economy has reached its stationary distribution, and initial conditions no longer matter, as a sensible measure for conducting my thought experiments. To target household and firm investing behavior, I introduce a proportional position cost τ^r on household trading and a linear subsidy on firm real investment τ^I. I construct these instruments so that the extracted revenue is returned to households as lump-sum transfers that households view as being proportional to their wealth. The position cost lets me manipulate households' trading decisions while the real investment subsidy lets me manipulate firms' investment decisions.

Solving for household's optimal investment in the presence of the position cost, it is straightforward to see from Proposition 1.2 that household i invests a fraction $x\left(i\right)$

$$x\left(i\right)=\frac{\frac{a-I}{q}+\frac{\partial_I q}{q}Ig+I\hat{\theta}\left(i\right)-r-\delta}{\left(1-\tau^r\right)\sigma_k^2},$$

of its wealth in firm claims when not hit by the liquidity shock. Households that are hit by the preference shock continue to take a fixed position $-\xi$ proportional to their wealth in the risky asset, regardless of the position cost. Then, by similar arguments to those in Section IV, one can arrive at the form for the riskless rate r when investment is unconstrained

$$r=\frac{a}{a-I}\rho-\delta+I\frac{\Sigma}{\Sigma+\sigma_s^2}\left(\theta-\hat{\theta}^c\right)-\left(1-\tau^r\right)\frac{\sigma_k^2}{1-\pi}\left(1+\pi\xi\right),$$

from which follows that

$$x\left(i\right)=\frac{1}{1-\pi}+\frac{\pi}{1-\pi}\xi+\frac{1}{1-\tau^r}\frac{I}{\sigma_k^2}\frac{\Sigma}{\Sigma+\sigma_s^2}\sigma_s Z^s\left(i\right).$$

The position cost has the counterintuitive property that it induces households to take larger positions in the risky asset based on their private information. This happens because households in continuous-time can rebalance their portfolios instantaneously to take a large enough position to offset the impact of the cost. Since the collateral is returned lump-sum, however, the cost introduces a distortion to household wealth. A higher position cost τ^r increases the amount of public information in the price by inducing households to trade more on their private information without affecting the position taken by households hit by the liquidity shock, but it also introduces more wealth inequality. There is then a tradeoff for welfare in increasing τ^r.

It is also straightforward to see from Proposition 1.4 that the real investment subsidy induces the firm to choose a growth rate for real investment g

$$g = \left(\left(a - I\right)\hat{\theta}^{c} - \left(1 - \tau^{I}\right)\rho\right)\mathbf{1}_{\left\{I > \underline{I} \cup \hat{\theta}^{c} \geq \frac{\left(1-\tau^{I}\right)\rho}{a - \underline{I}}\right\}}.$$

With these instruments in place, I now search for the probability law of the economy once it has reached its stationary distribution $p\left(\hat{\theta}^{c}, \Sigma, I\right)$, if it exists. I derive the Kolmogorov Forward Equation (KFE), or transport equation, which summarizes the (instantaneous) transition of the probability law of the economy $p_t\left(\hat{\theta}^{c}, \Sigma, I\right)$ and characterize the conditions under which $\partial_t p_t\left(\hat{\theta}^{c}, \Sigma, I\right) = 0$. This reduces to solving the appropriate boundary value problem for a second-order elliptic partial differential equation, summarized in the following proposition.

PROPOSITION 1.9: The stationary distribution of the economy $p\left(\hat{\theta}^{c}, \Sigma, I\right)$ satisfies the Kolmogorov Forward Equation

$$\begin{aligned} 0 &= -\partial_{\hat{\theta}^{c}}\left\{p\lambda\left(\bar{\theta} - \hat{\theta}^{c}\right)\right\} - \partial_{I}\left\{pI\left(\left(a - I\right)\hat{\theta}^{c} - \left(1 - \tau^{I}\right)\rho\right)\right\}\mathbf{1}_{\left\{I > \underline{I} \cup \hat{\theta}^{c} \geq \frac{\left(1-\tau^{I}\right)\rho}{a - \underline{I}}\right\}} - \partial_{\Sigma}\left\{p\frac{d\Sigma}{dt}\right\} \\ &\quad + \frac{1}{2}\partial_{\hat{\theta}^{c}\hat{\theta}^{c}}\left\{p\left(\sigma_{\hat{\theta}k}^{2} + \sigma_{\hat{\theta}r}^{2}\right)\right\}, \end{aligned}$$

with boundary conditions given in the Appendix.

The KFE that defines the stationary distribution is a conservation of mass law that has an intuitive interpretation. It states that the sum of the flows of probability through a cube in the $\left(\hat{\theta}^{c}, \Sigma, I\right)$ space must be zero for the probability mass of the cube to be conserved over time. The stochastic component of $\hat{\theta}^{c}$ introduces a second-order term in the KFE related to its volatility since the high variability of Wiener processes has a first-order effect on the law of motion of $\hat{\theta}^{c}$.[25] In the case where $\sigma_s \nearrow \infty$ and $\alpha = 0$, the economy is analogous to that

[25] To find the stationary distribution numerically, I follow the trick of rewriting the KFE in Proposition 1.9 as $\mathcal{D}^{g*}p = 0$, where $\mathcal{D}^{g*}$ is the adjoint of the infinitesimal generator $\mathcal{D}^{g}$ defined in the proof of the proposition. Discretizing the state space $\left(\hat{\theta}^{c}, \Sigma, I\right)$ into a $N_{\hat{\theta}} \times N_{\Sigma} \times N_{I}$ grid, one can stack the $N_{\hat{\theta}} \cdot N_{\Sigma} \cdot N_{I}$

of Van Nieuwerburgh and Veldkamp (2006) in which only a real investment signal provides information.

Given the KFE, I now construct my welfare measure. Let U_p^c be utilitarian welfare in the economy, normalized to initial wealth, and $E^p[\cdot]$ be the expectation operator with respect to the stationary distribution. Then I have the following corollary.

COROLLARY 1: Expected utilitarian welfare under the stationary distribution U_p^c with position cost and real investment subsidy τ^r and τ^I, respectively, is given by

$$\begin{aligned} U_p^c &= \frac{1}{\rho} E^p\left[\frac{I_0}{a-I_0}\right] - \frac{1-\pi}{2\rho^2}\left(\frac{\sigma_s}{\sigma_k}\right)^2 E^p\left[\left(\frac{I_0}{1-\tau^r}\frac{\Sigma_0}{\Sigma_0+\sigma_s^2}\right)^2\right] \\ &\quad -\frac{1}{2\rho^2}\frac{1-\pi}{\pi\sigma_k^2} E^p\left[\left(\frac{I_0}{1-\tau^r}\frac{\Sigma_0}{\Sigma_0+\sigma_s^2}\right)^2 \Sigma_0\right] - \frac{\delta}{\rho^2} \\ &\quad -\frac{1}{2\rho^2}\frac{\sigma_k^2}{1-\pi}\left(1+\frac{1}{\rho}\pi\sigma_\xi^2\right). \end{aligned}$$

The first two pieces again relate to the efficiency of real investment and cross-sectional inequality among households, while the third reflects uncertainty over the current size of the liquidity shock. The direct contributions to welfare from uncertainty about investment productivity Σ_0 are unambiguously negative, and it is unlikely that informational frictions can improve real investment efficiency since firms can only be distorted away from the level

linear equations for $\mathcal{D}^{g*}p = 0$ to construct the matrix equation

$$A'\mathbf{p} = \mathbf{0}_{N_{\hat{\theta}}\cdot N_\Sigma \cdot N_I \times 1},$$

where $\mathbf{p} = vec(p)$ and A is the $\left(N_{\hat{\theta}} \cdot N_\Sigma \cdot N_I\right) \times \left(N_{\hat{\theta}} \cdot N_\Sigma \cdot N_I\right)$ square matrix that approximates the derivative operator $\mathcal{D}^g$ constructed with the "upwind" method. Here A' denotes the transpose of A. Since the matrix equation defines the stationary distribution for a Markov chain with transition matrix A', it follows by the Frobenius-Perron Theorem for nonnegative compact operators that A' has a unique largest eigenvalue (in absolute value), called the principal eigenvalue, and an associated strictly positive eigenvector ϕ unique up to a scaling factor. Since A is singular, it is convenient to replace one row i of A' with $A_{ij} = \delta_{ij}$ and the i^{th} entry of the zero vector with 1. This allows me to update to the stationary distribution in one step after defining A.

In practice, I find it convenient to populate the matrix A imposing that $\hat{\theta}^c$ has reflecting boundaries on both sides, and then set the boundaries sufficiently far into the tails of the distribution that the choice is insensitive to my results.

of investment they would choose with perfect-information. Welfare is about 1.9% lower compared to the perfect-information benchmark, and modestly about .5% higher than in the economy analogous to that of Van Nieuwerburgh and Veldkamp (2006) where households do not aggregate private information in financial markets. This modest gain reflects the tradeoff between the increased informativeness of public signals and the cross-sectional inequality induced by households trading on their heterogeneous private information.

To highlight the presence of information externalities in the economy, I conduct several illustrative thought experiments varying the position cost and real investment subsidy. I report the gain in welfare in consumption equivalent λ in the tradition of Lucas (1987).[26]

τ^r	.05	.1	.15
$100 \times \lambda$	0.128	0.266	0.411

Table 1: Transaction Cost Experiment

From Table 1, the position cost improves welfare in the economy with informational frictions. The intuition for this is that the gain in informational provision by having households take larger positions, is larger than the cost of generating more inequality by having households trade more on their heterogeneous private information.[27] Since better public information lowers the average level of uncertainty in the economy, however, this mitigates the rise in inequality.

[26] Formally, the consumption equivalent λ for an alternative level of the transaction tax or real investment subsidy that raises welfare from $\tilde{U}_p^c$ to U_p^c is defined as the fractional increase in the consumption of all households under the baseline level that delivers the same gain.

For log utility, λ satisfies

$$U_p^c = \frac{1}{\rho} E^p \left[\int_0^1 \log\left((1+\lambda)\, \tilde{c}(i) \right) di \right]$$

for $\tilde{U}_p^c = \frac{1}{\rho} E^p \left[\int_0^1 \log \tilde{c}(i)\, di \right]$, from which follows that

$$\lambda = \exp\left(\rho \left(U_p^c - \tilde{U}_p^c \right) \right) - 1.$$

[27] An important caveat is that the experiment understates the extent to which heterogeneous information generates wealth inequality because household private information is short-lived, and therefore there is no persistence in positions. With long-lived private information, the net benefit is likely to be more modest.

To see if subsidizing real investment improves welfare by improving the informational content of public information, I give firms a proportional investment subsidy τ^I whenever investment is at least one standard deviation below its unconditional mean in the stationary distribution. This has the interpretation of being a countercyclical real investment subsidy. To capture the welfare impact of the subsidy through the informational channel, I modify the experiment by subtracting out expected welfare under the perfect-information benchmark U_p^{perf}, since the subsidy will mechanically impact welfare by raising the average level of investment in the economy. It is easy to derive the analogous KFE for the perfect-information benchmark economy

$$-\partial_\theta \left\{p\lambda\left(\bar{\theta}-\theta\right)\right\} - \partial_I \left\{pI\left(\left(a-I\right)\theta-\left(1-\tau^c\right)\rho\right)\right\} \mathbf{1}_{\left\{I>\underline{I}\cup\theta\geq\frac{(1-\tau^c)\rho}{a-\underline{I}}\right\}} + \frac{1}{2}\sigma_\theta^2\partial_{\theta\theta}p = 0,$$

which has similar boundary conditions. Subtracting out the expected welfare under the perfect-information benchmark captures the incremental benefit of the subsidy from mitigating informational frictions.

τ^I	.05	.10	.15
$100 \times \lambda$	1.875	3.709	5.502

Table 2: Investment Subsidy Experiment

Table 2 reveals that the real investment subsidy also improves welfare. Since the subsidy increases real investment, which increases the average position households take in asset markets, it also has a similar effect to implementing a position cost. Real investment subsidies, therefore, improve the provision of public information by increasing the informativeness of both real and financial signals, which might, in part, explain why the gains from this experiment are larger than for the position cost.

These two thought experiments are meant to illustrate that there is a role for welfare-improving policies that address an information externality that arises because of decentralization. If instead of continuums, there were only one trader or one firm in the economy, such an agent would internalize its impact on the formation of the endogenous public signals when choosing its investment policies. While also likely to be present in static settings of incomplete information, this externality has a dynamic dimension because households and firms learn from signals formed inefficiently because of decentralization in the past. Though Greenwald and Stiglitz (1986) demonstrate that there often exist welfare-improving policy for economies with incomplete information and incomplete markets, their analysis is silent as to what form these policies take, and whether there is an optimal policy. These thought experiments motivate a more systematic analysis of policy interventions to address such information externalities within an optimal policy framework, which is beyond the scope of my analysis.

1.7. Empirical Implications

In this section, I explore several empirical implications of my framework that build off the observation that financial prices provide useful signals about the state of the economy, and that the strength of these signals is strongest during downturns and recoveries. I first discuss the asset pricing implications of my analysis, and then turn to conceptual issues my framework implies for empirical analysis and other empirical implications.

1.7.1. Implications for Asset Pricing

In this section, I characterize the business cycle implications of macroeconomic uncertainty in financial markets for asset risk premia and asset turnover. My analysis illustrates that, in the presence of informational frictions, there is an additional component to asset risk premia and asset turnover that reflects uncertainty about the state of the economy. This informational piece appears because households have heterogeneous private information and the degree to which they have heterogeneous beliefs increases as uncertainty rises about investment productivity. Furthermore, it gives asset returns predictive power for future returns and macroeconomic growth. The strength of this predictive power, however, varies over the business cycle, and I show that this variation is related to the behavior of asset turnover from informational trading.

1.7.1.1 Risk Premia

When the true state of the economy is known, then from Proposition 1.6 firms pay a risk premium on their claims

$$RP_{perf} = \rho - \frac{I}{a-I}g + I\theta - \delta - r = \underbrace{\frac{\sigma_k^2}{1-\pi} + \frac{\pi\sigma_k^2}{1-\pi}\xi}_{variance\ and\ liquidity\ risk},$$

which compensates households for variance risk and liquidity shocks. From Proposition 1.1, however, in the presence of informational frictions this risk premium includes an additional piece

$$RP_{NREE} = \underbrace{\frac{\sigma_k^2}{1-\pi} + \frac{\pi\sigma_k^2}{1-\pi}\xi}_{variance\ and\ liquidity\ risk} + \underbrace{\frac{I}{1+\Sigma/\sigma_s^2}\left(\theta - \hat{\theta}^c\right)}_{informational\ risk}.$$

that compensates investors for informational risk. This piece arises because households overreact to liquidity and capital quality shocks, and underreact to news about real investment productivity, driving a wedge between θ and $\hat{\theta}^c$. Similar to the risky asset demand of each household x_i from Proposition 1.1, the price of informational risk I is increasing in the level of investment by firms, while the quantity of informational risk $\frac{1}{1+\Sigma/\sigma_s^2}\left(\theta - \hat{\theta}^c\right)$ is increasing in the "average" pessimism of economic agents $\theta - \hat{\theta}^c$ and the level of informational frictions σ_s through the relative precision of public-to-private information Σ/σ_s^2. Consequently, investors earn risk compensation not only because of financial shocks and variance risk, but also because of distorted beliefs.

Similar to the speculative risk premium in Nimark (2012), this additional informational piece is, by construction, orthogonal to all public information, since $E\left[\frac{I}{1+\Sigma/\sigma_s^2}\left(\theta - \hat{\theta}^c\right) \mid \mathcal{F}^c\right] = 0$. Unlike the conditional mean, however, the conditional variance of this informational piece

$CV = E\left[\left(\frac{I}{1+\Sigma/\sigma_s^2}\right)^2\left(\theta-\hat{\theta}^c\right)^2 \mid \mathcal{F}^c\right] = \left(\frac{I}{1+\Sigma/\sigma_s^2}\right)^2\Sigma$ is, in principle, measurable by the econometrician. This conditional variance is increasing in investment I and can be hump-shaped in the conditional variance of beliefs Σ (since $\frac{dCV}{d\Sigma} = \left(\frac{\sigma_s^2}{1+\Sigma/\sigma_s^2}\right)^2 \frac{\sigma_s^2-\Sigma}{\sigma_s^2+\Sigma}$). Consequently, this informational risk premium contributes most to the time-variation in risk premia when I is sufficiently large and Σ is in an intermediate range.

To see how this informational component of risk premia affects the predictive power of asset prices for output, $Y_t = aK_t$, I integrate equation (1) from t to $s \geq t$ to find that output growth $\log\frac{Y_s}{Y_t}$ is given by

$$\log\frac{Y_s}{Y_t} = \int_t^s I_u\theta_u du + \sigma_k\left(Z_s^k - Z_t^k\right).$$

Using only public information, the covariance between output growth and expected excess returns in asset prices is

$$\begin{aligned} Cov\left[\log\frac{Y_s}{Y_t}, RP_{NREEt} \mid \mathcal{F}_t^c\right] &= \frac{I_t}{1+\Sigma_t/\sigma_s^2}Cov\left[\int_t^s I_u\theta_u du, \theta_t - \hat{\theta}_t^c \mid \mathcal{F}_t^c\right] \\ &\quad + \frac{\pi\sigma_k^2}{1-\pi}Cov\left[\int_t^s I_u\theta_u du, \xi_t \mid \mathcal{F}_t^c\right]. \end{aligned}$$

Since the riskless rate r_t is observable, $r_t \in \mathcal{F}_t^c$, I substitute for $\frac{\pi\sigma_k^2}{1-\pi}\xi_t$ with r_t from Proposition 1.1 to find

$$Cov\left[\log\frac{Y_s}{Y_t}, RP_{NREEt} \mid \mathcal{F}_t^c\right] = I_t Cov\left[\int_t^s I_u\theta_u du, \theta_t - \hat{\theta}_t^c \mid \mathcal{F}_t^c\right].$$

To turn off any mechanical correlation between expected excess returns and output growth, I consider the case where investment productivity shocks and liquidity shocks are uncorrelated $\alpha = 0$. In the absence of informational frictions, then, the covariance between risk premia and output growth is zero, since there is no misperception among firms or investors about θ_t, so $\hat{\theta}_t^c \equiv \theta_t$.

In the presence of informational frictions, however, this covariance is nonzero. Informational frictions introduce a short-run positive correlation between output growth and current risk premia since the true future investment productivity θ_u $u \geq t$ and investment are positively correlated at short-horizons with the true current level of investment productivity θ_t.[28] At longer horizon, the correlation weakens because of the mean-reversion in investment productivity θ_t and the potential fall in investment as it approaches its upper bound a. Since uncertainty Σ_t is countercyclical in my economy, the covariance also weakens around the peaks of business cycles, contributing to the countercyclical properties of asset price predictability for output growth. Similar insights hold for the relationship between expected returns and the growth in real investment.[29]

Substituting with r_t from Proposition 1.1, and recognizing that ξ_s and $\theta_t - \hat{\theta}_t^c$ are correlated only insofar as $\theta_t - \hat{\theta}_t^c$ is correlated with ξ_t, I also find that

$$\begin{aligned} Cov\left[\int_t^s RP_{NREEu}du, RP_{NREEt} \mid \mathcal{F}_t^c\right] &= I_t Cov\left[\int_t^s \frac{I_u}{1+\Sigma_u/\sigma_s^2}\left(\theta_u - \hat{\theta}_u^c\right)du, \theta_t - \hat{\theta}_t^c \mid \mathcal{F}_t^c\right] \\ &\quad + \frac{I_t^2\Sigma_t^2}{\Sigma_t + \sigma_s^2}(s-t), \end{aligned}$$

from which follows that $Cov\left[\int_t^s RP_{NREEu}du, RP_{NREEt} \mid \mathcal{F}_t^c\right]$ is positive. The correlation weakens at longer horizons because θ_t and $\hat{\theta}_t^c$ are mean-reverting.

Though there is this persistence in returns, households do not trade to eliminate this predictability. By the Law of Total Covariance, I can manipulate $Cov\left[\int_t^s RP_{NREEu}du, RP_{NREEt} \mid \mathcal{F}_t^c\right]$ to arrive at

$$\begin{aligned} & E\left[Cov\left[\int_t^s RP_{NREEu}du, RP_{NREEt} \mid \mathcal{F}_t^i\right] \mid \mathcal{F}_t^c\right] \\ = \; & Cov\left[\int_t^s RP_{NREEu}du, RP_{NREEt} \mid \mathcal{F}_t^c\right] \\ & - Cov\left[E\left[\int_t^s RP_{NREEu}du \mid \mathcal{F}_t^i\right], E\left[RP_{NREEt} \mid \mathcal{F}_t^i\right] \mid \mathcal{F}_t^c\right], \end{aligned}$$

[28] That future investment I_u and θ_t are positively correlated when investment is not close to its upper bound follows since the growth of investment I_u is increasing $\hat{\theta}_u^c$ from Proposition 1.4, and $\hat{\theta}_u^c = \theta_u + \varepsilon_u$ for some ε_u such that $E\left[\varepsilon_u \mid \mathcal{F}_u^c\right] = 0$, since $\hat{\theta}_u^c$ is an unbiased estimator of θ_u.

[29] My focus in this section is on conditional covariances. It is less clear that the signs and strengths of these covariances also hold unconditionally, since for random variables X, Y, and Z, by the Law of Total Covariance

$$Cov\left[X, Y\right] = E\left(Cov\left[X, Y \mid Z\right]\right) + Cov\left[E\left(X \mid Z\right), E\left(Y \mid Z\right)\right].$$

This implies that empirical tests would ideally focus on these conditional relationships.

from which it is apparent that the "average" perceived covariance of expected returns by household i $Cov\left[\int_t^s RP_{NREEu}du, RP_{NREEt} \mid \mathcal{F}_t^i\right]$ differs from the "average" covariance of expected returns $Cov\left[\int_t^s RP_{NREEu}du, RP_{NREEt} \mid \mathcal{F}_t^c\right]$ because of heterogeneous information. Consequently, households differ not only in their beliefs about expected returns, but also in their beliefs about the persistence of returns, which gives them incentive to trade without eliminating the predictability found with only public information.

This exercise illustrates that, in the presence of informational frictions, asset risk premia inherently contain an informational component that reflects uncertainty over current macroeconomic conditions above and beyond the correlation between real and financial shocks (since ξ may, in practice, be correlated with θ). Such a positive relationship between returns and future real activity, which arises because of the underreaction of investors to changes in the prospects of firms, is consistent, for instance, with the findings of Barro (1990), Fama (1990), and Schwert (1990). Moreover, this additional informational component exhibits countercyclical behavior, since uncertainty about investment productivity is countercyclical in the economy, and larger when financial markets are dysfunctional (larger, negative ξ shocks which depress $\hat{\theta}^c$). This may help explain why studies such as Stock and Watson (2003) and Ng and Wright (2013) find that the predictive power of asset prices for macroeconomic outcomes is somewhat episodic over business cycles, since the informational content of asset prices displays business cycle variation.

In addition to providing a measure of market liquidity ξ, which is documented, for instance, in Gilchrist, Yankov, and Zakrajsek (2009), market risk measures reflect the average expectations of market participants about the strength of the economy. This provides a strong empirical prediction that asset returns have predictive power for future returns and macroeconomic aggregates that varies with the business cycle, which is strongest during downturns and recoveries, and motivates more tests of asset pricing predictability that take this explicitly into account. Henkel, Martin, and Nardari (2011) and Dangl and Halling

(2012), for instance, provide evidence of business cycle asymmetries in stock market return predictability.

Given the risk premia from the firm's perspective RP_{NREE}, one can construct the risk premium demanded by an individual household to hold firm claims

$$RP_{NREE}(i) = RP_{NREE} + \frac{I}{1+\Sigma/\sigma_s^2}\left(\hat{\theta}^c - \hat{\theta}(i)\right).$$

Since $RP_{NREE}(i)$ is increasing in the pessimism of household i, lower $\hat{\theta}(i)$ relative to the average $\hat{\theta}^c$, it follows that more pessimistic households demand higher compensation to hold firm claims, and for sufficient pessimism instead sit on their capital by investing it in the riskless asset. This pattern is consistent with the tightening of lending standards seen in the FRB Senior Loan Officer Survey during the recent recession and recovery. In support of this prediction, the survey respondents often cited a poor economic outlook, along with bank competition, as a key factor in shaping their lending standards.

1.7.1.2 Asset Turnover

Though trading volume and asset turnover have been studied extensively in the literature, relatively little attention has been given to their business cycle properties.[30] Sarolli (2013) and DeJong and Espino (2011), for instance, provide evidence of business cycle variation in turnover. My analysis aims to help understand how differential information influences asset turnover over the business cycle and provides new empirical predictions.

To explore these issues, I derive a measure $\mathcal{V}$ on asset turnover (trading volume / shares outstanding) from informational trading at any given instant in the economy. To do so, I recognize that households that trade because of preference shocks with take an aggregate position -$\pi\xi W$ in firm claims, and that households that trade for informational and market-making reasons each invest a fraction of their wealth

$$x(i) = \frac{1}{1-\pi} + \frac{\pi}{1-\pi}\xi + \frac{I}{\sigma_k^2}\frac{\Sigma}{\Sigma+\sigma_s^2}\sigma_s Z^s(i),$$

[30] See Lo and Wang (2009) for a survey of this literature.

and take an aggregate position $(1+\pi\xi)\,W$. Intuitively, informational and market-making households take the offsetting position against liquidity traders plus a directional bet on the prospects of the economy based on the noise in their private signals. I thus construct a pseudo liquidity trader that takes a position $-\pi\xi W$ each period, and pseudo informational and market-making traders of mass $1-\pi$ that start with wealth W and always receive the same signal noise $Z^s(i)$.

This construction of pseudo traders is meant to mitigate the trading that arises because of preference shocks and the OLG structure of households, which mechanically leads to large changes in individual trader positions. I do not view the simplification as material for my results since I are abstracting from changes in positions that occur because of preference shocks and large changes in beliefs because of the myopic nature of households, which are both static effects over the business cycle.

The informational and market-making traders each enter the market with a position $X_I = x(i)\,W$ and will trade to have a position

$$\begin{aligned} dX_I &= W\frac{I}{\sigma_k^2}\frac{\Sigma\sigma_s}{\Sigma+\sigma_s^2}\left(g+\frac{\sigma_s^2}{\Sigma+\sigma_s^2}\frac{1}{\Sigma}\frac{d\Sigma}{dt}\right)Z^s(i)\,dt + x(i)\,W\left(I\theta-\delta-\frac{I}{a-I}g\right)dt \\ &\quad + \frac{\pi}{1-\pi}W\sigma_\xi dZ^\xi + x(i)\,W\sigma_k dZ^k. \end{aligned}$$

Following the insights of Xiong and Yan (2010), I aggregate the local volatility of these position changes and normalize by the price / share of firm claims q as a measure of trading volume[31]

$$\frac{1}{dt}E\left[v \mid q, W, \hat{\theta}^c, I, \Sigma\right] = \left(\frac{W}{q}\right)^2 \int_0^1 \left(\left(\frac{\pi}{1-\pi}\sigma_\xi\right)^2 + \left(x(i)\,\sigma_k\right)^2\right) di.$$

Substituting for $x(i)$ and applying the weak LLN, I arrive at

$$\frac{1}{dt}E\left[v \mid q, W, \hat{\theta}^c, I, \Sigma\right] = K^2\left(\frac{\pi^2}{1-\pi}\sigma_\xi^2 + \frac{(1+\pi\xi)^2}{1-\pi}\sigma_k^2 + (1-\pi)\left(\frac{I}{\sigma_k}\frac{\Sigma}{\Sigma+\sigma_s^2}\sigma_s\right)^2\right).$$

[31] Xiong and Yan (2010) motivates this measure by recognizing that the absolute value of realized position changes over small intervals is finite and increasing, on average, in the volatility of the position change.

From Section IV, $W = qK$, and therefore $K = W/q$ is the total number of shares outstanding for firm claims. From Section IV, $W = qK$ is the total market capitalization of firms, and therefore $K = W/q$ is the total number of shares outstanding for firm claims.

When $\sigma_s \nearrow \infty$, and there is no private information, then this expression reduces to

$$\frac{1}{dt} E\left[v^* \mid q, W, \hat{\theta}^c, I, \Sigma\right] = K^2 \left(\frac{\pi^2}{1-\pi} \sigma_\xi^2 + \frac{(1+\pi\xi)^2}{1-\pi} \sigma_k^2 \right),$$

which represents the level of pseudo trading volume not driven by information. Thus the difference $\frac{1}{dt} E\left[v \mid W, \hat{\theta}^c, I, \Sigma\right] - \frac{1}{dt} E\left[v^* \mid q, W, \hat{\theta}^c, I, \Sigma\right]$ normalized by total shares outstanding K delivers me my measure of share turnover from informational trading

$$\mathcal{V} = (1-\pi) \left(\frac{I}{\sigma_k} \frac{\Sigma}{\Sigma + \sigma_s^2} \sigma_s \right)^2.$$

When there is no asymmetric information among households, either $\sigma_s \searrow 0$, and households all know the hidden investment productivity θ, $\sigma_s \nearrow \infty$, and all households are equally naíve, or $\Sigma \searrow 0$, and there is no uncertainty about θ, then $\mathcal{V} \searrow \prime$, and there is no informational trading. Intuitively, households trade when they have heterogeneous information on which to speculate against each other.

Asset turnover $\mathcal{V}$ from informational trading is increasing in both real investment I and the level of uncertainty Σ. Similar to Xiong and Yan (2010), this measure of turnover is increasing in the disagreement among investors, as measured by Σ, since $\Sigma(i)$ is increasing in Σ. Li and Li (2014) provide evidence that belief dispersion about macroeconomic conditions positively correlates with stock market turnover. Asset turnover from informational trading is, consequently, strongest when real investment and uncertainty are in an intermediate range. This pattern helps us understand why market prices are most informative about investment productivity during downturns and recoveries, which is when a negative financial shock can be particularly devastating. Market prices have their highest information content during these parts of the business cycle because they are when households are trading intensely on their private information, and asset markets have high turnover.

1.7.2. Implications for Econometric Models

My analysis has several conceptual implications for empirical models that I now explore in this section. Building off the discussion in the previous section of the business cycle properties of risk premia in financial markets that arises because of learning, my analysis motivates econometricians to take advantage of this behavior for macroeconomic forecasting. Since real signals are procyclical, and those of financial markets are strongest during downturns and recoveries, a weighting scheme that weighs financial market data more heavily around troughs and real data near peaks is likely to be fruitful. My analysis also stresses the importance of including measures of uncertainty as forecasting variables because of the information aggregation channel in financial markets, yet cautions that uncertainty is itself endogenous and driven by fluctuations in both the real economy and financial markets.

A second econometric issue my model highlights occurs when the econometrician tries to disentangle the channels by which financial market dysfunction propagates to the real economy in the presence of informational frictions using structural vector autoregressions (SVARs) or factor models.[32] Since a financial market shock impacts expectations about the real economy through learning from prices, it is, in part, perceived as a negative shock to real economic fundamentals. Specifically, the riskless rate in my economy is the sum of real investment productivity θ_t and the aggregate market liquidity shock ξ_t. In the presence of informational frictions, however, firms decompose θ_t and ξ_t instead into their perceived counterparts, $\hat{\theta}_t^c$ and $\hat{\xi}_t^c$, respectively. For them to react to the financial market shock, it must be the case that this decomposition results in $\hat{\theta}_t^c < \theta_t$ and $\hat{\xi}_t^c < \xi_t$, and thus the shock propagates to the real economy by depressing firm expectations about θ_t. This highlights an invertibility issue that arises when firms learn from prices when making real decisions that prevents the econometrician from finding an orthonormal rotation that can recover the true

[32] There are abundant similarities in recovering structural shocks from reduced-form VARs and from factor models, since factor innovations estimated by principal components are unique only up to orthonormal rotations of the $SO(n)$ group.

historical decomposition of structural financial market shocks from reduced-form VAR or factor model innovations.[33][34]

Finally, a third implication of learning from financial markets over the business cycle is that shocks to uncertainty are inherently entangled with shocks to financial markets. As illustrated in Section V, prices that measure financial distress, such as market risk premia and credit spreads, can contain an informational component in the presence of informational frictions that reflects uncertainty about current economic conditions. Since private agents learn from prices, adverse financial shocks will affect the conditional variance of their expectations, as can be seen in Proposition 1.3 in Section IV, and consequently will also propagate through the economy as uncertainty shocks back to prices. This makes it difficult to separate structural shocks stemming from financial market dislocation from innovations to uncertainty because of learning, and relates to the use of prices as external instruments in disentangling these structural shocks from reduced-form VAR and factor model innovations. Such a channel, for example, can help explain the high correlation between the recovered financial distress and uncertainty shocks found in Stock and Watson (2012).[35]

1.7.3. Other Emprical Implications

Several additional empirical predictions of the impact of feedback in learning merit mention. First, since uncertainty in my framework is countercyclical and downturns stem from real shocks to investment productivity, my model is consistent with the observations of Naka-

[33] This invertibility issue is different from the one that arises because private agents and the econometrician have nested information sets, as explored, for instance, in Hansen and Sargent (1991) and Leeper, Walker, and Yang (2013). There is a large literature on dealing with news shocks when agents have superior information to the economerician. See, for instance, Beaudry and Portier (2006), Fujiwara, Hirose, and Shintani (2011), and Schmitt-Grohé and Uribe (2012).

[34] Sockin and Xiong (2014) make a similar point about trying to disentangle supply and demand shocks in commodity markets in the presence of informational frictions.

[35] Stock and Watson (2012) use innovations to the VIX and the poliicy news uncertainty index of Baker, Bloom, and Davis (2013) as instruments for uncertainty shocks. The VIX, as a measure of market volatility, has a direct analogue with prices in my economy. Innovations to the policy uncertainty index have a correlation of about 0.2 with the forecast dispersion of the Survey of Professional Forecasters, which can be viewed as a noisy analogue of uncertainty in my economy.

mura et al (2012) that, unconditionally, first moment shocks are negatively correlated with movements in uncertainty.

Second, while not the central focus of my analysis, another implication of asymmetric learning over the business cycle with dispersed information is that my model predicts countercyclical dispersion in wealth across households, a feature consistent with evidence from the latest recession.[36] This arises because informational frictions are most severe at the trough, where agents have incentive to trade on their private information, whereas, at the peak, uncertainty about the underlying strength of the economy Σ is low and households coordinate around the common knowledge belief $\hat{\theta}^c$ (since Σ/σ_s^2 is small).

Finally, my model features asset prices as a coordination mechanism among firms in making their investment decisions. My model, therefore, offers an additional information channel through which learning by individual firms can give rise to the strong comovement in macroeconomic aggregates documented in Christiano and Fitzgerald (1998), and since heavily exploited through factor model analysis in the macroeconometric literature. This channel is distinct from the information externality channel informally discussed in Christiano and Fitzgerald (1998), as well as the mechanism of strategic complementarity in common information that arises because of costly sector-specific information acquisition featured in Veldkamp and Wolfers (2007).

1.8. Conclusion

In this paper, I develop a dynamic model of information aggregation in financial markets in a macroeconomic setting where both financial investors and firm managers learn about the productivity of investment from market prices. My dynamic framework features a feedback loop between investor trading behavior and firm real investment decisions by which noise in

[36] Since the noise in household private signals is unbiased, the wealth distribution is a mean-preserving spread of the wealth of an agent who has perfect-information. The wealth of this perfectly-informed pseudo-agent will, in general, not be the same as the wealth of the representative household in either benchmark because heterogeneous information impacts both investment decisions and the risk premia on firm claims.

financial prices can feed into real investment through learning by firm managers, and then feed back into financial prices through the impact of learning and investment on the trading incentives of market participants. This feedback loop highlights a possible amplification mechanism through which the financial crisis of 2008 contributed to the deep recession and anemic recovery in the US by distorting firm expectations about the strength of the US economy.

While the strength of signals from real activity is procyclical, that of financial signals is strongest during downturns and recoveries. This occurs because the value of private information that financial investors have increases with uncertainty about real investment productivity, which is countercyclical, and more information is aggregated into prices as investors start to trade against each other on their private information. As a result, financial signals are strongest when real investment and uncertainty are in an intermediate range.

I then explore the welfare and empirical implications of my model. Informational frictions introduce a role for policy to provide guidance to economic agents about the current state of the economy. As an empirical prediction of my model, informational frictions also give rise to an informational component in asset risk premia that has predictive power for future returns and real activity. This predictive power is greatest during downturns and recoveries when asset turnover from informational trading is highest. Finally, informational frictions make it difficult to disentangle the effects of financial and uncertainty shocks in the data, and confound attempts to recover historical structural shocks stemming from the financial crisis of 2008.

1.A: Proofs of Propositions

Proof of Proposition 1.2:

Households solve the optimization problem (4) subject to equation (9). In a recursive competitive equilibrium, all equilibrium objects are functions of the state of the economy from the household's perspective $\left(w\left(i\right),\hat{\theta}\left(i\right),\xi\left(i\right),h\right)$, where h is a list of general equilibrium objects including $\log K$ and r.[37] By the Martingale Representation Theorem for $\mathcal{L}^2$ processes, all these objects will be continuous Itô-semimartingales with respect to the smallest filtration on which they are measurable to the household. The Wiener processes to which they are adapted, which will be common to all households, are absolutely continuous with respect to the true processes for real investment productivity θ, household liquidity shocks ξ, and the aggregate diffusion for K.

Taking the limit of problem (4) as $\Delta t \searrow dt$, assuming v is twice differentiable in its arguments, I can differentiate v and take expectations to find

$$\rho v = \sup_{\{c,x\}} \log c + \partial_w v \frac{1}{dt} E\left[dw\left(i\right) \mid \mathcal{F}^i\right] + \frac{1}{2}\partial_{ww} v \frac{1}{dt} d\left\langle w\left(i\right) \mid \mathcal{F}^i_t\right\rangle + \frac{1}{dt}\partial_t v, \tag{A-1}$$

subject to the law of motion of $w\left(i\right)$ (9), and $\langle \cdot \mid \mathcal{F}^i_t\rangle$ indicates quadratic variation under the measure $\mathcal{F}^i_t$. The $\partial_t v$ term is meant to capture the additional dependence of the drift of the household's bequest utility v on the vector of general equilibrium objects h that the household takes as given. Equation (A-1) is the usual Hamilton-Jacobi-Bellman (HJB) equation for optimal control. Necessity and sufficiency of the FOCs for the optimal controls $\{c,x\}$ follows from the concavity of their programs.

Before deriving the FOCs of the HJB equation (A-1) for households, it is useful to first recognize that all Wiener processes $\tilde{Z}^{\xi}_t\left(i\right)$ and $\tilde{Z}^{k}_t\left(i\right)$ will be uncorrelated under each household $i's$ measure since the true processes are uncorrelated and the change of measure under Girsanov's Theorem is equivalent to a change in drift.

Suppressing arguments for the bequest utility v, the FOCs of the HJB equation (A-1)

[37] Since the household treats prices as exogenous, the price of firm claims q and the riskless rate r are additional states for the household. This, however, only affects their optimal consumption and portfolio choices, in which they do not see the dependence of these prices on the Markov states.

are given by

$$\begin{aligned} c(i) &: \frac{1}{c(i)} - \partial_w v \leq 0 \ \ (= \ \ if \ c > 0), \\ x(i) &: 0 = w(i)\,\partial_w v\left(\frac{a-I}{q} + \frac{\partial_I q}{q} Ig + I\hat{\theta}(i) - r - \delta\right) + x(i)\,w(i)^2\,\partial_{ww} v \sigma_k^2 \\ & \quad + w(i)\,\partial_{wh} v d\left\langle \tilde{Z}^k(i), h \mid \mathcal{F}^i \right\rangle, \end{aligned}$$

when household i is not hit by the liquidity shock $l(i) = 0$, from which follows that

$$x(i) = -\frac{\partial_w v\left(\frac{a-I}{q} + \frac{\partial_I q}{q} Ig + I\hat{\theta}(i) - r - \delta\right)}{w(i)\,\partial_{ww} v \sigma_k^2} - \frac{\partial_{wh} v d\left\langle \tilde{Z}^k(i), h \mid \mathcal{F}^i \right\rangle}{w(i)\,\partial_{ww} v \sigma_k^2}.$$

While objects in h like r all have Itô-semimartingale representations by the Martingale Representation Theorem, I do not expand out the quadratic covariation expressions for brevity.

Given that households have log utility, I conjecture that $v\left(w(i), \hat{\theta}(i), l(i), h\right) = A \log w(i) + f\left(\hat{\theta}(i), l(i), h\right)$. This conjecture implies that

$$\begin{aligned} c(i) &= \frac{w(i)}{A}, \\ x(i) &= \begin{cases} \frac{\frac{a-I}{q} + \frac{\partial_I q}{q} Ig + I\hat{\theta}(i) - r - \delta}{\sigma_k^2} & l(i) = 0 \\ -\xi & l(i) = 1 \end{cases}. \end{aligned}$$

Substituting this conjecture and the controls into the maximized HJB equation

$$\begin{aligned} \rho v &= \log c + \partial_w v\left(x(i)\left(\frac{a-I}{q} + \frac{\partial_I q}{q} Ig + I\hat{\theta}(i) - r - \delta\right) w(i) + rw(i) - c(i)\right) \\ & \quad + \frac{1}{2}\partial_{ww} v x(i)^2\, w(i)^2\, \sigma_k^2 + \partial_t f\left(\hat{\theta}(i), l(i), h\right), \end{aligned}$$

where $\partial_t f\left(\hat{\theta}(i), l(i), h\right)$ is shorthand for remaining terms in the HJB equation, it follows that $A = \frac{1}{\rho}$, $c(i) = \rho w(i)$, and that $f\left(\hat{\theta}(i), l(i), h\right)$ implicitly satisfies

$$\begin{aligned} \rho f\left(\hat{\theta}(i), l(i), h\right) &= \log \rho + \frac{1}{\rho}\left(r - \rho + x(i)\left(\frac{a-I}{q} + \frac{\partial_I q}{q} Ig + I\hat{\theta}(i) - r - \delta\right) - \frac{1}{2} x(i)^2\, \sigma_k^2\right) \\ & \quad + \partial_t f\left(\hat{\theta}(i), l(i), h\right), \end{aligned}$$

which confirms the conjecture since $x(i)$ does not depend on $w(i)$.

When the household is hit by the liquidity shock, $l(i) = 1$, then $x(i) = -\xi$. Direct verification of the value function $v\left(w(i), \hat{\theta}(i), l(i), h\right) = A \log w(i) + f\left(\hat{\theta}(i), l(i), h\right)$

in the maximized HJB equation again confirms the conjectured functional form and that $c(i) = \rho w(i)$.

Recognizing that $v\left(w(i), \hat{\theta}(i), l(i), h\right) = A \log w(i) + f\left(\hat{\theta}(i), h\right)$, the envelope condition for the maximized HJB equation (A-1) evaluated at the optimal controls takes the form

$$\begin{aligned}
\rho \partial_w v &= \partial_{ww} v \left(x(i) w(i) \left(\frac{a-I}{q} + \frac{\partial_I q}{q} Ig + I\hat{\theta}(i) - r - \delta \right) + rw(i) - c(i) \right) \\
&\quad + \frac{1}{2} \partial_{www} v x(i)^2 w(i)^2 \sigma_k^2 + \partial_{ww} v x(i)^2 w(i) \sigma_k^2 \\
&\quad + \partial_w v \left(x(i) \left(\frac{a-I}{q} + \frac{\partial_I q}{q} Ig + I\hat{\theta}(i) - r - \delta \right) + r \right).
\end{aligned}$$

Applying Itô's Lemma directly to $\partial_w v$, one also has that

$$\begin{aligned}
d(\partial_w v) &= \partial_{ww} v \left(x(i) w(i) \left(\frac{a-I}{q} + \frac{\partial_I q}{q} Ig + I\hat{\theta}(i) - r - \delta \right) + rw(i) - c(i) \right) dt \\
&\quad + \frac{1}{2} \partial_{www} v x(i)^2 w(i)^2 \sigma_k^2 + \partial_{ww} v x(i) w(i) \sigma_k d\tilde{Z}^k.
\end{aligned}$$

Taking expectations and substituting the envelope condition, it follows that

$$\frac{1}{dt} E\left[\frac{d(\partial_w v)}{\partial_w v} \mid \mathcal{F}^i \right] = \rho - r - x(i) \left(\frac{a-I}{q} + \frac{\partial_I q}{q} Ig + I\hat{\theta}(i) - r - \delta \right) - \frac{\partial_{ww} v}{\partial_w v} x(i)^2 w(i) \sigma_k^2.$$

Given $\partial_w v = \frac{1}{w}$, the solution for $x(i)$ when the household is not hit by the liquidity shock, and defining $\Lambda_t(i) = e^{-\rho t} \frac{1}{w_t(i)}$ to be the pricing kernel of household i, it follows that

$$r = -\frac{1}{dt} E\left[\frac{d\Lambda(i)}{\Lambda(i)} \mid \mathcal{F}^i \right]. \tag{A.2}$$

From $\Lambda_t(i) = e^{-\rho t} \frac{1}{w_t(i)}$, the optimal choice of $x(i)$, and equation (A.2), it follows that

$$\frac{a-I}{q} dt + E\left[\frac{d\Lambda(i)}{\Lambda(i)} + \frac{d(qK)}{qK} \mid \mathcal{F}^i \right] = x(i) \sigma_k^2 dt. = -Cov\left[\frac{d(qK)}{qK}, \frac{d\Lambda(i)}{\Lambda(i)} \mid \mathcal{F}^i \right],$$

from which one arrives at

$$\frac{a-I}{qK} K dt + E\left[\frac{d(\Lambda(i) qK)}{\Lambda(i) qK} \mid \mathcal{F}^i \right] = 0,$$

for household i not hit by the liquidity shock, which completes the proof.

Proof of Proposition 1.3:

Define $\bar{R}_\theta(\zeta_t) = R_\theta(I_t, \Sigma_t)$, and $\bar{g}_t(\zeta_t) = g_t$. Given ζ_t, one can express the law of motion of the vector of public signals as

$$\begin{aligned} d\zeta_t &= A_0(\zeta_t)\,dt + \begin{bmatrix} I_t \\ \partial_\Sigma \bar{R}_\theta(\zeta_t)\frac{d\Sigma_t}{dt} + \partial_I \bar{R}_\theta(\zeta_t)\, I_t \bar{g}_t(\zeta_t) - \lambda \bar{R}_\theta(\zeta_t) \end{bmatrix} \theta_t dt \\ &\quad + \bar{b}_t(\zeta_t)\, dZ_t^\theta + \bar{B}_t(\zeta_t)\, dZ_t, \end{aligned}$$

where $Z_t = \left[Z_t^k, Z_t^\xi\right]'$ and

$$\begin{aligned} A_0(\zeta_t) &= \begin{bmatrix} -\delta - \frac{1}{2}\sigma_k^2 \\ \bar{R}_\theta(\zeta_t)\lambda\bar{\theta} \end{bmatrix} \\ \bar{b}_t(\zeta_t) &= \begin{bmatrix} 0 \\ \bar{R}_\theta(\zeta_t)\sigma_\theta + \alpha\sigma_\xi \end{bmatrix}, \\ \bar{B}_t(\zeta_t) &= \begin{bmatrix} \sigma_k & 0 \\ 0 & \sqrt{1-\alpha^2}\sigma_\xi \end{bmatrix}, \end{aligned}$$

with $\bar{R}_\theta(\zeta_t)$ uniformly bounded and $\bar{R}_\theta(\zeta_t) > 0 \;\forall\; \zeta_t$. By Theorem 7.17 of Lipster and Shiryaev (1977), then one can construct the vector of standard Wiener processes $\tilde{Z} = \left(\tilde{Z}_t, \mathcal{F}_t^c\right)$ where $\tilde{Z}_t = \left[\tilde{Z}_t^k, \tilde{Z}_t^r\right]'$ admits the representation

$$\begin{aligned} \tilde{Z}_t &= \int_0^t \left[\bar{b}_t(\zeta_t)\,\bar{b}_t(\zeta_t)' + \bar{B}_s(\zeta_s)\,\bar{B}_s(\zeta_s)'\right]^{-1/2} \times \\ &\quad \left(d\zeta_s - A_0(\zeta_t)\,dt - \begin{bmatrix} I_t \\ \partial_\Sigma \bar{R}_\theta(\zeta_t)\frac{d\Sigma_t}{dt} + \partial_I \bar{R}_\theta(\zeta_t)\, I_t \bar{g}_t(\zeta_t) - \lambda \bar{R}_\theta(\zeta_t) \end{bmatrix} \hat{\theta}_t^c dt\right), \end{aligned}$$

where $\hat{\theta}_t^c = E[\theta_t \mid \mathcal{F}_t^c]$ is the conditional expectation of θ_t w.r.t. $\mathcal{F}_t^c$. That $\tilde{Z}$ are standard Wiener processes can be verified directly from Levy's three properties that uniquely identify Wiener processes. That $\tilde{Z}$ is a martingale generator for $\mathcal{F}_t^c$ follows since $\tilde{Z}$ generates K and r trivially, from which the other objects of $\mathcal{F}_t^c$ can be generated, and Lemma 4.9 guarantees the existence of a representation for the driver (which possibly depends on the unobservable θ_t) in the Martingale Representation Theorem (Theorem 5.8) that is measurable w.r.t ζ_t P$-a.s.$

Given that θ_t has the representation

$$\theta_t = \int_0^t \lambda\left(\bar{\theta} - \theta_s\right) ds + \int_0^t \sigma_\theta dZ_s^\theta,$$

it follows from similar arguments that lead to the proof of Theorem 12.7 that $\hat{\theta}_t^c$ has the representation

$$\begin{aligned} \hat{\theta}_t^c &= \int_0^t \left(d\left\langle \begin{matrix} S \\ Q \end{matrix}, \sigma_\theta Z^\theta \right\rangle_s + Cov\left[\theta_s, \begin{bmatrix} I_s \\ \partial_\Sigma \bar{R}_\theta(\zeta_t)\frac{d\Sigma_t}{dt} + \partial_I \bar{R}_\theta(\zeta_t)\, I_s \bar{g}_t(\zeta_t) - \lambda \bar{R}_\theta(\zeta_t) \end{bmatrix} \theta_s \mid \mathcal{F}_s^c\right]'\right) \times \\ &\quad \left[\bar{b}_t(\zeta_t)\,\bar{b}_t(\zeta_t)' + \bar{B}_s(\zeta_s)\,\bar{B}_s(\zeta_s)'\right]^{-1/2} d\tilde{Z}_s + \int_0^t \lambda\left(\bar{\theta} - \hat{\theta}_s^c\right) ds, \end{aligned} \tag{A.3}$$

where $d\left\langle \xi, Z^{\theta}\right\rangle_t$ is the quadratic covariation of ξ_t and Z_t^{θ}. It is easy to see that $Cov\left[\theta_s, \theta_s \mid \mathcal{F}_s^c\right] = Var\left[\theta_s \mid \mathcal{F}_s^c\right] = \Sigma_s$. The covariance matrix in equation (A.3) is given by

$$\bar{b}_t\left(\zeta_t\right)\bar{b}_t\left(\zeta_t\right)' + \bar{B}_s\left(\zeta_s\right)\bar{B}_s\left(\zeta_s\right)' = \begin{bmatrix} \sigma_k^2 & 0 \\ 0 & \left(\bar{R}_\theta\left(\zeta_t\right)\sigma_\theta + \alpha\sigma_\xi\right)^2 + \left(1-\alpha^2\right)\sigma_\xi^2 \end{bmatrix},$$

from which follows that

$$\left[\bar{b}_t\left(\zeta_t\right)\bar{b}_t\left(\zeta_t\right)' + \bar{B}_s\left(\zeta_s\right)\bar{B}_s\left(\zeta_s\right)'\right]^{-1/2} = \begin{bmatrix} \frac{1}{\sigma_k} & 0 \\ 0 & \frac{1}{\sqrt{\left(\bar{R}_\theta(\zeta_t)\sigma_\theta + \alpha\sigma_\xi\right)^2 + (1-\alpha^2)\sigma_\xi^2}} \end{bmatrix}$$

Thus it follows that $\hat{\theta}_t^c$ follows the law of motion

$$\begin{aligned} d\hat{\theta}_t^c &= \frac{\bar{R}_\theta\left(\zeta_t\right)\sigma_\theta^2 + \alpha\sigma_\xi\sigma_\theta + \left(\partial_\Sigma\bar{R}_\theta\left(\zeta_t\right)\frac{d\Sigma_t}{dt} + \partial_I\bar{R}_\theta\left(\zeta_t\right)I_t\bar{g}_t\left(\zeta_t\right) - \lambda\bar{R}_\theta\left(\zeta_t\right)\right)\Sigma_t}{\sqrt{\left(\bar{R}_\theta\left(\zeta_t\right)\sigma_\theta + \alpha\sigma_\xi\right)^2 + \left(1-\alpha^2\right)\sigma_\xi^2}} d\tilde{Z}_t^r \\ &\quad + I_t\frac{\Sigma_t}{\sigma_k}d\tilde{Z}_t^k + \lambda\left(\bar{\theta} - \hat{\theta}_t^c\right)dt. \end{aligned}$$

Given the common Gaussian prior of households $N\left(\hat{\theta}_0^c, \Sigma_0\right)$, establishing the conditional Gaussianity of the posterior $\theta_t \mid \mathcal{F}_t^c$ can be done through similar arguments to those made in Chapter 11 of Lipster and Shiryaev (1977) with the appropriate regularity conditions. Similar to the arguments of Theorem 12.7, one can the also establish that the conditional variance of beliefs $\Sigma_t = Var\left[\theta_t \mid \mathcal{F}_t^c\right]$ follows the deterministic law of motion

$$\begin{aligned} \frac{d\Sigma_t}{dt} &= \sigma_\theta^2 - \frac{\left(\bar{R}_\theta\left(\zeta_t\right)\sigma_\theta^2 + \alpha\sigma_\xi\sigma_\theta + \left(\partial_\Sigma\bar{R}_\theta\left(\zeta_t\right)\frac{d\Sigma_t}{dt} + \partial_I\bar{R}_\theta\left(\zeta_t\right)I_t\bar{g}_t\left(\zeta_t\right) - \lambda\bar{R}_\theta\left(\zeta_t\right)\right)\Sigma_t\right)^2}{\left(\bar{R}_\theta\left(\zeta_t\right)\sigma_\theta + \alpha\sigma_\xi\right)^2 + \left(1-\alpha^2\right)\sigma_\xi^2} \\ &\quad - 2\lambda\Sigma_t - I_t^2\frac{\Sigma_t^2}{\sigma_k^2}, \end{aligned} \tag{A.4}$$

which is a second-order polynomial in $\frac{d\Sigma_t}{dt}$, from which follows from equation (A.4) that

$$\frac{d\Sigma_t}{dt} = -\frac{B\left(\zeta_t\right)}{2A\left(\zeta_t\right)} \pm \frac{1}{2A\left(\zeta_t\right)}\sqrt{2B\left(\zeta_t\right) - 4A\left(\zeta_t\right)\left(2\lambda\Sigma_t - \sigma_\theta^2 + I_t^2\frac{\Sigma_t^2}{\sigma_k^2}\right) - 1},$$

where

$$\begin{aligned} A\left(\zeta_t\right) &= \frac{\left(\partial_\Sigma\bar{R}_\theta\left(\zeta_t\right)\Sigma_t\right)^2}{\left(\bar{R}_\theta\left(\zeta_t\right)\sigma_\theta + \alpha\sigma_\xi\right)^2 + \left(1-\alpha^2\right)\sigma_\xi^2}, \\ B\left(\zeta_t\right) &= 1 + 2\partial_\Sigma\bar{R}_\theta\left(\zeta_t\right)\Sigma_t\frac{\bar{R}_\theta\left(\zeta_t\right)\sigma_\theta^2 + \alpha\sigma_\xi\sigma_\theta + \left(\partial_I\bar{R}_\theta\left(\zeta_t\right)I_t\bar{g}_t\left(\zeta_t\right) - \lambda\bar{R}_\theta\left(\zeta_t\right)\right)\Sigma_t}{\left(\bar{R}_\theta\left(\zeta_t\right)\sigma_\theta + \alpha\sigma_\xi\right)^2 + \left(1-\alpha^2\right)\sigma_\xi^2}. \end{aligned}$$

Substituting $R_\theta = -\frac{1-\pi}{\pi}\frac{I}{\sigma_k^2}\frac{\Sigma}{\Sigma+\sigma_s^2}$ into the above expressions delivers the laws of motion stated in the proposition.

The conditional variance of beliefs Σ is trivially bounded from below by 0. To find the upper bound, consider the case when all public signals are completely uninformative $\forall\ t$, then Σ follows the law of motion

$$\frac{d\Sigma}{dt} = \sigma_\theta^2 - 2\lambda\Sigma_t,$$

which has the steady-state solution $\Sigma_t = \frac{\sigma_\theta^2}{2\lambda}$. Since any informativeness of the public signals reduces the conditional variance of beliefs, $\Sigma_t \leq \frac{\sigma_\theta^2}{2\lambda}$.

To find the relationship between $\hat{\theta}_t^c$ and $\hat{\theta}_t^c(i)$ for households, I make use of the Law of Iterated Expectations to write

$$\hat{\theta}_t^c(i) = E\left[\theta_t \mid \mathcal{F}_t^i\right] = E\left[\theta_t^c \mid s_t(i)\right],$$

where $\theta_t^c = \theta_t \mid \mathcal{F}_t^c$. Consider the common knowledge estimate $\hat{\theta}_t^c$, one I can arrive at the estimate of household i $\hat{\theta}_t(i)$ by updating $\mathcal{F}_t^c$ with household i's private signal $s_t(i)$. Since both the average household estimate $\hat{\theta}_t^c$ and the signal $s_t(i)$ are jointly Gaussian, which is apparent from the linearity of the Kalman Filter in the data $\{\zeta_s, \theta_s\}_{s\leq t}$, the process of updating the conditional mean is an exercise in the updating of two sets of Gaussian random variables. It then follows that

$$\begin{aligned}
\hat{\theta}_t(i) &= \hat{\theta}_t^c + Cov\left[\theta_t, s_t(i) \mid \mathcal{F}_t^c\right] Var\left[s_t(i) \mid \mathcal{F}_t^c\right]^{-1}\left(s_t(i) - E\left[s_t(i) \mid \mathcal{F}_t^c\right]\right) \\
&= \hat{\theta}_t^c + \frac{\Sigma_t}{\Sigma_t + \sigma_s^2}\left(s_t(i) - \hat{\theta}_t^c\right).
\end{aligned}$$

Similarly, the conditional variance of household $i's$ estimate of θ is

$$\begin{aligned}
\Sigma_t(i) &= \Sigma_t - Cov\left[\theta_t, s_t(i) \mid \mathcal{F}_t^c\right] Var\left[s_t(i) \mid \mathcal{F}_t^c\right]^{-1} Cov\left[\theta_t, s_t(i) \mid \mathcal{F}_t^c\right] \\
&= \Sigma_t - \frac{\Sigma_t^2}{\Sigma_t + \sigma_s^2} = \frac{\sigma_s^2}{\Sigma_t + \sigma_s^2}\Sigma_t.
\end{aligned}$$

Proof of Proposition 1.4:

To find the optimal level of investment I, let me conjecture that $E = E(t, K, I)$. Then, by the Feyman-Kac Theorem and $\frac{\Lambda_t}{\Lambda_0}, E_t > 0$, the function E that solves each manager's problem (5) must solve the necessary condition

$$0 \geq \sup_{g_t} \frac{\left(a - I_t - \frac{1}{\rho}g_t I_t + \tau_t\right)K_t}{E_t}dt + E\left[\frac{d\left(\Lambda_t E_t\right)}{E\left[\Lambda_t \mid \mathcal{F}_t^c\right]E_t} \mid \mathcal{F}_t^c\right],$$

which can be rewritten as

$$0 \geq \sup_{g_t} \frac{\left(a - I_t - \frac{1}{\rho} g_t I_t + \tau_t\right) K_t}{E_t} dt + E\left[\frac{dE_t}{E_t} \mid \mathcal{F}_t^c\right] + \frac{E\left[d\Lambda_t \mid \mathcal{F}_t^c\right]}{E\left[\Lambda_t \mid \mathcal{F}_t^c\right]} + \frac{d\left\langle \Lambda_t, E_t \mid \mathcal{F}_t^c\right\rangle}{E\left[\Lambda_t \mid \mathcal{F}_t^c\right] E_t}. \quad \text{(A.6)}$$

By Proposition 1.2, the pricing kernel of investor j $\Lambda_t(j)$ satisfies $\frac{1}{dt} E\left[\frac{d\Lambda_t(j)}{\Lambda_t(j)} \mid \mathcal{F}^j\right] = -r_t$. Thus, by the Law of Iterated Expectations, $E\left[\frac{d\Lambda_t}{\Lambda_t} \mid \mathcal{F}^c\right] = -r_t$, regardless of the distribution of ownership among households. Then, applying Itô's Lemma to E, equation (A.6) becomes

$$\begin{aligned} 0 \quad \geq \quad & \sup_{g_t} \frac{a - I_t}{E_t} K_t - \frac{I_t}{\rho} g_t K_t + \tau_t K_t + \frac{\partial_K E_t}{E_t} \left(I_t \hat{\theta}_t^c - \delta\right) K_t + \frac{1}{2} \frac{\partial_{KK} E_t}{E_t} \sigma_k^2 K_t^2 + \partial_t E_t \\ & -r_t + \frac{1}{dt} \frac{d\left\langle \Lambda_t, E_t \mid \mathcal{F}_t^c\right\rangle}{E\left[\Lambda_t \mid \mathcal{F}_t^c\right] E_t}, \end{aligned} \quad \text{(A.7)}$$

where $\frac{1}{dt} \frac{d\langle \Lambda_t, E_t \mid \mathcal{F}_t^c\rangle}{\Lambda_t E_t}$ is the risk premium on firm claims. Since firms are perfectly competitive, they do not recognize, in equilibrium, that their actions affect the riskless rate r_t or the pricing kernel of shareholders Λ_t.

Firm effort g_t is chosen by the firm to achieve its optimal level of investment. Since equation (A.7) is (locally) riskless and linear in investment I_t, firm managers are effective risk-neutral and it follows that it must be the case that g_t satisfies

$$-1 + \partial_K E_t \hat{\theta}_t^c - \frac{1}{\rho} g_t = 0, \quad \text{(A.8)}$$

or else there is a riskless gain to changing g if the marginal return to investment for firm value is positive or negative. By market clearing, the value of firm claims must be such that $E_t = q_t K_t$, where $q_t = \frac{a - I_t}{\rho}$. To see that $E_t = q_t K_t$ satisfies the maximized form of equation (A.6), recall from Proposition 1.2 that $E_t = q_t K_t$ satisfies at the optimal I_t

$$\Lambda_t(i) \frac{a - I_t}{E_t} K_t dt + E\left[\frac{d\left(\Lambda_t(i) E_t\right)}{E_t} \mid \mathcal{F}_t^i\right] = 0.$$

Let $u_t(i)$ be the share of the firm owned by household i that has not experienced a preference shock, such that $\Lambda_t = \int u_t(i) \Lambda_t(i)\, di$. Assuming that the firm equal weights the pricing kernels of investing households $\Lambda_t = \int u_t(i) \Lambda_t(i)\, di = \int e^{-\rho t} u_t(i) \frac{1}{w_t(i)} di$, then it follows, by linearity and the finiteness of Λ_t, that

$$\frac{1}{dt} \frac{d\left\langle \Lambda_t, E_t \mid \mathcal{F}_t^c\right\rangle}{E\left[\Lambda_t \mid \mathcal{F}_t^c\right] E_t} = \int u_t(i) \frac{1}{dt} \frac{d\left\langle \frac{1}{w_t(i)}, K_t \mid \mathcal{F}_t^c\right\rangle}{K_t E\left[\int u_t(i) \frac{1}{w_t(j)} dj \mid \mathcal{F}_t^c\right]} di = -\sigma_k^2 \frac{\int u_t(i) E\left[\frac{x_t(i)}{w_t(i)} \mid \mathcal{F}_t^c\right] di}{E\left[\int u_t(i) \frac{1}{w_t(j)} dj \mid \mathcal{F}_t^c\right]}.$$

Given the optimal position of investing households from Proposition 1.2, and that $w_t(i)$ is independent of $\hat{\theta}(i)$ because of the generational structure of the economy, it follows that

$$-\frac{1}{dt}\frac{d\langle \Lambda_t, E_t \mid \mathcal{F}_t^c\rangle}{E[\Lambda_t \mid \mathcal{F}_t^c]\,E_t} = \frac{a_t - I_t}{q_t} - \frac{I_t}{a - I_t}g_t + I_t\hat{\theta}_t^c - r_t - \delta.$$

Thus by direct integration, the linearity of the expectation and covariance operators, and the Law of Iterated Expectations, it follows that

$$\frac{a - I_t}{E_t}K_t + \frac{1}{dt}E\left[\frac{dE_t}{E_t} \mid \mathcal{F}_t^c\right] + \frac{1}{dt}\frac{E[d\Lambda_t \mid \mathcal{F}_t^c]}{E[\Lambda_t \mid \mathcal{F}_t^c]} + \frac{1}{dt}\frac{d\langle \Lambda_t, E_t \mid \mathcal{F}_t^c\rangle}{E[\Lambda_t \mid \mathcal{F}_t^c]\,E_t} = 0.$$

Therefore, if E_t satisfies each household's Euler equation, then $E_t = q_t K_t$ solves each manager's problem.

Thus from equation (A.8), it follows that

$$g = \rho\left(q\hat{\theta}^c - 1\right)\mathbf{1}\left\{I > \underline{I} \,\cup\, \hat{\theta}^c \geq \frac{\rho}{a - \underline{I}}\right\}.$$

Proof of Proposition 1.5:

By the second part of Proposition 1.3

$$\int_{\mathcal{D}_t^c}\frac{w_t(i)}{W_t}\left(\hat{\theta}_t(i) - \hat{\theta}_t^c\right)di = \frac{\Sigma_t}{\Sigma_t + \sigma_s^2}\left(\theta_t - \hat{\theta}_t^c\right)\int_{\mathcal{D}_t^c}\frac{w_t(i)}{W_t}di + \frac{\Sigma_t}{\Sigma_t + \sigma_s^2}\int_{\mathcal{D}_t^c}\frac{w_t(i)}{W_t}Z_t^s(i)\,di. \tag{A.5}$$

Let me define the integral X_t

$$X_t = \int_0^1 \psi_t(i)\,dZ_t^s(i)\,di.$$

where $\psi_t(i) = \frac{w_t(i)}{W_t} > 0$ is now a weight function, with $\psi_t(i) \in (0,1)$ on a set of full measure, whose integral is bounded on any set of positive measure and is 1 over the set $i \in [0,1]$.

Importantly, since the law of motion of the price of firm equity q and the riskless rate r by conjecture do not depend on the wealth share or signal noise of any one household, the only difference in the wealth shares of households at time t are the histories of the fraction of wealth invested in firm equity $\{x_u(i)\}_{u\leq t}$, which differ across households only because of differences in signal noise. Therefore, conditional on the initial wealth share of households and the history of the fundamentals $\mathcal{G}_t = \sigma\left(\{\theta_u, K_u, \xi_u\}_{u\leq t} \vee w_0\right)$, the weights $\psi_i(t)$ are independent across households.

First, I establish that X_t converges to its cross-sectional expectation $E[X_t \mid \mathcal{G}_t]$ in the $\mathcal{L}^2 - norm$. As an aside, I do not require convergence $a.s.$ and rely on a weaker notion of convergence because of the issues discussed in Judd (1985).

Similar to Uhlig (1996), one can discretize the integral across i into a Riemann sum $\Sigma(t,\varphi)$ with a partition φ with $0 = i_0 < ...i_j < ...i_m = 1$ and midpoints $\phi_j \in [i_{j-1}, i_j]$, $j \in \{1, ..., m\}$

$$\Sigma(t,\varphi) = \sum_{j=1}^{m} \psi_t(\phi_j) Z_t^s(\phi_j)(i_j - i_{j-1}).$$

Conditional on $\mathcal{G}_t$, $E[X_t \mid \mathcal{G}_t]$ is a constant, and one recognizes by Chebychev's Inequality that

$$\begin{aligned}
& E\left[(\Sigma(t,\varphi) - E[X_t \mid \mathcal{G}_t])^2 \mid \mathcal{G}_t\right] \\
= \; & E\left[\left(\sum_{j=1}^{m} \left(\psi_t(\phi_j) Z_t^s(\phi_j) - E[X_t \mid \mathcal{G}_t]\right)(i_j - i_{j-1})\right)^2 \mid \mathcal{G}_t\right] \\
= \; & E\left\{\sum_{j=1}^{m} E\left[\left(\psi_t(\phi_j) Z_t^s(\phi_j) - E[X_t \mid \mathcal{G}_t]\right)^2 \mid \mathcal{G}_t\right](i_j - i_{j-1})^2\right\} \\
\leq \; & \sum_{j=1}^{m} (i_j - i_{j-1})^2 \\
\leq \; & \varepsilon(\varphi),
\end{aligned}$$

where $\varepsilon(\varphi) = \max_j (i_j - i_{j-1})$. As $\varepsilon(\varphi) \searrow 0$, the above integral converges to the $\mathcal{L}^2$ distance between $\Sigma(t,\varphi)$ and $E[X_t \mid \mathcal{G}_t]$ on the LHS and 0 on the RHS.

Therefore

$$\lim_{\varepsilon(\varphi)\searrow 0} E\left[(\Sigma(t,\varphi) - E[X_t \mid \mathcal{G}_t])^2 \mid \mathcal{G}_t\right] = 0.$$

By Dominated Convergence and Slusky's Theorem

$$\lim_{\varepsilon(\varphi)\searrow 0} E\left[(\Sigma(t,\varphi) - E[X_t \mid \mathcal{G}_t])^2 \mid \mathcal{G}_t\right] = E\left[(X_t - E[X_t \mid \mathcal{G}_t])^2 \mid \mathcal{G}_t\right].$$

Therefore

$$E\left[(X_t - E[X_t \mid \mathcal{G}_t])^2 \mid \mathcal{G}_t\right] = 0,$$

which does not depend on the wealth share or signal noise of any individual household because $E[X_t \mid \mathcal{G}_t] = g(\tilde{\omega}_t)$ for some $\tilde{\omega}_t \in \mathcal{G}_t$.

Since the choice of partition φ was arbitrary, the convergence result did not depend on my choice of partition, and therefore X_t and its convergence to $g(\tilde{\omega}_t)$ in $\mathcal{L}^2$ are well-defined. Furthermore, since convergence is in $\mathcal{L}^2$, the integral is $g(\tilde{\omega}_t)$ $a.s.$ and I can choose a modification of the process, if need be, under which it is always 0.[38] Given that this convergence is ex-post the realized sample path of the aggregate state variables $\mathcal{G}_t$, this convergence also holds unconditionally.

Recognizing that $E\left[\bar{Z}(i) \mid \mathcal{G}_t\right] = 0$, it follows that

$$g(\tilde{\omega}_t) = \frac{\Sigma_t}{\Sigma_t + \sigma_s^2} E\left[\psi_t(i) Z_t^s(i) \mid \mathcal{G}_t\right] = \frac{\Sigma_t}{\Sigma_t + \sigma_s^2} E\left[\psi_t(i) \mid \mathcal{G}_t\right] E\left[Z_t^s(i) \mid \mathcal{G}_t\right] = 0,$$

since $\psi_t(i)$ is independent of $Z_t^s(i)\ \forall\, i$ and $E\left[Z_t^s(i) \mid \mathcal{G}_t\right] = 0$. Similarly, I can apply a weak LLN to $\int_0^1 \frac{w_t(i)}{W_t} di$, which holds on subintervals of $[0,1]$ $a.s.$, to arrive at

$$\begin{aligned} W_t &= E\left[w_t(i) \mid \mathcal{G}_t\right], \\ \int_{\mathcal{D}_t^c} \frac{w_t(i)}{W_t} di &= 1 - \pi, \\ \int_{\mathcal{D}_t} \frac{w_t(i)}{W_t} di &= \pi. \end{aligned}$$

Thus equation (A.5) becomes

$$\int_{\mathcal{D}_t^c} \frac{w_t(i)}{W_t} \left(\hat{\theta}_t(i) - \hat{\theta}_t^c\right) di - \int_{\mathcal{D}_t} \frac{w_t(i)}{W_t} \xi_t di = (1-\pi) \frac{\Sigma_t}{\Sigma_t + \sigma_s^2} \left(\theta_t - \hat{\theta}_t^c\right) - \pi \xi_t.$$

Proof of Proposition 1.1:

Substituting $q = \frac{a-I}{\rho}$, optimal household demand for firm claims $x(i)$ from Proposition 1.2, and optimal firm investment g from Proposition 1.4 into the market clearing condition for the market for riskless debt (8), and imposing $W > 0$ and Proposition 1.5, one has, when $I > \underline{I}$, that

$$r = \frac{a}{a-I}\rho - \delta + I \frac{\Sigma}{\Sigma + \sigma_s^2} \left(\theta - \hat{\theta}^c\right) - \frac{1 + \pi \xi}{1 - \pi} \sigma_k^2,$$

and therefore, matching this with the conjectured representation equation (13), it follows that

$$\begin{aligned} r_0 &= \frac{a}{a-I}\rho - \delta - I \frac{\Sigma}{\Sigma + \sigma_s^2} \hat{\theta}^c - \frac{1}{1-\pi} \sigma_k^2, \\ r_\theta &= I \frac{\Sigma}{\Sigma + \sigma_s^2}, \\ r_\xi &= -\frac{\pi}{1-\pi} \sigma_k^2, \end{aligned}$$

[38] Though the convergence implies that the variance of X_t is zero over time, X_t can deviate from its expected value on a negligible subset of times.

which confirms the conjecture. Given optimal firm equity demand $x(i)$ from Proposition 1.2, it follows that $x(i)$ can be decomposed as

$$x(i) = x_c + x_i \left(\hat{\theta}(i) - \hat{\theta}^c\right)$$

where

$$\begin{aligned} x_c &= \frac{\frac{a}{a-I}\rho - r - \delta}{\sigma_k^2}, \\ x_i &= \frac{I}{\sigma_k^2}. \end{aligned}$$

When $I = \underline{I}$ and $g = 0$, then r is instead given by

$$r = \rho - \delta + I\hat{\theta} + I\frac{\Sigma}{\Sigma + \sigma_s^2}\left(\theta - \hat{\theta}^c\right) - \frac{1+\pi\xi}{1-\pi}\sigma_k^2,$$

and x_c is insteady

$$x_c = \frac{\rho + I\hat{\theta} - r}{\sigma_k^2}.$$

Proof of Proposition 1.6:

When θ, then optimal investment I and the firm equity price q are given by equations (12) and (15)

$$q = \frac{a - I}{\rho}$$

and

$$g = \rho(q\theta - 1).$$

Since all households are now perfectly informed, it follows that the only heterogeneity among them is whether they are hit by liquidity shocks. Following the arguments of Proposition 1.2, their optimal policies are

$$\begin{aligned} c(i) &= \rho w(i), \\ x(i) &= \frac{\frac{a-I}{q} - \frac{I}{a-I}g + I\theta - r - \delta}{\sigma_k^2}. \end{aligned}$$

By the market clearing condition for riskless debt (8), it follows that

$$r = \frac{a}{a-I}\rho - \delta - \frac{1+\pi\xi}{1-\pi}\sigma_k^2.$$

Proof of Proposition 1.7:

When households are perfectly informed about θ, they consume a fixed fraction of their wealth and follow identical investment strategies

$$\begin{aligned} c(i) &= \rho w(i), \\ x(i) &= \frac{\frac{a-I}{q} - \frac{I}{a-I} g + I\theta - r - \delta}{\sigma_k^2}, \end{aligned}$$

when not hit by the preference shock. Since managers still learn from prices, it follows that the optimal g still satisfies

$$g = \rho \left(q\hat{\theta}^c - 1 \right).$$

It follows by market clearing condition for riskless debt (8) that the riskless rate satisfies

$$r = \frac{a}{a-I}\rho - \delta + I\left(\theta - \hat{\theta}^c\right) - \frac{1+\pi\xi}{1-\pi}\sigma_k^2.$$

As $\sigma_s \searrow 0$, from the law of motion of $\hat{\theta}^c$ and $\hat{\theta}(i)$ from Proposition 1.3, it follows that $\Sigma(i) \searrow 0$ while

$$\frac{d\Sigma_t}{dt} \to \sigma_\theta^2 - 2\lambda\Sigma - I^2\frac{\Sigma^2}{\sigma_k^2} - \frac{\left(\alpha\sigma_\xi\sigma_\theta + R_\theta\sigma_\theta^2 + R_\theta(g-\lambda)\Sigma_t\right)^2}{\left(R_\theta\sigma_\theta + \alpha\sigma_\xi\right)^2 + (1-\alpha^2)\sigma_\xi^2}$$

where R_θ, the loading of the riskless rate r on the household expectational error $\theta - \hat{\theta}^c$ converges to

$$R_\theta \to -\frac{1-\pi}{\pi}\frac{I}{\sigma_k^2}.$$

Thus it follows that Σ does not converge to 0 as $\sigma_s \searrow 0$, reflecting the uncertainty that firm managers still face about θ from observing only $\log K$ and r, and g does not converge to its perfect-information benchmark value.

Since $\hat{\theta}^c \to \theta$, $\hat{\theta}(i) \to \theta$, it follows that the investment strategy of households $x(i)$ converges to

$$x(i) = \frac{\frac{a-I}{q} - \frac{I}{a-I} g + I\theta - r - \delta}{\sigma_k^2},$$

from which it follows that the riskless rate r approaches its representative agent benchmark value. Thus beliefs, prices, and optimal policies in the economy with informational frictions approach their representative agent benchmark values as $\sigma_s \searrow 0$.

Proof of Proposition 1.8:

From Proposition 1.1, it follows that each household's demand for the risky asset when not hit by the liquidity shock can be rewritten as

$$x(i) = \frac{1+\pi\xi}{1-\pi} + \frac{I}{\sigma_k^2}\frac{\Sigma_t}{\Sigma_t+\sigma_s^2}\sigma_s Z^s(i). \tag{A.11}$$

Substituting my expressions for q and g into the law of motion of household wealth $w_t(i)$ equation (9), it follows by Itô's Lemma that

$$d\log w(i) = (1-x(i))(r-\rho)\,dt + x(i)\left(\left(\frac{\rho I}{a-I} - \delta + I\left(\theta - \hat{\theta}^c\right)\right)dt + \sigma_k dZ^k\right) - \frac{1}{2}x(i)^2\sigma_k^2 dt,$$

Substituting for $x(i)$ with equation (A.11) and aggregate across households, one then has that

$$\begin{aligned}\int_0^1 d\log w(i)\,di &= \left(\frac{\rho I}{a-I} - \delta + I\left(\theta - \hat{\theta}^c\right)\right)dt + \sigma_k dZ^k \\ &\quad -\frac{1}{2}\left(\pi\sigma_k^2\xi^2 + \sigma_k^2\frac{(1+\pi\xi)^2}{1-\pi} + (1-\pi)\left(\frac{I}{\sigma_k}\frac{\Sigma}{\Sigma+\sigma_s^2}\sigma_s\right)^2\right)dt \end{aligned}\tag{A.12}$$

With equation (A.12), one can then express aggregate flow utility $\int_0^1 \log c_s(i)\,di = \log\rho + \log w_0 + \int_0^1\int_0^s d\log w_u(i)\,du\,di$ as

$$\begin{aligned}\int_0^1 \log c_t(i)\,di &= \int_0^t\left(\frac{\rho I_s}{a-I_s} - \delta + I_s\left(\theta_s - \hat{\theta}_s^c\right)\right)ds + \sigma_k Z_t^k + \log\rho + \log w_0 \\ &\quad -\frac{1}{2}\int_0^t\left(\sigma_k^2 + \frac{\pi\sigma_k^2}{1-\pi}(1+\xi_s)^2 + (1-\pi)\left(\frac{I_s}{\sigma_k}\frac{\Sigma_s}{\Sigma_s+\sigma_s^2}\sigma_s\right)^2\right)ds.\end{aligned}$$

It follows then that utilitarian welfare at time 0 $U = E\left[\int_0^\infty e^{-\rho t}\int_0^1 \log\frac{c_t(i)}{c_0(i)}\,di\,dt \mid \mathcal{F}_0\right]$ in the economy under the physical measure $\mathcal{P}$ defined on $\mathcal{F}_0$ is given by

$$\begin{aligned}U &= E\left[\int_0^\infty e^{-\rho t}\left[\int_0^t\left(\frac{\rho I_s}{a-I_s} - \delta + I_s\left(\theta_s - \hat{\theta}_s^c\right)\right)ds\right]dt \mid \mathcal{F}_0\right] + E\left[\int_0^\infty e^{-\rho t}\sigma_k Z_t^k dt \mid \mathcal{F}_0\right] \\ &\quad -\frac{1}{2}E\left[\int_0^\infty e^{-\rho t}\left[\int_0^t\left(\sigma_k^2 + \frac{\pi\sigma_k^2}{1-\pi}(1+\xi_s)^2 + (1-\pi)\left(\frac{I_s}{\sigma_k}\frac{\Sigma_s}{\Sigma_s+\sigma_s^2}\sigma_s\right)^2\right)ds\right]dt \mid \mathcal{F}_0\right].\end{aligned}$$

Taking expectations under $\mathcal{P}$, it follows that

$$\begin{aligned}U &= E\left[\int_0^\infty e^{-\rho t}\left[\int_0^t\left(\frac{\rho I_s}{a-I_s} - \delta + I_s\left(\theta_s - \hat{\theta}_s^c\right) - \frac{1-\pi}{2}\left(\frac{I_s}{\sigma_k}\frac{\Sigma_s}{\Sigma_s+\sigma_s^2}\sigma_s\right)^2\right)ds\right]dt \mid \mathcal{F}_0\right] \\ &\quad -\frac{1}{2}\left(\sigma_k^2 + \frac{\pi\sigma_k^2}{1-\pi}(1+\xi_0)^2\right)E\left[\int_0^\infty e^{-\rho(s-t)}t\,dt \mid \mathcal{F}_0\right] \\ &\quad -\frac{1}{4}\frac{\pi\sigma_k^2}{1-\pi}\sigma_\xi^2 E\left[\int_0^\infty e^{-\rho(s-t)}t^2 dt \mid \mathcal{F}_0\right].\end{aligned}$$

Recognizing that $\int_0^\infty e^{-\rho\tau}\tau d\tau = \frac{1}{\rho^2}$ and $\int_0^\infty e^{-\rho\tau}\tau^2 d\tau = \frac{2}{\rho^3}$, one arrives at

$$\begin{aligned} U &= E\left[\int_0^\infty e^{-\rho t}\int_0^t \left(\frac{\rho}{a-I_s} + \theta_s - \hat{\theta}^c_s\right) I_s ds dt \mid \mathcal{F}_0\right] - \frac{1}{2\rho^2}\left(\sigma_k^2 + \frac{\pi\sigma_k^2}{1-\pi}\left(1+\xi_0\right)^2\right) \\ &\quad -\frac{1-\pi}{2}\left(\frac{\sigma_s}{\sigma_k}\right)^2 E\left[\int_0^\infty e^{-\rho t}\int_0^t \left(\frac{I_s\Sigma_s}{\Sigma_s+\sigma_s^2}\right)^2 ds dt \mid \mathcal{F}_0\right] \\ &\quad -\frac{\delta}{\rho^2} - \frac{1}{2\rho^3}\frac{\pi\sigma_k^2}{1-\pi}\sigma_\xi^2. \end{aligned} \tag{A.13}$$

By stacking the terms in the two double integrals in equation (A.13), I can rewrite them to arrive at

$$\begin{aligned} U &= \frac{1}{\rho}E\left[\int_0^\infty e^{-\rho t}\left(\frac{\rho}{a-I_t} + \theta_t - \hat{\theta}^c_t\right) I_t dt \mid \mathcal{F}_0\right] \\ &\quad -\frac{1-\pi}{2\rho}\left(\frac{\sigma_s}{\sigma_k}\right)^2 E\left[\int_0^\infty e^{-\rho t}\left(\frac{I_t\Sigma_t}{\Sigma_t+\sigma_s^2}\right)^2 dt \mid \mathcal{F}_0\right] \\ &\quad -\frac{\delta}{\rho^2} - \frac{1}{2\rho^2}\left(\sigma_k^2 + \frac{\pi\sigma_k^2}{1-\pi}\left(1+\xi_0\right)^2\right) - \frac{1}{2\rho^3}\frac{\pi\sigma_k^2}{1-\pi}\sigma_\xi^2. \end{aligned}$$

Defining w such that $d\log w = \int_0^1 d\log w(i)\, di$, from equation (A.12) and Itô's Lemma it follows that w has the law of motion

$$\frac{dw}{w} = \left(\frac{\rho I}{a-I} - \delta + I\left(\theta - \hat{\theta}^c\right) - \frac{1}{2}\left(\frac{\pi\sigma_k^2}{1-\pi}\left(1+\xi\right)^2 + (1-\pi)\left(\frac{I}{\sigma_k}\frac{\Sigma}{\Sigma+\sigma_s^2}\sigma_s\right)^2\right)\right) dt + \sigma_k dZ^k. \tag{A.14}$$

Thus I can think of the economy as having a representative household who holds all firm claims in the economy and whose wealth evolves according to the law of motion (A.14).

Proof of Proposition 1.9:

To find the law of motion of the probability law of the economy $p_t\left(\hat{\theta}^c, \Sigma, I\right)$, I find the probability law implied by households and firms whose optimization is consistent with their HJB equations. This is commonly referred to as the Kolmogorov Forward Equation. To find this, I recognize that, under the optimal control for the change in investment $g\left(\hat{\theta}^c_s, \Sigma_s, I_s\right)_{s\geq 0}$, $\mathcal{D}^g f = 0$ where $\mathcal{D}^g$ is the infinitesimal generator that satisfies

$$\begin{aligned} \mathcal{D}^g f &= \partial_{\hat{\theta}^c} f \lambda\left(\bar{\theta} - \hat{\theta}^c\right) + \partial_\Sigma f \frac{d\Sigma}{dt} + \partial_I f I\left((a-I)\hat{\theta}^c - \left(1-\tau^I\right)\rho\right)\mathbf{1}_{\left\{I>\underline{I}\cup\hat{\theta}^c \geq \frac{(1-\tau^I)\rho}{a-\underline{I}}\right\}} \\ &\quad + \frac{1}{2}\partial_{\hat{\theta}^c\hat{\theta}^c} f\left(\sigma^2_{\hat{\theta}k} + \sigma^2_{\hat{\theta}r}\right), \end{aligned}$$

where $\sigma_{\hat{\theta}k}$ and $\sigma_{\hat{\theta}r}$ are given in Proposition 1.3 appropriately modified for the position cost τ^r.

Let $z\left(\hat{\theta}^c, \Sigma, I\right) \in \mathcal{C}_0^\infty\left(\mathbb{R} \times \left[0, \frac{\sigma_\theta^2}{2\lambda}\right] \times [\underline{I}, a]\right)$ be an arbitrarily, infinitely differentiable test function with compact support. Then $E\left[z\left(\hat{\theta}_t^c, \Sigma_t, I_t\right)\right] = \int z\left(\hat{\theta}^c, \Sigma, I\right) p_t\left(\hat{\theta}^c, \Sigma, I\right) d\hat{\theta}^c d\Sigma dI$ can be written as

$$\begin{aligned} E\left[z\left(\hat{\theta}_t^c, \Sigma_t, I_t\right)\right] &= E\left[\int_0^t dz\left(\hat{\theta}_s^c, \Sigma_s, I_s\right)\right] \\ &= E\left[\int_0^t \mathcal{D}^g z\left(\hat{\theta}_s^c, \Sigma_s, I_s\right) ds\right] \\ &= \int \int_0^t \mathcal{D}^g z\left(\hat{\theta}^c, \Sigma, I\right) p_t\left(\hat{\theta}^c, \Sigma, I\right) d\hat{\theta} d\Sigma dI. \end{aligned}$$

Differentiating w.r.t t, one finds that

$$\int z\left(\hat{\theta}^c, \Sigma, I\right) \partial_t p_t\left(\hat{\theta}^c, \Sigma, I\right) d\hat{\theta} d\Sigma dI = \int \mathcal{D}^g z\left(\hat{\theta}^c, \Sigma, I\right) p_t\left(\hat{\theta}^c, \Sigma, I\right) d\hat{\theta}^c d\Sigma dI.$$

Since z has compact support, I can perform integration by parts to arrive at

$$\int z\left(\hat{\theta}^c, \Sigma, I\right) \partial_t p_t\left(\hat{\theta}^c, \Sigma, I\right) d\hat{\theta}^c d\Sigma dI = \int z\left(\hat{\theta}^c, \Sigma, I\right) \mathcal{D}^{g*} p_t\left(\hat{\theta}^c, \Sigma, I\right) d\hat{\theta}^c d\Sigma dI,$$

where $\mathcal{D}^{g*}$ is the adjoint of $\mathcal{D}^g$ and is the time-homogeneous infinitesimal generator associated with the Koopman operator. Assuming $\partial_t p_t\left(\hat{\theta}^c, \Sigma, I\right) - \mathcal{D}^{g*} p_t\left(\hat{\theta}^c, \Sigma, I\right)$ is continuous, it follows, since z is arbitrary, that

$$\partial_t p_t\left(\hat{\theta}^c, \Sigma, I\right) = \mathcal{D}^{g*} p_t\left(\hat{\theta}^c, \Sigma, I\right), \tag{A.9}$$

Importantly, $\mathcal{D}^{g*}$ is a (uniformly) elliptic operator that has divergence form. When p_t has reached its stationary distribution p, where $p = \lim_{t \nearrow \infty} p_t$, it follows that $\partial_t p_t = 0$. Thus equation (A.9) is a second-order parabolic equation and can can be rewritten when p_t has reached its stationary distribution, suppressing arguments, as

$$\begin{aligned} 0 &= -\partial_{\hat{\theta}^c}\left\{p\lambda\left(\bar{\theta} - \hat{\theta}^c\right)\right\} - \partial_I\left\{pI\left((a - I)\hat{\theta}^c - \left(1 - \tau^I\right)\rho\right)\right\} \mathbf{1}_{\left\{I > \underline{I} \cup \hat{\theta}^c \geq \frac{(1-\tau^I)\rho}{a - \underline{I}}\right\}} - \partial_\Sigma\left\{p\frac{d\Sigma}{dt}\right\} \\ &\quad + \frac{1}{2}\partial_{\hat{\theta}^c\hat{\theta}^c}\left\{p\left(\sigma_{\hat{\theta}k}^2 + \sigma_{\hat{\theta}r}^2\right)\right\}, \end{aligned} \tag{A.10}$$

which is the expression given in the proposition.

That $p_t\left(\hat{\theta}^c, \Sigma, I\right)$ will satisfy the conservation of mass law $\int p_t\left(\hat{\theta}^c, \Sigma, I\right) d\hat{\theta}^c d\Sigma dI = 1$, where the integral is understood to be taken over the entire space $\mathbb{R} \times \left[0, \frac{\sigma_\theta^2}{2\lambda}\right] \times [\underline{I}, a]$, gives rise to my spatial boundary conditions. Notice that I can rewrite equation (A.10) as

$$\nabla \cdot S\left(\hat{\theta}^c, \Sigma, I\right) = 0,$$

where

$$S\left(\hat{\theta}^c, \Sigma, I\right) = \begin{bmatrix} S^{\hat{\theta}^c}\left(\hat{\theta}^c, \Sigma, I\right) \\ S^{\Sigma}\left(\hat{\theta}^c, \Sigma, I\right) \\ S^{I}\left(\hat{\theta}^c, \Sigma, I\right) \end{bmatrix} = \begin{bmatrix} \lambda\left(\bar{\theta} - \hat{\theta}^c\right) p\left(\hat{\theta}^c, \Sigma, I\right) - \frac{1}{2}\partial_{\hat{\theta}^c}\left\{\left(\sigma^2_{\hat{\theta}k} + \sigma^2_{\hat{\theta}r}\right) p\left(\hat{\theta}^c, \Sigma, I\right)\right\} \\ \frac{d\Sigma}{dt} p\left(\hat{\theta}^c, \Sigma, I\right) \\ \left((a - I)\hat{\theta}^c - \left(1 - \tau^I\right)\rho\right) I p\left(\hat{\theta}^c, \Sigma, I\right) \mathbf{1}_{\left\{I > \underline{I} \cup \hat{\theta}^c \geq \frac{(1-\tau^I)\rho}{a - \underline{I}}\right\}} \end{bmatrix}.$$

Here $S\left(\hat{\theta}^c, \Sigma, I\right)$ is the "probability flux" representing the flow or flux of particles through the point $\left(\hat{\theta}^c, \Sigma, I\right)$. Consequently, a reflecting boundary condition will ultimately impose that the flux through boundary points must be zero.

For $\hat{\theta}^c$, which have unbounded support, one has that for $\varepsilon > 0$ arbitrary that

$$\lim_{\hat{\theta}^c \nearrow \infty} \left(\hat{\theta}^c\right)^{2(1+\varepsilon)} p\left(\hat{\theta}^c, \Sigma, I\right) = 0 \;\forall\; I,$$

while for $\Sigma = 0$, one has that $\partial_\Sigma p\left(\hat{\theta}^c, \frac{\sigma^2_\theta}{2\lambda}, I\right) = 0$, since $\frac{\sigma^2_\theta}{2\lambda}$ is a reflecting boundary, and $\lim_{\Sigma \searrow 0} p\left(\hat{\theta}^c, \Sigma, I\right) = 0$, since arbitrarily small precision becomes arbitrarily unlikely given that new unobservable innovations to θ_t occur at each instant.

Integrating this expression over the entire space, imposing that $\int \partial_t p_t\left(\hat{\theta}^c, \Sigma, I\right) d\hat{\theta}^c d\Sigma dI = \partial_t \int p_t\left(\hat{\theta}^c, \Sigma, I\right) d\hat{\theta}^c d\Sigma dI = 0$, applying the Divergence Theorem, it follows that the appropriate "reflecting" boundary condition for I is $\hat{n}_{I=\underline{I}} \cdot S\left(\hat{\theta}^c, \Sigma, \underline{I}\right) = \hat{n}_{I=a} \cdot S\left(\hat{\theta}^c, \Sigma, a\right) = 0 \;\forall\; \left(\hat{\theta}^c, \Sigma\right)$, where $\hat{n}_{I=i}$ is the unit (outward) normal vector perpendicular to the $I = i$ boundary. The intuition for these two boundary conditions is that the probability flux, or flow, through the two walls $I = \underline{I}$ and $I = a$ must be zero for probability mass not to leak out through them.

Proof of Corollary 1.1:

Let U^c be ex-ante utilitarian welfare under the common knowledge filtration. Then U^c satisfies

$$\begin{aligned} U^c &= E\left[\int_0^\infty e^{-\rho t} \int_0^1 \log c_t(i)\, di dt \mid \mathcal{F}_0^c\right] \\ &= E\left[E\left[\int_0^\infty e^{-\rho t} \int_0^1 \log c_t(i)\, di dt \mid \mathcal{F}_0\right] \mid \mathcal{F}_0^c\right] \\ &= E\left[U \mid \mathcal{F}_0^c\right], \end{aligned}$$

from which follows from the expression for U from 1.9, and since $\hat{\theta}_t^c \sim \mathcal{N}(0, \Sigma)$, that the above reduces by the LIE to

$$\begin{aligned} U^c &= E\left[\int_0^\infty e^{-\rho t}\frac{I_t}{a-I_t}dt \mid \mathcal{F}_0^c\right] - \frac{1-\pi}{2\rho}\left(\frac{\sigma_s}{\sigma_k}\right)^2 E\left[\int_0^\infty e^{-\rho t}\left(\frac{I_t}{1-\tau^r}\frac{\Sigma_t}{\Sigma_t+\sigma_s^2}\right)^2 dt \mid \mathcal{F}_0^c\right] \\ &\quad -\frac{1}{2\rho^2}\frac{1-\pi}{\pi\sigma_k^2}\left(\frac{I_0}{1-\tau^r}\frac{\Sigma_0}{\Sigma_0+\sigma_s^2}\right)^2\Sigma_0 - \frac{1}{2\rho^2}\left(\sigma_k^2 + \frac{\pi\sigma_k^2}{1-\pi}\left(1+E\left[\xi_0 \mid \mathcal{F}_0^c\right]\right)^2\right) \\ &\quad -\frac{\delta}{\rho^2} - \frac{1}{2\rho^3}\frac{\pi\sigma_k^2}{1-\pi}\sigma_\xi^2, \end{aligned} \tag{A.12}$$

since $\left(r, \hat{\theta}^c, \hat{\xi}^c\right) \in \mathcal{F}^c \subseteq \mathcal{F}$, from the expression for the riskless rate r in Proposition 1.1, and it follows

$$\xi - \hat{\xi}^c = \frac{1-\pi}{\pi\sigma_k^2}\frac{I}{1-\tau^r}\frac{\Sigma}{\Sigma+\sigma_s^2}\left(\theta - \hat{\theta}^c\right),$$

and therefore

$$E\left[\left(\xi_0 - E\left[\xi_0 \mid \mathcal{F}_0^c\right]\right)^2 \mid \mathcal{F}_0^c\right] = \left(\frac{1-\pi}{\pi\sigma_k^2}\frac{I_0}{1-\tau^r}\frac{\Sigma_0}{\Sigma_0+\sigma_s^2}\right)^2\Sigma_0,$$

Assume now that the economy is initialized from the stationary distribution $p\left(\hat{\theta}^c, I, \Sigma\right)$ and that the stationary distribution is bounded $p\left(\hat{\theta}^c, I, \Sigma\right) \in \mathcal{L}^\infty\left(\mathbb{R}, \left[0, \frac{\sigma_\theta^2}{2\lambda}\right], \left[\underline{I}, a\right]\right)$. Let U_p^c be the expected welfare in economy under the stationary distribution, and $E^p[\cdot]$ be the expectation operator w.r.t. the stationary distribution. Then the first expectation, when taken w.r.t. the stationary distribution, can be rewritten as

$$\begin{aligned} E^p\left[\int_0^\infty e^{-\rho t}\frac{I_t}{a-I_t}dt\right] &= \int_0^\infty e^{-\rho t}\int P_t\frac{I_0}{a-I_0}p\left(\hat{\theta}_0^c, \Sigma_0, I_0\right)d\hat{\theta}d\Sigma dI dt \\ &= \int_0^\infty e^{-\rho t}\int \frac{I_0}{a-I_0}P_t^* p\left(\hat{\theta}_0^c, \Sigma_0, I_0\right)d\hat{\theta}d\Sigma dI dt, \end{aligned} \tag{A.13}$$

where $P_t = e^{t\mathcal{D}^g}$ is the Ruelle-Frobenius-Perron operator and $P_t^* = e^{t\mathcal{D}^{g*}}$ is its adjoint, often called the Koopman operator. P_t^* is defined such that, for a bounded, Borel measurable function f and measure ν $\langle\ P_t f,\ \ \nu\ \rangle = \int_{\mathbb{R}\times\mathbb{R}_+\times\mathbb{R}_+} P_t f d\nu = \langle\ f,\ \ P_t^*\nu\ \rangle$. Probabilistically, P_t^* corresponds to time-reversal and acts on measures, whereas P_t acts on functions. By construction, since $\mathcal{D}^{g*}p = 0$,

$$e^{t\mathcal{D}^{g*}}p\left(\hat{\theta}_0^c, \Sigma_0, I_0\right) = p\left(\hat{\theta}_0^c, \Sigma_0, I_0\right),$$

and therefore equation (A.13) simplies to

$$E^p\left[\int_0^\infty e^{-\rho t}\frac{\rho}{a-I_t}I_t dt\right] = \frac{1}{\rho}E^p\left[\frac{I_0}{a-I_0}\right].$$

A similar result obtains for the second expectation, under the assumption that Σ is essentially bounded. Since $\Sigma \leq \frac{\sigma_\theta^2}{2\lambda}$ from Proposition 1.3, this assumption is justified for σ_θ finite and $\lambda > 0$. It follows from these results, $E^p[\xi_0] = 0$, and equation (A.12), that U_p^c takes the form

$$\begin{aligned} U_p^c &= \frac{1}{\rho} E^p \left[\frac{I_0}{a - I_0} \right] - \frac{1-\pi}{2\rho^2} \left(\frac{\sigma_s}{\sigma_k} \right)^2 E^p \left[\left(\frac{I_0}{1-\tau^r} \frac{\Sigma_0}{\Sigma_0 + \sigma_s^2} \right)^2 \right] \\ &\quad - \frac{1}{2\rho^2} \frac{1-\pi}{\pi \sigma_k^2} E^p \left[\left(\frac{I_0}{1-\tau^r} \frac{\Sigma_0}{\Sigma_0 + \sigma_s^2} \right)^2 \Sigma_0 \right] \\ &\quad - \frac{\delta}{\rho^2} - \frac{1}{2\rho^2} \frac{\sigma_k^2}{1-\pi} \left(1 + \frac{1}{\rho} \pi \sigma_\xi^2 \right). \end{aligned}$$

1.C: Figures and Tables

Figure 1.1: Structure of the Model

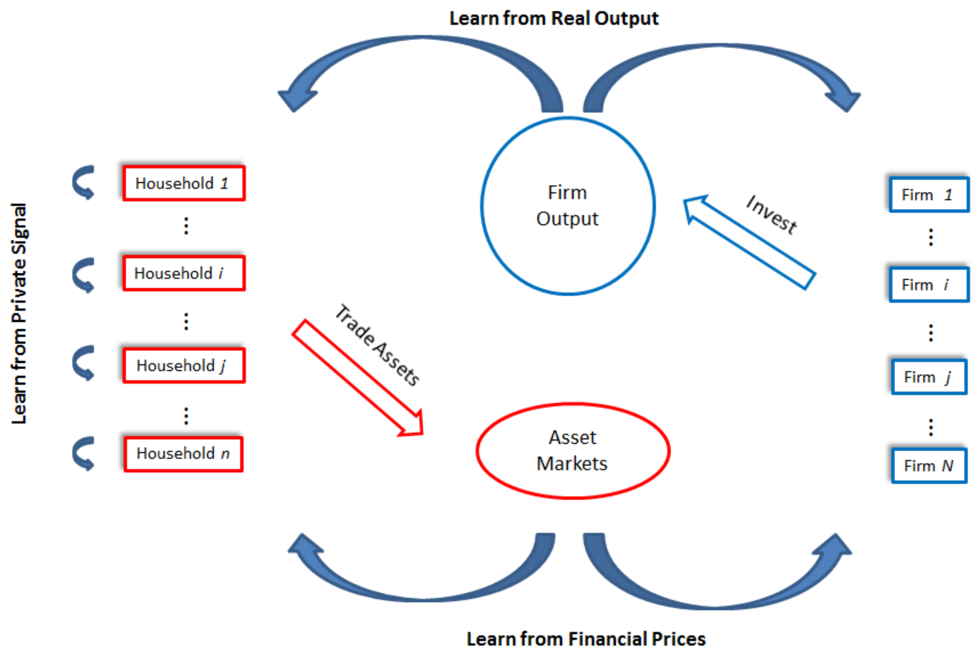

In the numerical experiments that follow, I treat one time unit (t.u) as a year. I set the subjective discount rate ρ to be .02 and depreciation δ to be .10. I choose a to be .2 so that the maximum level of investment I in my model is three standard deviations above its mean of the ratio of US private nonresidential fixed investment to real GDP since 1973. Given the stylized structure of my model, I choose reasonable values for the remaining parameters.

I set the mean-reversion and standard deviation of investment productivity shocks, λ and σ_θ, respectively, to both be .02. I set the long-run mean $\bar{\theta}$ to be .3. I set the standard deviations of capital and financial shocks to be the same $\sigma_k = \sigma_\xi = .05$ so that the exogenous noise in both the real and financial signals are the same. I set the standard deviation of private information σ_s to .03 and the fraction of households hit by the preference shock π to .4. Finally, to shut off any mechanical learning from market prices, I set the correlation between investment productivity and financial shocks α to zero.

Figure 1.2: Loading on Market Signal for Fixed Perceived Investment Productivity $\hat{\theta}^c = 0.3$

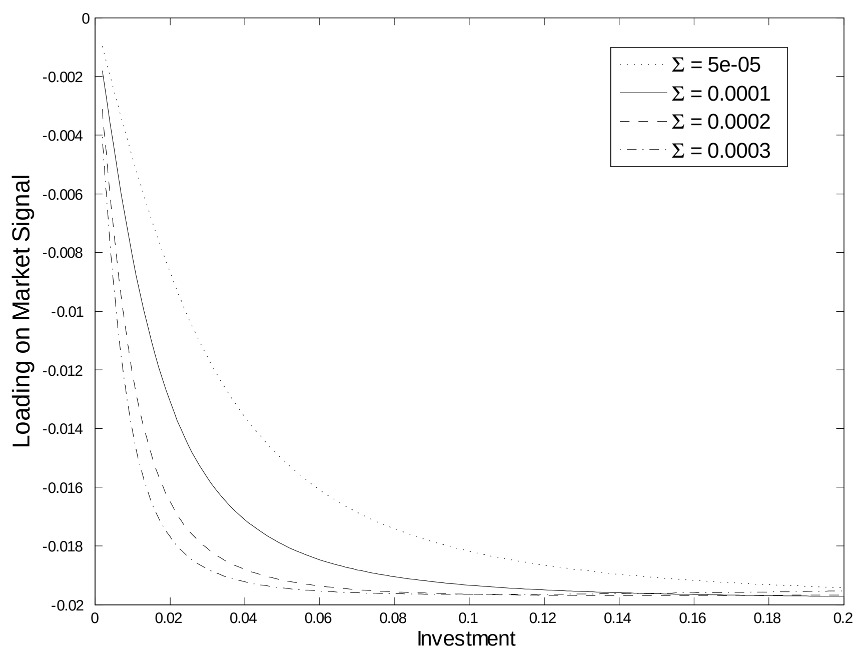

Figure 1.3: Impulse Response of the Economy to a One Standard Deviation Negative Financial Shock in a Boom (Panel 1) and a Bust (Panel 2) (Output is normalized to 1 at time 0).

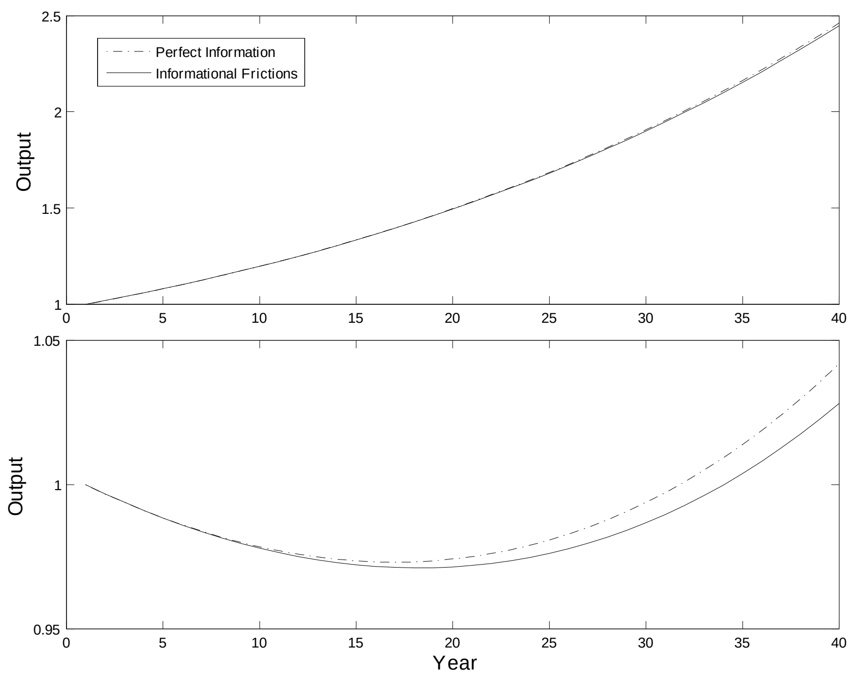

Chapter 2

Informational Frictions and Commodity Markets[39]

2.1. Introduction

In the aftermath of the dramatic boom and bust cycle of commodity prices in 2007 to 2008, there has been renewed interest among academics and policy makers regarding the drivers of commodity price fluctuations. In particular, whether fundamental demand and supply shocks are sufficient to explain the observed price cycles and whether speculation in commodity futures markets exacerbated these cycles are subjects of debate. In this debate, it is common for academic and policy studies to treat different types of shocks (such as supply, demand, and financial market shocks) as observable to market participants.[40] In doing so, however, these studies ignore a key aspect of commodity markets, namely, severe informational frictions faced by market participants. The markets for major commodities, such as crude oil and copper, have become globalized in recent decades, with supply and demand now stemming from across the world. This globalization exposes market participants to heightened informational frictions regarding the global supply, demand, and inventory of these commodities.

The economics literature has developed an elegant theoretical framework to analyze how trading in centralized asset markets facilitates both information aggregation among market participants and helps them overcome the informational frictions they face (e.g., Grossman and Stiglitz (1980) and Hellwig (1980)). This framework, however, crucially relies on the combination of constant absolute risk aversion (CARA) utility functions for agents and Gaussian distributions for asset prices to ensure a tractable linear equilibrium, and thus one cannot readily adopt this framework to analyze commodity markets, in which both CARA utility and Gaussian distributions are unrealistic. It is challenging to analyze information

[39] A version of this chapter, which is joint work with Wei Xiong at Princeton University, is forthcoming the Journal of Finance.

[40] See a recent review by Cheng and Xiong (2014).

aggregation in settings without the tractable linear equilibrium. This technical challenge is common in analyzing how asset prices affect real activity, such as firm investment and central bank policies, through an informational channel.[41]

In this paper, we aim to confront this challenge by developing a tractable model to analyze how informational frictions affect commodity markets. Our model integrates the standard framework of asset market trading with asymmetric information into an international macro setting (e.g., Obstfeld and Rogoff (1996) and Angeletos and La'O (2013)). In this global economy, a continuum of specialized goods producers whose production has complementarity—which emerges from their need to trade produced goods with each other—demand a key commodity, such as copper, as a common production input. Through trading the commodity, the goods producers aggregate dispersed information regarding unobservable global economic strength, which ultimately determines their commodity demand.

Our main model focuses on a centralized spot market through which the goods producers acquire the commodity from a group of suppliers, who are subject to an unobservable supply shock. The supply shock prevents the commodity price from perfectly aggregating the goods producers' information with respect to the strength of the global economy. Nevertheless, the commodity price provides a useful signal to guide the producers' production decisions and commodity demand. Despite the nonlinearity in the producers' production decisions, we derive a unique log-linear equilibrium in closed form. In this equilibrium, each producer's commodity demand is a log-linear function of its private signal and the commodity price, while the commodity price is a log-linear function of global economic strength and the supply shock. This tractable log-linear equilibrium builds on a combination of Cobb-Douglas utility functions for households, log-normal distributions for commodity prices, and a key aggregation property: the aggregate demand of a continuum of producers remains log-linear as a result of the Law of Large Numbers. We also extend the model to incorporate a futures market to further characterize the role of futures market trading.

It is common for empirical studies of commodity markets to rely on conventional wisdom generated from settings without any informational frictions (i.e., agents directly observing both supply and demand shocks). According to such wisdom, 1) a higher price leads to lower commodity demand as a result of the standard cost effect, 2) a positive supply shock reduces the commodity price, which in turn stimulates greater commodity demand, and 3)

[41]See a recent review by Bond, Edmans, and Goldstein (2012).

the futures price of the commodity simply tracks the spot price, and trading in the futures market does not affect either commodity demand or the spot price.

Our model allows us to contrast the effects of informational frictions with this conventional wisdom. First, through its informational role, a higher commodity price signals a stronger global economy and motivates each goods producer to produce more goods. This leads to greater demand for the commodity as an input, which offsets the usual cost effect. The complementarity in production among goods producers magnifies this informational effect through their incentives to coordinate production decisions. Under certain conditions, our model shows that the informational effect can dominate the cost effect and lead to a positive price elasticity of producers' demand for the commodity.

Second, our model illustrates a feedback effect of supply shocks. In the presence of informational frictions, supply shocks also act as informational noise, which prevents the commodity price from fully revealing the strength of the global economy. As goods producers partially attribute the lower commodity price caused by a positive supply shock to a weak global economy, this inference induces them to reduce their commodity demand. This feedback effect thus further amplifies the negative price impact of the supply shock and undermines its impact on commodity demand.

Third, futures markets serve as a useful platform, in addition to spot markets, for aggregating information regarding demand and supply of commodities. As futures markets attract a different group of participants from spot markets, the futures price is not simply a shadow of the spot price, and instead may have its own informational effects on commodity demand and the spot price.

Based on these results, our analysis offers important implications for the empirical analysis of commodity markets. In estimating the effects of supply and demand shocks in commodity markets, it is common for the empirical literature to adopt structural models that ignore informational frictions by simply assuming that agents can directly observe both demand and supply shocks. As highlighted by our analysis, this common practice is likely to understate the effect of supply shocks and overstate the effect of demand shocks. Our model provides the basic ingredients for expanding these structural models to account for how commodity prices impact agents' expectations.

Our analysis also cautions against a commonly used empirical strategy based on commodity inventory to detect speculative effects (e.g., Kilian and Murphy (2014), Juvenal and

Petrella (2012), and Knittel and Pindyck (2013)). This strategy is premised on the widely held argument that if speculators distort the price of a commodity upward, consumers will find the commodity too expensive and thus reduce consumption, causing inventory of the commodity to spike. By assuming that consumers are able to recognize the commodity price distortion, this argument again ignores realistic informational frictions faced by consumers, which are particularly relevant in times of great economic uncertainty. In contrast, our model shows that informational frictions may cause consumers to react to the distorted price by increasing rather than decreasing their consumption. In this light, the lack of any pronounced oil inventory spike before the peak of oil prices in July 2008, as highlighted by the recent empirical literature, cannot be taken as evidence to reject the presence of any speculative effect during the period.

Finally, by systematically illustrating that prices of key industrial commodities can serve as price signals for the strength of the global economy and that informational noise in commodity prices can feed back to commodity demand and spot prices,[42] our analysis provides a coherent argument for how the large inflow of investment capital to commodity futures markets during the 2000s might have amplified the boom and bust of commodity prices in 2007 to 2008. By interfering with the price signals, informational noise from the investment flow may have temporarily led market participants to increase their commodity demand despite a weakening global economy. This confusion helped sustain the commodity price boom until information arrived later to correct their expectations, which then caused commodity prices to collapse.

Our paper contributes to the emerging literature that analyzes the causes of the commodity market cycle of the 2000s, for example Hamilton (2009), Stoll and Whaley (2010), Tang and Xiong (2012), Singleton (2014), Cheng, Kirilenko, and Xiong (2012), Hamilton and Wu (2012), Kilian and Murphy (2014), and Henderson, Pearson, and Wang (2012). The mechanism illustrated by our model echoes Singleton (2014), who emphasizes the importance of accounting for agents' expectations to explain this commodity market cycle. In particular, our analysis highlights the weakness common in empirical studies on the effects of supply

[42] Consistent with this notion, in explaining the decision of the European Central Bank (ECB) to raise its key interest rate in March 2008 on the eve of the worst economic recession since the Great Depression, ECB policy reports cite high prices of oil and other commodities as a key factor, suggesting the significant influence of commodity prices on the expectation of central bankers. Furthermore, Hu and Xiong (2013) provide evidence that in recent years, stock prices across East Asian economies have displayed significant and positive reactions to overnight futures price changes of a set of commodities traded in the U.S., suggesting that people across the world regard commodity futures prices as barometers of the global economy.

and demand shocks and speculation in commodity markets of assuming that different types of shocks are publicly observable to market participants.

Our model complements the recent macro literature that analyzes the role of informational frictions on economic growth. Lorenzoni (2009) shows that by influencing agents' expectations, noise in public news can generate sizable aggregate volatility. Angeletos and La'O (2013) focus on endogenous economic fluctuations that result from the lack of centralized communication channels to coordinate the expectations of different households. Our model adopts the setting of Angeletos and La'O (2013) to model goods market equilibrium and derive endogenous complementarity in goods producers' production decisions. We analyze information aggregation through centralized commodity trading, which is absent from their model, and the feedback effects of the commodity price.

The literature has long recognized that trading in financial markets aggregates information and the resulting prices can feed back to real activity (e.g., Bray (1981) and Subrahmanyam and Titman (2001)). Furthermore, recent literature points out that such feedback effects can be particularly strong in the presence of strategic complementarity in agents' actions. Morris and Shin (2002) show that in such a setting, noise in public information has an amplified effect on agents' actions and thus on equilibrium outcomes. In our model, commodity prices serve such a role in feeding back noise to goods producers' production decisions. Similar feedback effects are also modeled in several other contexts, such as from stock prices to firm capital investment decisions and from exchange rates to policy choices of central banks (e.g., Ozdenoren and Yuan (2008), Angeletos, Lorenzoni, and Pavan (2010), and Goldstein, Ozdenoren, and Yuan (2011, 2013)). The log-linear equilibrium derived in our model accommodates the nonlinearity induced by goods producers' production decisions and, at the same time, is tractable for the analysis of feedback effects of commodity prices. This tractable log-linear equilibrium can be adapted by future studies to analyze feedback effects in settings outside commodity markets.[43]

The paper is organized as follows. We first present the model setting in Section I and derive the equilibrium in Section II. Section III analyzes the effects of informational frictions.

[43] It is also worth noting that our setting is different from existing settings adopted by the literature to analyze real consequences of asset prices. For example, Goldstein, Ozdenoren, and Yuan (2013) develop a model to analyze stock market trading with asymmetric information and the feedback effect from the equilibrium stock price to firm investment. The equilibrium derived in their model is also nonlinear. They ensure tractability by imposing a set of assumptions, including that each trader is risk-neutral and faces upper and lower position limits and that noisy stock supply follows a rigid functional form involving the cumulative standard normal distribution function. Our setting does not require these nonstandard assumptions.

Section IV provides a brief summary of a model extension to include a futures market. We discuss the implications of our analysis in Section V and conclude the paper in Section VI. We provide all the technical proofs in the Appendix and provide a separate online appendix to provide the details of the model extension summarized in Section IV.[44]

2.2. Model Setting

In this section we develop a baseline model with two dates $t = 1, 2$ to analyze the effects of informational frictions on the market equilibrium related to a commodity. One can think of this commodity as crude oil or copper, which is used across the world as a key production input. We model a continuum of islands of total mass one. Each island produces a single good, which can be either consumed at "home" or traded for another good produced "away" by another island. A key feature of the baseline model is that the commodity market is not only a place for market participants to trade the commodity but also a platform to aggregate private information about the strength of the global economy, which ultimately determines the global demand for the commodity.

Table I summarizes the timeline of the model. There are three types of agents: households on the islands, goods producers on the islands, and a group of commodity suppliers. The goods producers trade the commodity with commodity suppliers at $t = 1$ and use the commodity to produce goods at $t = 2$. Their produced goods are distributed to the households on their respective islands at $t = 2$. The households then trade their goods with each other and consume.

Table 1.1 Time Line of the Model

	t=1	t=2
	spot market	goods market
Households		trade/consume goods
Producers	observe signals acquire commodity produce goods	
Com Suppliers	observe supply shock supply commodity	

[44] The Internet Appendix is available in the online version of the article on the Journal of Finance website.

2.2.1 Island Households

Each island has a representative household. Following Angeletos and La'O (2013), we assume a particular structure for goods trading between households on different islands. Each island is randomly paired with another island at $t = 2$. The households on the two islands trade their goods with each other and consume both goods produced by the islands. For a pair of matched islands, we assume that the preference of the households on these islands over the consumption bundle (C_i, C_i^*), where C_i represents consumption of the "home" good while C_i^* consumption of the "away" good, is determined by a utility function $U(C_i, C_i^*)$. The utility function increases in both C_i and C_i^*. This utility function specifies all "away" goods as perfect substitutes, so that the utility of the household on each island does not depend on the matched trading partner. The households on the two islands thus trade their goods to maximize the utility of each. We assume that the utility function of the island households takes the Cobb-Douglas form

$$U(C_i, C_i^*) = \left(\frac{C_i}{1-\eta}\right)^{1-\eta} \left(\frac{C_i^*}{\eta}\right)^{\eta} \tag{16}$$

where $\eta \in [0, 1]$ measures the utility weight of the away good. A greater η means that each island values more of the away good and thus relies more on trading its good with other islands. Thus, η eventually determines the degree of complementarity in the islands' goods production.

2.2.2. Goods Producers

Each island has a locally owned representative firm to organize its goods production. We refer to each firm as a producer. Production requires use of the commodity as an input. To focus on the commodity market equilibrium, we exclude other inputs such as labor from production. Each island has the following decreasing-returns-to-scale production function:[45]

$$Y_i = AX_i^{\phi}, \tag{17}$$

where Y_i is the output produced by island i and X_i is the commodity input. Parameter $\phi \in (0, 1]$ measures the degree to which the production function exhibits decreasing returns

[45] One can also specify a Cobb-Douglas production function with both commodity and labor as inputs. The model remains tractable although the formulas become more complex and harder to interpret.

to scale. When $\phi = 1$, the production function has constant returns to scale. The variable A is the common productivity shared by all islands. For simplicity, we assume that each island's productivity does not have an idiosyncratic component. This simplification is innocuous for our qualitative analysis of how information frictions affect commodity demand.

For an individual goods producer, A has a dual role—it determines its own output as well as other producers' output. To the extent that demand for the producer's good depends on other producers' output, A represents the strength of the global economy. We assume that A is a random variable, which becomes observable only when the producers complete their production at $t = 2$. This is the key informational friction in our setting. We assume that A has a lognormal distribution,

$$\log A \backsim \mathcal{N}\left(\bar{a}, \tau_A^{-1}\right),$$

where $\bar{a}$ is the mean of $\log A$ and τ_A^{-1} is its variance. At $t = 1$, the goods producer on each island observes a private signal about $\log A$,

$$s_i = \log A + \varepsilon_i,$$

where $\varepsilon_i \backsim \mathcal{N}\left(0, \tau_s^{-1}\right)$ is random noise independent of $\log A$ and independent of noise in other producers' signals, and τ_s is the precision of the signal. The signal allows the producer to form its expectation of the strength of the global economy and determine its production decision and commodity demand. The commodity market serves to aggregate the private signals dispersed among the producers. As each producer's private signal is noisy, the publicly observed commodity price also serves as a useful price signal to form its expectation.

At $t = 1$, the producer on island i maximizes its expected profit by choosing its commodity input X_i,

$$\max_{X_i} \; E\left[P_i Y_i \mid \mathcal{I}_i\right] - P_X X_i, \tag{18}$$

where P_i is the price of the good produced by the island. The producer's information set $\mathcal{I}_i = \{s_i, P_X\}$ includes its private signal s_i and the commodity price P_X. The goods price P_i, which one can interpret as the terms of trade, is determined at $t = 2$ based on the matched trade with another island.

2.2.3. Commodity Suppliers

We assume there is a group of commodity suppliers who face a convex labor cost

$$\frac{k}{1+k} e^{-\xi/k} \left(X_S\right)^{\frac{1+k}{k}}$$

in supplying the commodity, where X_S is the quantity supplied, $k \in (0,1)$ is a constant parameter, and ξ represents random noise in the supply. As a key source of information frictions in our model, we assume that ξ is observable to the suppliers themselves but not by other market participants. We assume that from the perspective of goods producers, ξ has Gaussian distribution $\mathcal{N}\left(\bar{\xi}, \tau_{\xi}^{-1}\right)$, where $\bar{\xi}$ is its mean and τ_{ξ}^{-1} as its variance. The mean captures the part that is predictable to goods producers, while the variance represents uncertainty in supply that is outside goods producers' expectations.

Based on the above, given a spot price P_X, suppliers face the following optimization problem:

$$\max_{X_S} \; P_X X_S - \frac{k}{1+k} e^{-\xi/k} \left(X_S\right)^{\frac{1+k}{k}}. \tag{19}$$

It is easy to determine from (19) that the suppliers' optimal supply curve is

$$X_S = e^{\xi} P_X^k, \tag{20}$$

where ξ is uncertainty in the commodity supply and k the price elasticity.

2.2.4. Joint Equilibrium of Different Markets

Our model features the joint equilibrium of a number of markets: the goods markets between each pair of matched islands and the market for the commodity. Equilibrium requires clearing of each of these markets:

- At $t = 2$, for each pair of randomly matched islands $\{i, j\}$, the households of these islands trade their produced goods and clear the market for each good,

$$\begin{aligned} C_i + C_j^* &= AX_i^{\phi}, \\ C_i^* + C_j &= AX_j^{\phi}. \end{aligned}$$

- At $t = 1$, in the commodity market, the goods producers' aggregate demand equals the supply,

$$\int_{-\infty}^{\infty} X_i\left(s_i, P_X\right) d\Phi\left(\varepsilon_i\right) = X_S\left(P_X\right),$$

  where each producer's commodity demand $X_i\left(s_i, P_X\right)$ depends on its private signal s_i and the commodity price P_X. The demand from producers is integrated over the noise ε_i in their private signals.

2.3 The Equilibrium

2.3.1. Goods Market Equilibrium

We begin our analysis of the equilibrium with the goods markets at $t = 2$. For a pair of randomly matched islands, i and j, the representative household of island i possesses Y_i units of the good produced by the island while the representative household of island j holds Y_j units of the other good.[46] They trade the two goods with each other to maximize the utility function of each given in (16). The following proposition, which resembles a similar proposition in Angeletos and La'O (2013), describes the goods market equilibrium between these two islands.

PROPOSITION 2.1: For a pair of randomly matched islands, i and j, their representative households' optimal consumption of the two goods is

$$C_i = (1-\eta) Y_i, \; C_i^* = \eta Y_j, \; C_j = (1-\eta) Y_j, \; C_j^* = \eta Y_i.$$

The price of the good produced by island i is

$$P_i = \left(\frac{Y_j}{Y_i}\right)^{\eta}. \tag{21}$$

As a direct implication of the Cobb-Douglas utility function, each household divides its consumption between the home and away goods with fractions $1-\eta$ and η, respectively. When $\eta = 1/2$, the household consumes the two types of goods equally. The price of each good is determined by the relative output of the two matched islands.[47] One island's good is more valuable when the other island produces more. This feature is standard in the international macroeconomics literature (e.g., Obstfeld and Rogoff (1996)) and implies that each goods producer needs to take into account the production decisions of producers of other goods.[48]

[46] Here we treat a representative household as representing different agents holding stakes in an island's goods production, such as workers, managers, suppliers of inputs, etc. We agnostically group their preferences for the produced goods of their own island and other islands into the preferences of the representative household.

[47] The goods price P_i given in (21) is the price of good i normalized by the price of good j produced by the other matched island.

[48] Decentralized goods market trading is not essential to our analysis. This feature allows us to conveniently capture endogenous complementarity in goods producers' production decisions with tractability. Alternatively, one can adopt centralized goods markets and let island households consume goods produced by all producers. See Angeletos and La'O (2009) for such a setting. We expect our key insight to carry over to this alternative setting.

2.3.2. Production Decision and Commodity Demand

By substituting the production function in (17) into (18), which gives the expected profit of the goods producer on island i, we obtain the following objective:

$$\max_{X_i} \; E\left[AP_i X_i^{\phi} \middle| \, s_i, P_X \right] - P_X X_i.$$

In a competitive goods market, the producer will produce to the level that marginal revenue equals marginal cost:

$$\phi E\left[AP_i \middle| \, s_i, P_X \right] X_i^{\phi-1} = P_X.$$

By substituting in P_i from Proposition 1, we obtain

$$X_i = \left\{ \frac{\phi E\left[AX_j^{\phi\eta} \middle| \, s_i, P_X \right]}{P_X} \right\}^{1/(1-\phi(1-\eta))}, \tag{22}$$

which depends on the producer's expectation $E\left[AX_j^{\phi\eta} \middle| \, s_i, P_X \right]$ regarding the product of global productivity A and the production decision $X_j^{\phi\eta}$ of its randomly matched trading partner, island j. This expression demonstrates the complementarity in the producers' production decisions. A larger η makes the complementarity stronger as the island households engage more in trading the produced goods with each other and the price of each good depends more on the output of other goods.

The commodity price P_X is a source of information for the producer to form its expectation of $E\left[AX_j^{\phi\eta} \middle| \, s_i, P_X \right]$, which serves as a channel for the commodity price to feed back into each producer's commodity demand. The presence of complementarity strengthens this feedback effect relative to standard models of asset market trading with asymmetric information.

2.3.3. Commodity Market Equilibrium

By clearing the aggregate demand of goods producers with the supply of suppliers, we derive the commodity market equilibrium. As is common in settings with real investment, equation (22) shows that each producer's commodity demand is a nonlinear function of the price. Despite the nonlinearity, we manage to derive a tractable and unique log-linear equilibrium in closed form. The following proposition summarizes the commodity price and each producer's commodity demand in this equilibrium.

PROPOSITION 2.2: At $t = 1$, the commodity market has a unique log-linear equilibrium: 1) The commodity price is a log-linear function of $\log A$ and ξ,

$$\log P_X = h_A \log A + h_\xi \xi + h_0, \tag{23}$$

with the coefficients h_A and h_ξ given by

$$h_A = -\frac{(1-\phi)\, b + (1-\phi(1-\eta))\, \tau_s^{-1} \tau_\xi b^3}{1 + k(1-\phi)} > 0, \tag{24}$$

$$h_\xi = -\frac{1-\phi + (1-\phi(1-\eta))\, \tau_s^{-1} \tau_\xi b^2}{1 + k(1-\phi)} < 0, \tag{25}$$

where $b < 0$ is given in equation (A.19) and h_0 in equation (A.20). 2) The commodity purchased by goods producer i is a log-linear function of its private signal s_i and $\log P_X$,

$$\log X_i = l_s s_i + l_P \log P_X + l_0, \tag{26}$$

with the coefficients l_s and l_P given by

$$l_s = -b > 0, \;\; l_P = k + h_\xi^{-1}, \tag{27}$$

and l_0 by equation (A.21).

Proposition 2.2 shows that each producer's commodity demand is a log-linear function of its private signal and the commodity price, while the commodity price $\log P_X$ aggregates the producers' dispersed private information to partially reveal the global productivity $\log A$. The commodity price does not depend on any producer's signal noise as a result of the aggregation across a large number of producers with independent noise. This feature is similar to Hellwig (1980). The commodity price also depends on the supply shock ξ, which serves the same role as noise trading in the standard models of asset market trading with asymmetric information.

It is well known that asset market equilibria with asymmetric information are often intractable due to the difficulty in analyzing each agent's learning from the equilibrium asset price and in aggregating different agents' asset demands. Existing literature commonly adopts the setting of Grossman and Stiglitz (1980) and Hellwig (1980), which features CARA utility for agents and Gaussian distributions for asset fundamentals and noise trading. This setting ensures a linear equilibrium in which the asset price is a linear function of asset fundamental and noise trading, while each agent's asset demand is a linear function of the price and its own signal. One cannot directly adopt this setting, however, to analyze informational feedback effects of asset prices to real activity, such as firm investment and central bank policies, which typically involve asset fundamentals with non-Gaussian distributions and agents with non-CARA utility.

Our model presents a tractable setting to analyze real consequences of asset prices. Despite the commodity price and each producer's commodity demand both having non-Gaussian distributions, the log-linear equilibrium derived in Proposition 2.2 maintains similar tractability as the linear equilibrium derived by Grossman and Stiglitz (1980) and Hellwig (1980). A key feature contributing to this tractability is that the producers' aggregate demand remains log-normal as a result of the Law of Large Numbers.

2.4 Effects of Informational Frictions

2.4.1. Perfect-Information Benchmark

To facilitate our analysis of the effects of informational frictions, we first establish a benchmark without any informational friction. Suppose that the global fundamental A and commodity supply shock ξ are both observable by all market participants. Then the goods producers can choose their optimal production decisions without any noise interference. The following proposition characterizes this benchmark.

PROPOSITION 2.3: When both A and ξ are observed by all market participants, there is a unique equilibrium. In this equilibrium, 1) the goods producers share an identical commodity demand curve, $X_i = X_j = \left(\frac{\phi A}{P_X}\right)^{\frac{1}{1-\phi}}$, $\forall i$ and j, and 2) the commodity price is given by

$$\log P_X = \frac{1}{1+k(1-\phi)} \log A - \frac{1-\phi}{1+k(1-\phi)} \xi + \frac{1}{1+k(1-\phi)} \log \phi,$$

while the goods producers' aggregate commodity demand is given by

$$\log X_S = \frac{k}{1+k(1-\phi)} \log A + \frac{1}{1+k(1-\phi)} \xi + \frac{k}{1+k(1-\phi)} \log \phi.$$

In the absence of any informational frictions, the benchmark features a unique equilibrium despite the complementarity in the goods producers' production decisions because competition between goods producers leads to a downward-sloping demand curve for the commodity. This demand curve intersects the suppliers' upward-sloping supply curve at the unique commodity price P_X given in the proposition. As a result, the complementarity between goods producers does not lead to multiple equilibria in which goods producers coordinate on certain high or low demand levels.

Proposition 3 derives the equilibrium commodity price and aggregate demand. Intuitively, the global fundamental $\log A$ increases both the commodity price and aggregate demand, while the supply shock ξ reduces the commodity price but increases aggregate demand.

The following proposition compares the equilibrium derived in Proposition 2.2 with the perfect-information benchmark.

PROPOSITION 2.4: In the presence of informational frictions, the commodity price coefficients with the global fundamental $h_A > 0$ and the commodity supply shock $h_\xi < 0$, as derived in Proposition 2.2, are both lower than their corresponding values in the perfect-information benchmark, and converge to these values as $\tau_s \to \infty$.

In the presence of informational frictions, the commodity price deviates from that in the perfect-information benchmark, with the supply shock having a greater price impact (i.e., h_ξ being more negative) and the global fundamental having a smaller impact (i.e., h_A being less positive). Through these price impacts, informational frictions eventually affect goods producers' production decisions and island households' goods consumption, which we analyze step-by-step below.

2.4.2. Price Informativeness

In the presence of informational frictions, the equilibrium commodity price $\log P_X = h_A \log A + h_\xi \xi + h_0$ serves as a public signal of the global fundamental $\log A$. This price signal is contaminated by the presence of the supply noise ξ. The informativeness of the price signal is determined by the ratio of the contributions to the price variance of $\log A$ and ξ:

$$\pi = \frac{h_A^2/\tau_A}{h_\xi^2/\tau_\xi}.$$

The following proposition characterizes how the price informativeness measure π depends on several key model parameters: τ_s, τ_ξ, and η.

PROPOSITION 2.5: π is monotonically increasing in τ_s and τ_ξ, and is decreasing in η.

As τ_s increases, each goods producer's private signal becomes more precise. The commodity price aggregates the goods producers' signals through their demand for the commodity and therefore becomes more informative. The parameter τ_ξ measures the amount of noise in the supply shock. As τ_ξ increases, there is less noise from the supply side interfering with the commodity price reflecting $\log A$. Thus, the price also becomes more informative.

The effect of η is more subtle. As η increases, there is greater complementarity in each goods producer's production decision. Consistent with the insight of Morris and Shin (2002), such complementarity induces each producer to put greater weight on the publicly observed price signal and lesser weight on its own private signal, which makes the equilibrium price less informative.

2.4.3. Price Elasticity

The coefficient l_P, derived in (27), measures the price elasticity of each goods producer's commodity demand. The standard cost effect suggests that a higher price leads to a lower demand. The producer's optimal production decision in equation (22), however, also indicates a second effect through the term in the numerator—a higher price signals a stronger global economy and greater production by other producers. This informational effect motivates each producer to increase its production and thus demand more of the commodity. The price elasticity l_P nets these two offsetting effects. The following proposition shows that under certain necessary and sufficient conditions, the informational effect dominates the cost effect and leads to a positive l_P.

PROPOSITION 2.6: Two necessary and sufficient conditions ensure that $l_P > 0$: first,

$$\tau_\xi / \tau_A > 4k^{-1}\left(1 - \phi + k^{-1}\right),$$

and, second, parameter η is within the range

$$1 - \frac{1}{\phi} + \frac{k\tau_\xi\tau_s}{4\phi\tau_A^2}\left(1-\rho\right)^2 < \eta \ < 1 - \frac{1}{\phi} + \frac{k\tau_\xi\tau_s}{4\phi\tau_A^2}\left(1+\rho\right)^2,$$

where $\rho = \tau_A^{1/2}\tau_\xi^{-1/2}\sqrt{\tau_\xi/\tau_A - 4k^{-1}\left(1-\phi+k^{-1}\right)}$.

For the informational effect to be sufficiently strong, the commodity price has to be sufficiently informative. The conditions in Proposition 6 reflect this observation. First, the supply noise needs to be sufficiently small (i.e., τ_ξ sufficiently large relative to τ_A) so that the price can be sufficiently informative. Second, η needs to be within an intermediate range, which results from two offsetting forces. On the one hand, a larger η implies greater complementarity in producers' production decisions and thus each producer cares more about other producers' production decisions and assigns greater weight to the public price signal in its own decision making. On the other hand, a larger η also implies a less informative price signal (Proposition 2.5), which motivates each producer to be less responsive to the price. Netting out these two forces dictates that η needs to be in an intermediate range for $l_P > 0$.[49]

This second condition implies that when $\eta = 0$, $l_P < 0$. Therefore, in the absence of production complementarity, the price elasticity is always negative, that is, the cost effect always dominates the informational effect.

2.4.4. Feedback Effect on Demand

In the perfect-information benchmark (Proposition 2.3), the supply shock ξ decreases the commodity price and increases the aggregate demand through the standard cost effect. In the presence of informational frictions, however, the supply shock, by distorting the price signal, has a more subtle effect on commodity demand.

By substituting equation (23) into (26), the commodity demand of producer i is

$$\log X_i = l_s s_i + l_P h_A \log A + l_P h_\xi \xi + l_P h_0 + l_0.$$

The producers' aggregate commodity demand is then

$$\log\left[\int_{-\infty}^{\infty} X_i\left(s_i, P_X\right) d\Phi\left(\varepsilon_i\right)\right] = l_P h_\xi \xi + \left(l_s + l_P h_A\right)\log A + l_0 + l_P h_0 + \frac{1}{2} l_s^2 \tau_s^{-1}.$$

Note that $h_\xi < 0$ (Proposition 2.2) and the sign of l_P is undetermined (Proposition 2.6). Thus, the effect of ξ on the aggregate demand is also undetermined.

Under the conditions given in Proposition 6, an increase in ξ decreases the aggregate demand, which is the opposite of the perfect-information benchmark. This effect arises through the informational channel. As ξ rises, the commodity price falls. Since goods

[49] Upward-sloping demand for an asset may also arise from other mechanisms even in the absence of informational frictions highlighted in our model, such as income effects, complementarity in production, and complementarity in information production (e.g., Hellwig, Kohls, and Veldkamp (2012)).

producers cannot differentiate a price decrease caused by ξ from one caused by a weaker global economy, they partially attribute the reduced price to a weaker economy. This, in turn, motivates them to cut their commodity demand. Under the conditions given in Proposition 6, this informational effect is sufficiently strong to dominate the effect of a lower cost to acquire the commodity, leading to a lower aggregate commodity demand.

Furthermore, through its informational effect on aggregate demand, ξ can further push down the commodity price in addition to its price effect in the perfect-information benchmark. This explains why h_ξ is more negative in this economy than in the benchmark (Proposition 2.4): informational frictions amplify the negative price impact of ξ.[50]

2.4.5. Social Welfare

By distorting the commodity price and aggregate demand, informational frictions distort producers' production decisions and households' goods consumption. We now evaluate the unconditional expected social welfare at time 1:

$$W = E\left[\int_0^1 \left(\frac{C_i}{1-\eta}\right)^{1-\eta}\left(\frac{C_i^*}{\eta}\right)^{\eta} di\right] - E\left[\frac{k}{1+k}e^{-\xi/k}X_S^{\frac{1+k}{k}}\right],$$

which contains two parts. The first part comes from aggregating the expected utility from goods consumption of all island households, and the second part comes from the commodity suppliers' cost of supplying labor.

The next proposition proves that informational frictions reduce the expected social welfare relative to the perfect-information benchmark.

PROPOSITION 2.7: In the presence of informational frictions, the expected social welfare is strictly lower than that in the perfect-information benchmark.

[50] One can also evaluate this informational feedback effect of the supply noise by comparing the equilibrium commodity price relative to another benchmark case, in which each goods producer makes his production decision based only on his private signal s_i without conditioning it on the commodity price P_X. In this benchmark, the commodity price $\log P_X$ is also a log-linear function of $\log A$ and ξ. Interestingly, despite the presence of informational frictions, the price coefficient on ξ is $-\frac{1-\phi}{1+k(1-\phi)}$, which is the same as that derived in Proposition 3 for the perfect-information benchmark. This outcome establishes the informational feedback mechanism as the driver for h_ξ to be more negative than that in the perfect-information benchmark.

2.5 A Model Extension

Stimulated by the large inflow of investment capital to commodity futures markets in recent years, there is an ongoing debate about whether speculation in futures markets might have affected commodity prices.[51] In this debate, an influential argument posits that as the trading of financial traders in futures markets does not directly affect the supply and demand of physical commodities, there is no need to worry about them affecting commodity prices. This argument ignores the informational role of futures prices. In practice, the lower cost of trading futures contracts compared with trading physical commodities encourages greater participation and facilitates aggregation of dispersed information among market participants.[52] To reduce this confusion, we extend our model to incorporate a futures market. For the sake of brevity, we briefly summarize the extended model and the key result in this section and relegate a more detailed model description and analysis Appendix 2.B.

2.5.1. Model Setting

The objective of this extension is not to provide a general model of information aggregation with both spot and futures markets. Instead, we use a specific yet realistic setting to highlight the conceptual point that informational noise introduced by futures market trading can feed back to commodity demand and spot prices.

We introduce a new date $t = 0$ before the two dates $t = 1$ and 2 in the main model, and a centralized futures market at $t = 0$ for delivery of the commodity at $t = 1$. All agents can take positions in the futures market at $t = 0$, and can choose to revise or unwind their positions before delivery at $t = 1$. The ability to unwind positions before delivery reduces transaction costs and makes futures market trading appealing in practice.

We maintain all of the agents in the main model - island households, goods producers, and commodity suppliers - and add a group of financial traders. These financial traders take

[51] Since the mid-2000s, commodity futures has become a new asset class for portfolio investors such as pension funds and endowments, which regularly allocate a fraction of their portfolios to investing in commodity futures and swap contracts. As a result, capital on the order of hundreds of billions of dollars flowed to the long side of commodity futures markets. This process is also called the financialization of commodity markets (e.g., Cheng and Xiong (2014)).

[52] Roll (1984) systematically analyzes the futures market of orange juice in efficiently aggregating information about weather in Central Florida, which produces more than 98% of the U.S. orange output. Garbade and Silber (1983) provide evidence that futures markets play a more important role in information discovery than cash markets for a set of commodities.

a position in the futures market at $t = 0$ and then unwind this position at $t = 1$ without taking delivery. We assume that there is no spot market trading at $t = 0$. A spot market naturally emerges at $t = 1$ through commodity delivery for the futures market.

Table 1.2. Time Line of the Extended Model

	t=0	t=1	t=2
	futures market	spot market	goods market
Households			trade/consume goods
Producers	observe signals long futures	take delivery produce goods	
Com Suppliers	short futures	observe supply shock deliver commodity	
Fin Traders	long/short futures	unwind position	

Table 1.2 specifies the timeline of the extended model. The timing of information flow is key to our analysis. We assume that goods producers receive their respective private signals $\{s_i\}$ about the global productivity at $t = 0$ and commodity suppliers observe their supply shock ξ only at $t = 1$. This structure leads to two rounds of information aggregation: trading in the futures market at $t = 0$ serves as the first round with informational noise originating from the trading of financial traders, and trading in the spot market at $t = 1$ serves as the second round with financial traders unwinding their futures position and commodity suppliers observing their supply shock.

We keep the same specification for the island households, who trade and consume both home and away goods at $t = 2$ as described in Section 2.2.1.

We allow the goods producers to have the same production technology and private signals as specified in Section 2.2.2. At $t = 1$, the producer optimizes its production decision X_i based on the objective function given in (18) and an expanded information set $I_i^1 = \{s_i, F, P_X\}$, where F is the futures price traded at $t = 0$ and P_X is the spot price traded at $t = 1$:

$$X_i = \left\{ \phi E\left[A X_j^{\phi\eta} \middle| \mathcal{I}_i^1 \right] \Big/ P_X \right\}^{1/(1-\phi(1-\eta))}.$$

At $t = 0$, the producer chooses a futures position $\tilde{X}_i$ to maximize the following expected production profit based on its information set $I_i^0 = \{s_i, F\}$:

$$\max_{\tilde{X}_i} \; E\left[P_i Y_i \middle| \, \mathcal{I}_i^0\right] - F\tilde{X}_i.$$

In specifying this objective function, we adopt a simplification by treating the producer as myopic at $t = 0$ (i.e., it treats $\tilde{X}_i$ as its production input at $t = 1$.)[53] Then, the producer's futures position is

$$\tilde{X}_i = \left\{ \phi E\left[A\tilde{X}_j^{\phi\eta} \middle| \mathcal{I}_i^0 \right] \Big/ F \right\}^{1/(1-\phi(1-\eta))}.$$

We assume that in the futures market at $t = 0$, the aggregate long position of financial traders and goods producers is given by the aggregate position of producers multiplied by a factor $e^{\kappa \log A + \theta}$: $e^{\kappa \log A+\theta} \int_{-\infty}^{\infty} \tilde{X}_i\left(s_i, F\right) d\Phi\left(\varepsilon_i\right)$, where the factor $e^{\kappa \log A+\theta}$ represents the contribution of financial traders. This multiplicative specification is useful for ensuring the tractable log-linear equilibrium of our model. The component $\kappa \log A$ with $\kappa > 0$ captures the possibility that the trading of financial traders is partially driven by their knowledge of the global fundamental $\log A$, while the other component $\theta \backsim N\left(\overline{\theta}, \tau_\theta^{-1}\right)$, a random Gaussian variable with a mean of $\overline{\theta}$ and variance τ_θ^{-1}, captures non-fundamental related trading induced by diversification motives and is unobservable by other market participants.

We allow the commodity suppliers to have the same convex cost function specified in Section 2.2.3. At $t = 1$, they observe their supply shock ξ and their marginal cost of supplying the commodity determines the spot price P_X. At $t = 0$, the suppliers take a short position in the futures market. To simplify the analysis, we assume that the suppliers are also myopic in believing that goods producers will take full delivery of their futures positions. Thus, the suppliers choose an initial short position to maximize the profit from making delivery of $e^{-(\kappa \log A+\theta)}\tilde{X}_S$ units of the commodity to goods producers:

$$\max_{\tilde{X}_S} E\left[Fe^{-(\kappa \log A+\theta)}\tilde{X}_S \middle| \, \mathcal{I}_S^0 \right] - E\left[\frac{k}{1+k} e^{-\xi/k} \left(e^{-(\kappa \log A+\theta)}\tilde{X}_S \right)^{\frac{1+k}{k}} \middle| \, \mathcal{I}_S^0 \right],$$

which gives that

$$\tilde{X}_S = e^{\bar{\xi}-\sigma_\xi^2/2k} \left\{ E\left[e^{-(\kappa \log A+\theta)} \middle| \, \mathcal{I}_S^0\right] / E\left[e^{-\frac{1+k}{k}(\kappa \log A+\theta)} \middle| \, \mathcal{I}_S^0 \right] \right\}^k F^k.$$

[53] In other words, at $t = 0$ each producer chooses a futures position as if it commits to taking full delivery and using the good for production, even though the producer can revise its production decision based on the updated information at $t = 1$. This simplifying assumption, while it affects each producer's trading profit, is innocuous for our analysis of how the futures price feeds back to the producers' later production decisions because each producer still makes good use of its information and the futures price is informative by aggregating each producer's information.

2.5.2. The Equilibrium and Key Result

We analyze the joint equilibrium of all markets: the goods markets between each pair of matched islands at $t = 2$, the spot market for the commodity at $t = 1$, and the futures market at $t = 0$. We derive a unique log-linear equilibrium of these markets in Appendix 2.B, and summarize only the key features of the equilibrium here.

During the first round of trading in the futures market at $t = 0$, the futures price aggregates the goods producers' private signals and is a log-linear function of $\log A$ and θ:

$$\log F = \tilde{h}_A \log A + \tilde{h}_\theta \theta + \tilde{h}_0, \tag{28}$$

where $\tilde{h}_A > 0$ and $\tilde{h}_\theta > 0$. The futures price does not fully reveal the global productivity $\log A$ because of the noise θ originated from the trading of financial traders.

The spot price that emerges from the commodity delivery at $t = 1$ represents another round of information aggregation by pooling together the goods producers' demand for delivery. As a result of the arrival of the supply shock ξ, the spot price $\log P_X$ does not fully reveal $\log A$ or θ, but instead reflects a linear combination of $\log A$, $\log F$, and ξ:

$$\log P_X = h_A \log A + h_F \log F + h_\xi \xi + h_0, \tag{29}$$

where $h_A > 0$, $h_F > 0$, and $h_\xi < 0$.

Despite the updated information from the spot price at $t = 1$, the informational content of $\log F$ is not subsumed by the spot price, and still has an influence on goods producers' expectations of the global productivity. As a result of this informational role, the commodity consumed by producer i at $t = 1$ is increasing with $\log F$:

$$\log X_i = l_s s_i + l_F \log F + l_P \log P_X + l_0, \tag{30}$$

where $l_s > 0$ and $l_F > 0$. The coefficient on the spot price l_P has an undetermined sign, which reflects the offsetting cost effect and informational effect of the spot price, similar to our characterization of the main model.

While the trading of financial traders does not have any direct effect on commodity supply and demand, it affects the futures price, through which it can further impact commodity demand and the spot price. By substituting equation (28) into (29), we express the spot price $\log P_X$ as a linear combination of the primitive shocks $\log A$, θ, and ξ:

$$\log P_X = \left(h_A + h_F \tilde{h}_A\right) \log A + h_F \tilde{h}_\theta \theta + h_\xi \xi + h_F \tilde{h}_0 + h_0. \tag{31}$$

This expression shows that θ, the noise from financial traders' futures position, has a positive effect on the spot price. Furthermore, by substituting (31) and (28) into (30) and then integrating the individual producers' commodity demands, their aggregate demand is

$$\log\left[\int_{-\infty}^{\infty} X\left(s_i, F, P_X\right) d\Phi\left(\varepsilon_i\right)\right] = \left[l_s + l_P h_A + l_F \tilde{h}_A + l_P h_F \tilde{h}_A\right] \log A + \left(l_F + l_P h_F\right) \tilde{h}_\theta \theta + l_P h_\xi \xi + \left(l_F + l_P h_F\right) \tilde{h}_0 + l_P h_0 + l_0 + \frac{1}{2} l_s^2 \tau_s^{-1}. \tag{32}$$

We can further derive that the coefficient on θ in the aggregate commodity demand is

$$l_F + l_P h_F = k h_F > 0.$$

Thus, θ also has a positive effect on aggregate commodity demand.

The effects of θ on commodity demand and the spot price clarify the simple yet important conceptual point that traders in commodity futures markets, who never take or make physical delivery, can nevertheless impact commodity markets through the informational feedback channel of commodity futures prices. Information frictions in the futures market, originating from the unobservability of the positions of different participants, are essential for this feedback effect. In Appendix 2.B, we further derive that as $\tau_\theta \to \infty$ (i.e., the position of financial traders becomes publicly observable), the spot market equilibrium converges to the perfect-information benchmark. This result highlights the importance of improving transparency in futures markets.

2.6 Implications

In this section, we discuss implications of our model for several empirical issues: estimating the effects of supply and demand shocks, detecting speculative effects in commodity markets, and understanding the puzzling commodity price boom in 2007 to 2008.

2.6.1. Estimating Effects of Supply and Demand Shocks

The feedback effect of commodity prices has important implications for studies of the effects of supply and demand shocks in commodity markets. For example, Hamilton (1983) emphasizes that disruptions to oil supply and resulting oil price increases can have a significant impact on the real economy, while Kilian (2009) argues that aggregate demand shocks have a bigger impact on the oil market than previously thought. As supply and demand shocks have opposite effects on oil prices, it is important to isolate their respective effects.

Existing literature commonly uses structural vector autoregressions (SVARs) to decompose historical commodity price dynamics. The premise of these structural models is that, while researchers cannot directly observe the shocks that hit commodity markets, agents in the economy are able to observe the shocks and optimally respond to them. As highlighted by our model, it is unrealistic to assume that agents can perfectly differentiate different types of shocks. In particular, our model shows that in the presence of informational frictions, supply shocks and demand shocks can have effects in sharp contrast to standard intuition developed from perfect-information settings. These contrasts render the structural models that ignore informational frictions unreliable and potentially misleading.

We now use the popular SVAR model developed by Kilian (2009) for the global oil market as an example. This model specifies the dynamics for a vector z_t as

$$\mathbf{A}_0 \mathbf{z}_t = \boldsymbol{\alpha} + \sum_{i=1}^{24} \mathbf{A}_i \mathbf{z}_{t-i} + \boldsymbol{\varepsilon}_t.$$

The vector z_t contains three variables: global crude oil production, a measure of real activity, and the spot price for oil. The vector ε_t contains serially uncorrelated and mutually independent structural shocks that hit the global oil market from different sources, such as an oil supply shock, a global demand shock, and an oil-specific demand shock. By imposing various restrictions on the matrix A_0 (which is assumed to be invertible), the model recovers the structural shocks ε_t from shocks e_t estimated from a reduced-form VAR model for z_t according to $\varepsilon_t = A_0 e_t$. Without going through the specific restrictions, the restrictions imposed in the literature are typically motivated by conventional wisdom regarding how supply, demand, and the spot price should react to the structural shocks under the implicit assumption that agents can directly observe them. Under this assumption, the structural shocks ε_t and the innovations to z_t (i.e., e_t) are informationally equivalent.

It is important to recognize that, in practice, agents observe neither the structural shocks ε_t nor the full vector z_t. While agents can observe oil prices in a timely fashion, they observe quantity variables such as global oil production and GDP with a substantial delay on the order of several quarters.[54] This delay in observing the full vector z_t makes it impossible for

[54] This delay results from the fact that it often takes several quarters for different countries to report both their GDP and their supply of and demand for crude oil, and some countries may even choose not to report at all. The measure of real activity used by Kilian (2009) builds on an index from dry bulk cargo freight rates. This index, while useful, is more a measure of expectations than a direct indicator of real activity.

agents to fully recover the structural shocks. Instead, they have to rely on what they can observe at the time to partially infer these shocks. Thus, by assuming that agents can directly observe the structural shocks, the model by Kilian (2009) ignores the realistic informational frictions that agents face in the global oil market. Without a systematic comparison using a correctly specified model, it is difficult to precisely determine the consequences of the misspecification. According to our model, since agents cannot disentangle supply and demand shocks, they partially attribute the observed price change caused by a positive supply shock to a weaker global economy. As a result, they reduce their own commodity demand, which amplifies the price impact of the initial supply shock. Therefore, by ignoring this learning effect induced by informational frictions, the misspecified SVAR model is likely to understate the effect of supply shocks and overstate of the effect of demand shocks.

Our model provides the basic ingredients for constructing more complete empirical models that account for informational frictions faced by economic agents. Ideally, one would want to build a full economic model that systematically accounts for how commodity prices aggregate agents' dispersed information and how each agent forms its expectations based on publicly observed commodity prices together with its own private signal. Even without such a model, one can still extend the more practical SVAR approach to explicitly account for the information set available to agents at the time they make their decisions. According to our analysis, the key is to account for how commodity prices impact agents' expectations.

2.6.2. Detecting Speculative Effects

In the ongoing debate on whether speculation has affected commodity prices during the commodity market boom and bust of 2007 to 2008, many studies (e.g., Kilian and Murphy (2014), Juvenal and Petrella (2012), and Knittel and Pindyck (2013)) adopt an inventory-based detection strategy. This strategy builds on the widely held argument that if speculators artificially drive up the commodity price, consumers will find consuming the commodity too expensive and thus reduce consumption, causing inventory of the commodity to spike.

Under this argument, price increases in the absence of inventory increases are explained by fundamental demand. Consequently, price effects induced by speculation should be limited to price increases that are accompanied by contemporaneous increases in inventory. Motivated by this argument, the literature, as reviewed by Fattouh, Kilian, and Mahadeva (2012), tends to use the lack of pronounced oil inventory spike before the July 2008 peak in oil prices as evidence ruling out any significant role played by speculation during the oil price boom.

Despite the intuitive appeal of this inventory-based detection strategy, it ignores important informational frictions faced by consumers in reality. Like the SVAR models we discussed earlier, it crucially relies on the assumption that oil consumers observe global economic fundamentals and are therefore able to recognize whether current oil prices are too high relative to fundamentals in making their consumption decisions. This assumption is unrealistic during periods with great economic uncertainty, especially during 2007 to 2008 when consumers faced severe informational frictions in inferring the strength of the global economy.

Our model illustrates a counter example to this widely used detection strategy. Under the conditions specified by Proposition 6, the price elasticity of the goods producers' commodity demand is positive.[55] In such an environment where goods producers have a positive demand elasticity, if speculation drives up the commodity price, the increased price will also cause goods producers to consume more rather than less of the commodity by influencing their expectations about the strength of the global economy. Our model therefore shows that in the presence of severe informational frictions, speculation can drive up commodity prices without necessarily reducing commodity consumption and boosting inventory. This insight points to the weak power of the widely used inventory-based strategy in detecting speculative effects. In this light, the absence of a pronounced oil inventory spike before the July 2008 peak in oil prices cannot be taken as evidence rejecting the presence of a speculative effect during this period.

2.6.3. Understanding the Commodity Price Boom of 2007 to 2008

In the aftermath of the synchronized price boom and bust of major commodities in 2007 to 2008, the price boom has been attributed to the combination of rapidly growing demand from emerging economies and stagnant supply (e.g., Hamilton (2009)). This argument is popular for explaining the commodity price increases before 2008. However, oil prices continued to rise over 40%, peaking at $147 per barrel, from January to July 2008 at a time when the U.S. had already entered a recession (in November 2007 as dated by the NBER), Bear Stearns had collapsed (in March 2008), and most other developed economies were already showing signs of weakness. While emerging economies remained strong at the time, it is difficult to

[55] We can also provide similar conditions for the extended model.

argue, in hindsight, that their growth sped up so much to be able to offset the weakness of the developed economies and cause oil prices to rise another 40%.

The informational frictions faced by market participants can help us understand this puzzling price boom. As a result of the lack of reliable data on emerging economies, it was difficult to precisely measure their economic strength in real time. The prices of crude oil and other commodities were regarded as important price signals (see the evidence referenced in Footnote 42). This environment makes our model particularly appealing for linking the large commodity price increases in early 2008 to the concurrent large inflow of investment capital, motivated by many portfolio managers seeking to diversify their portfolios out of declining stock markets and into the more promising commodity futures markets (e.g., Tang and Xiong (2012)). By pushing up commodity futures prices and sending a wrong price signal, the large investment flow might have confused goods producers across the world into believing that emerging economies were stronger than they actually were. This distorted expectation could have prevented the producers from reducing their commodity demand despite the high commodity prices, which in turn made the high prices sustainable. Even though more information would eventually correct the producers' expectations, the high commodity prices persisted for several months before their collapse in the second half of 2008. Interestingly, after oil prices dropped from their peak of $147 to $40 per barrel at the end of 2008, oil demand largely evaporated and inventory piled up, despite the much lower prices.

Taken together, the commodity price boom of 2007 to 2008 was not necessarily a price bubble detached from economic fundamentals. Instead, it is plausible to argue that, in the presence of severe informational frictions in early 2008, the large inflow of investment capital might have distorted signals coming from commodity prices and led to confusion among market participants about the strength of emerging economies. This confusion, in turn, could have amplified the boom and bust of commodity prices, which echoes Singleton's (2014) emphasis on accounting for agents' expectations in explaining this price cycle. To systematically examine this hypothesis would require estimating a structural model that explicitly accounts for the informational feedback effect of commodity prices.

2.7 Conclusion

This paper develops a tractable model to analyze effects of informational frictions in commodity markets. Our model shows that, through the informational role of commodity prices, goods producers' commodity demand can increase with the price, and supply shocks can have an amplified effect on the price and an undetermined effect on producers' demand. By further incorporating one round of futures market trading, our extended model shows that futures prices can also serve as important price signals, even when goods producers also observe spot prices. Thus, through the same informational channel, noise in futures market trading can also interfere with goods producers' expectations and affect their commodity demand. Our analysis highlights the weakness common in empirical and policy studies of assuming that different shocks are publicly observable to market participants. Our analysis also provides a coherent argument for how the large inflow of investment capital to commodity futures markets, by jamming commodity price signals and leading to confusion about the strength of emerging economies, might have amplified the boom and bust of commodity prices in the 2007 to 2008 period.

2.A: Proofs of Propositions

Proof of Proposition 2.1:

Consider the maximization problem of the household on island i:

$$\max_{\{C_i\}_{i\in[0,1]}} \left(\frac{H_i}{1-\eta_H}\right)^{1-\eta_H} \left\{\frac{1}{\eta_H}\left(\frac{C_i\left(i\right)}{1-\eta_c}\right)^{1-\eta_c}\left(\frac{\int_{[0,1]/i} C_j\left(i\right)dj}{\eta_c}\right)^{\eta_c}\right\}^{\eta_H}$$

subject to the budget constraint

$$P_H + \int_0^1 P_j C_j\left(i\right) dj = P_i A l_i. \tag{A.1}$$

The first order conditions with respect to C_i and C_i^* are

$$\left(\frac{C_i^*}{C_i}\right)^{\eta}\left(\frac{1-\eta}{\eta}\right)^{\eta} = \lambda_i P_i \tag{A.2}$$

$$\left(\frac{C_i}{C_i^*}\right)^{1-\eta}\left(\frac{\eta}{1-\eta}\right)^{1-\eta} = \lambda_i P_j \tag{A.3}$$

where λ_i is the Lagrange multiplier for his budget constraint. Dividing equations (A.2) and (A.3) leads to $\frac{\eta}{1-\eta}\frac{C_i}{C_i^*} = \frac{P_j}{P_i}$, which is equivalent to $P_j C_i^* = \frac{\eta}{1-\eta} P_i C_i$. By substituting this equation back to the household's budget constraint in (A.1), we obtain $C_i = \left(1-\eta\right) Y_i$.

Market clearing of the island's produced goods requires $C_i + C_j^* = Y_i$, which implies that $C_j^* = \eta Y_i$. The symmetric problem of the household of island j implies that $C_j = \left(1-\eta\right) Y_j$, and market clearing of the goods produced by island j implies $C_i^* = \eta Y_j$.

The first-order condition in equation (A.2) also gives the price of the goods produced by island i. Since the household's budget constraint in (A.1) is entirely in nominal terms, the price system is only identified up to λ_i, the Lagrange multiplier. Following Angeletos and La'O (2013), we normalize λ_i to one. Then,

$$P_i = \left(\frac{C_i^*}{C_i}\right)^{\eta}\left(\frac{1-\eta}{\eta}\right)^{\eta} = \left(\frac{\eta Y_j}{\left(1-\eta\right) Y_i}\right)^{\eta}\left(\frac{1-\eta}{\eta}\right)^{\eta} = \left(\frac{Y_j}{Y_i}\right)^{\eta}.$$

Proof of Proposition 2.2:

We first conjecture that the commodity price and each goods producer's commodity demand take the following log-linear forms:

$$\log P_X = h_0 + h_A \log A + h_\xi \xi, \tag{A.4}$$

$$\log X_i = l_0 + l_s s_i + l_P \log P_X, \tag{A.5}$$

where the coefficients h_0, h_A, h_ξ, l_0, l_s, and l_P will be determined by equilibrium conditions.

Define

$$z \equiv \frac{\log P_X - h_0 - h_\xi \bar{\xi}}{h_A} = \log A + \frac{h_\xi}{h_A}\left(\xi - \bar{\xi}\right),$$

which is a sufficient statistic of information contained in the commodity price P_X. Then, conditional on observing its private signal s_i and the commodity price P_X, goods producer i's expectation of $\log A$ is

$$E\left[\log A \mid s_i, \log P_X\right] = E\left[\log A \mid s_i, z\right] = \frac{1}{\tau_A + \tau_s + \frac{h_A^2}{h_\xi^2}\tau_\xi}\left(\tau_A \bar{a} + \tau_s s_i + \frac{h_A^2}{h_\xi^2}\tau_\xi z\right),$$

and its conditional variance of $\log A$ is

$$Var\left[\log A \mid s_i, \log P_X\right] = \left(\tau_A + \tau_s + \frac{h_A^2}{h_\xi^2}\tau_\xi\right)^{-1}.$$

According to equation (22),

$$\log X_i = \frac{1}{1-\phi(1-\eta)}\left\{\log\phi + \log\left(E\left[AX_j^{\phi\eta} \mid s_i, \log P_X\right]\right) - \log P_X\right\}. \tag{A.6}$$

By using equation (A.5), we obtain

$$\begin{aligned}
& E\left[AX_j^{\phi\eta} \mid s_i, \log P_X\right] \\
= \; & E\left\{\exp\left[\log A + \phi\eta\left(l_0 + l_s s_j + l_P \log P_X\right) \mid s_i, z\right]\right\} \\
= \; & \exp\left[\phi\eta\left(l_0 + l_P \log P_X\right)\right] \cdot E\left[\exp\left((1+\phi\eta l_s)\log A + \phi\eta l_s \varepsilon_j\right) | s_i, \log P_X\right] \\
= \; & \exp\left\{\phi\eta\left(l_0 + l_P\log P_X\right) + (1+\phi\eta l_s)E\left[\log A \mid s_i, \log P_X\right] + \frac{(1+\phi\eta l_s)^2}{2}Var\left[\log A \mid s_i, \log P_X\right]\right. \\
& \left. + \frac{\phi^2\eta^2 l_s^2}{2}Var\left[\varepsilon_j \mid s_i, \log P_X\right] + (1+\phi\eta l_s)\phi\eta l_s Cov\left[\varepsilon_j \log A \mid s_i, \log P_X\right]\right\}.
\end{aligned}$$

By recognizing that $Cov\left[\varepsilon_j \log A \mid s_i, \log P_X\right] = 0$ and substituting in the expressions of $E\left[\log A \mid s_i, \log P_X\right]$, $Var\left[\log A \mid s_i, \log P_X\right]$, and $Var\left[\varepsilon_j \mid s_i, \log P_X\right]$, we can further simplify the expression of $E\left[AX_j^{\phi\eta} \mid s_i, \log P_X\right]$. Equation (A.6) then gives

$$\begin{aligned}
\log X_i = \; & \frac{1}{1-\phi(1-\eta)}\log\phi + \frac{\phi\eta}{1-\phi(1-\eta)}l_0 + \frac{1}{1-\phi(1-\eta)}\left(\phi\eta l_P - 1\right)\log P_X \\
& + \left(\frac{1+\phi\eta l_s}{1-\phi(1-\eta)}\right)\left(\tau_A + \tau_s + \frac{h_A^2}{h_\xi^2}\tau_\xi\right)^{-1}\left(\tau_A\bar{a} + \tau_s s_i + \frac{h_A^2}{h_\xi^2}\tau_\xi \frac{\log P_X - h_0 - h_\xi\bar{\xi}}{h_A}\right) \\
& + \frac{(1+\phi\eta l_s)^2}{2(1-\phi(1-\eta))}\left(\tau_A + \tau_s + \frac{h_A^2}{h_\xi^2}\tau_\xi\right)^{-1} + \frac{\phi^2\eta^2 l_s^2}{2(1-\phi(1-\eta))}\tau_s^{-1}.
\end{aligned}$$

For the above equation to match the conjectured equilibrium position in (A.5), the constant term and the coefficients of s_i and $\log P_X$ have to match. We thus obtain the following equations for determining the coefficients in (A.5):

$$\begin{aligned} l_0 &= \left(\frac{1+\phi\eta l_s}{1-\phi(1-\eta)}\right)\left(\tau_A+\tau_s+\frac{h_A^2}{h_\xi^2}\tau_\xi\right)^{-1}\left(\tau_A\bar{a}-\frac{h_A}{h_\xi^2}\tau_\xi\left(h_0+h_\xi\bar{\xi}\right)\right) \qquad (A.7)\\ &+\frac{(1+\phi\eta l_s)^2}{2(1-\phi(1-\eta))}\left(\tau_A+\tau_s+\frac{h_A^2}{h_\xi^2}\tau_\xi\right)^{-1}\frac{\phi\eta}{1-\phi(1-\eta)}l_0\\ &+\frac{\phi^2\eta^2 l_s^2}{2(1-\phi(1-\eta))}\tau_s^{-1}+\frac{1}{1-\phi(1-\eta)}\log\phi,\\ l_s &= \left(\frac{1+\phi\eta l_s}{1-\phi(1-\eta)}\right)\left(\tau_A+\tau_s+\frac{h_A^2}{h_\xi^2}\tau_\xi\right)^{-1}\tau_s, \qquad (A.8)\\ l_P &= \frac{\phi\eta}{1-\phi(1-\eta)}l_P+\left(\frac{1+\phi\eta l_s}{1-\phi(1-\eta)}\right)\left(\tau_A+\tau_s+\frac{h_A^2}{h_\xi^2}\tau_\xi\right)^{-1}\frac{h_A}{h_\xi^2}\tau_\xi \qquad (A.9)\\ &-\frac{1}{1-\phi(1-\eta)}. \end{aligned}$$

By substituting (A.8) into (A.9), we have

$$l_s = \frac{1+(1-\phi)\,l_P}{1-\phi(1-\eta)}\frac{h_\xi^2}{h_A}\tau_s\tau_\xi^{-1}. \qquad (A.10)$$

By manipulating (A.8), we also have that

$$l_s = \left(\tau_A+\frac{1-\phi}{1-\phi(1-\eta)}\tau_s+\frac{h_A^2}{h_\xi^2}\tau_\xi\right)^{-1}\frac{\tau_s}{1-\phi(1-\eta)}. \qquad (A.11)$$

We now use the market-clearing condition for the commodity market to determine three other equations for the coefficients in the conjectured log-linear commodity price and demand. Aggregating (A.5) gives the aggregate commodity demand of the goods producers:

$$\begin{aligned} \int_{-\infty}^{\infty} X(s_i,P_X)\,d\Phi(\varepsilon_i) &= \int_{-\infty}^{\infty}\exp\left[l_0+l_s s_i+l_P\log P_X\right]d\Phi(\varepsilon_i)\\ &= \int_{-\infty}^{\infty}\exp\left[l_0+l_s(\log A+\varepsilon_i)+l_P(h_0+h_A\log A+h_\xi\xi)\right]d\Phi(\varepsilon_i)\\ &= \exp\left[(l_s+l_P h_A)\log A+l_P h_\xi\xi+l_0+l_P h_0+\frac{1}{2}l_s^2\tau_s^{-1}\right]. \qquad (A.12) \end{aligned}$$

Equation (20) implies that $\log X_S = k\log P_X+\xi$. Thus, the market-clearing condition

$$\log\left[\int_{-\infty}^{\infty}X(s_i,P_X)\,d\Phi(\varepsilon_i)\right] = \log X_S(P_X)$$

requires that the coefficients on $\log A$ and ξ and the constant term be identical on both sides:

$$l_s + l_P h_A = k h_A, \tag{A.13}$$

$$l_P h_\xi = 1 + k h_\xi, \tag{A.14}$$

$$l_0 + l_P h_0 + \frac{1}{2} l_s^2 \tau_s^{-1} = k h_0. \tag{A.15}$$

Equation (A.14) directly implies that

$$l_P = k + h_\xi^{-1}. \tag{A.16}$$

Equations (A.13) and (A.14) together imply that

$$l_s = -h_\xi^{-1} h_A. \tag{A.17}$$

By combining this equation with (A.11), and defining $b = -l_s = h_\xi^{-1} h_A$, we arrive at

$$b^3 + \left(\tau_A + \frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right)\tau_\xi^{-1} b + \frac{\tau_\xi^{-1}\tau_s}{1-\phi(1-\eta)} = 0, \tag{A.18}$$

where b is a real root of a depressed cubic polynomial of the form $x^3 + px + q = 0$, which has one real and two complex roots. As p and q are both positive, the left-hand side (LHS) is monotonically increasing in b while the right-hand side (RHS) is fixed. Thus, the real root b is unique and has to be negative: $b < 0$.

Following Cardano's method, the one real root of equation (A.18) is given by

$$\begin{aligned} b &= \left(\frac{\tau_\xi^{-1}\tau_s}{2(1-\phi(1-\eta))}\right)^{1/3} \sqrt[3]{-1 + \sqrt{1 + \frac{4}{27}\left(\frac{\tau_\xi^{-1}\tau_s}{1-\phi(1-\eta)}\right)^{-2}\left(\tau_A + \frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right)^3}} \\ &+ \left(\frac{\tau_\xi^{-1}\tau_s}{2(1-\phi(1-\eta))}\right)^{1/3} \sqrt[3]{-1 - \sqrt{1 + \frac{4}{27}\left(\frac{\tau_\xi^{-1}\tau_s}{1-\phi(1-\eta)}\right)^{-2}\left(\tau_A + \frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right)^3}} \end{aligned} \tag{A.19}$$

Since $b = h_\xi^{-1} h_A$, we have $h_\xi = b^{-1} h_A$, which together with our expression for l_s and equations (A.10) and (A.16) implies that expressions for h_A and h_ξ be given as in (24) and (25). With h_A and h_ξ determined, l_s is then given by (A.11), l_P by (A.16), h_0 by (A.7) as

$$\begin{aligned} h_0 &= \frac{1}{1+k(1-\phi)}\log\phi - \frac{1-\phi(1-\eta)}{1+k(1-\phi)} b\tau_s^{-1}\left(\tau_A\bar{a} - b\tau_\xi\bar{\xi}\right) \\ &+ \frac{1}{2}\frac{1-\phi(1-\eta)}{1+k(1-\phi)}\left(\left(\frac{1-\phi+\phi^2\eta^2}{1-\phi(1-\eta)} + \phi\eta\right)b - 1\right)\tau_s^{-1} b. \end{aligned} \tag{A.20}$$

and l_0 by equation (A.15) as

$$l_0 = (k - l_P)\, h_0 - \frac{1}{2} l_s^2 \tau_s^{-1}. \tag{A.21}$$

Proof of Proposition 2.3:

We keep the same setting outlined in the main model, except we let A and ξ be observable by all market participants. We first derive the equilibrium. In this setting, each producer's private signal s_i becomes useless as A is directly observable. We can still use equation (22) to derive producer i's optimal commodity demand. As the producers now share the same information about A, they must have the same expectation about their future trading partners' production decisions. As a result, $X_i = X_j$ for any i and j. Equation (22) therefore implies that in equilibrium $X_i = \left(\frac{\phi A}{P_X}\right)^{\frac{1}{1-\phi}}$.

Market clearing of the commodity market requires that the producers' aggregate demand equals the commodity supply, that is, $X_i = X_S$. From equation (20), we must have that $\log X_i = k \log P_X + \xi$. We then obtain $\log P_X$ and $\log X_i$ stated in Proposition 3. It is clear that this equilibrium is unique.

Proof of Proposition 2.4:

As $\tau_s \to \infty$, equation (A.18) implies that b goes to $-\frac{1}{1-\phi}$. Consequently, as $\tau_s \to \infty$, equation (24) gives that $h_A \to \frac{1}{1+k(1-\phi)}$, and equation (25) gives that $h_\xi \to -\frac{1-\phi}{1+k(1-\phi)}$. Therefore, both h_A and h_ξ converge to their corresponding values in the perfect-information benchmark.

That $|h_\xi|$ is larger than it is in the perfect-information benchmark is apparent since the numerator of $|h_\xi|$ in equation (25) is positive and larger than $1-\phi$. That h_A is lower follows by substituting equation (A.18) into equation (24) to arrive at

$$h_A = \frac{1 + \tau_A \tau_s^{-1} (1 - \phi (1 - \eta))\, b}{1 + k (1 - \phi)}.$$

Since $b < 0$, it follows that $h_A < \frac{1}{1+k(1-\phi)}$, which is the value of h_A in the perfect-information benchmark.

Proof of Proposition 2.5:

As $l_s = -h_\xi^{-1} h_A$ from (A.17), $\pi = \frac{h_A^2/\tau_A}{h_\xi^2/\tau_\xi} = l_s^2 \frac{\tau_\xi}{\tau_A}$. Since $l_s > 0$, it is sufficient to study the behavior of how l_s varies with τ_s and η to understand how π changes with τ_s and η. To see that l_s is monotonically increasing in τ_s, we note that $l_s = -b$, where b is the only real and negative root of equation (A.18). Then, by the Implicit Function Theorem it is apparent that

$$\frac{\partial b}{\partial \tau_s} = -\frac{\frac{1-\phi}{1-\phi(1-\eta)}\tau_\xi \tau_s b + \frac{\tau_s \tau_\xi}{1-\phi(1-\eta)}}{3b^2 + \left(\tau_A + \frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right)\tau_\xi^{-1}}\tau_s^{-1} = \frac{b^3 + \tau_A \tau_\xi^{-1} b}{3b^2 + \left(\tau_A + \frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right)\tau_\xi^{-1}}\tau_s^{-1} < 0.$$

Similarly, we have

$$\frac{\partial b}{\partial \eta} = \phi \tau_s \tau_\xi^{-1} \frac{\left(1-(1-\phi)\,l_s\right)\frac{1}{(1-\phi(1-\eta))^2}}{3l_s^2 + \left(\tau_A + \frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right)\tau_\xi^{-1}} > 0.$$

Thus, l_s is increasing in τ_s and decreasing in η, which in turn implies that π is increasing in τ_s and decreasing in η.

To analyze the dependence of π on τ_ξ, we have

$$\frac{\partial \pi}{\partial \tau_\xi} = l_s^2 \frac{1}{\tau_A} + 2l_s \frac{\tau_\xi}{\tau_A}\frac{\partial l_s}{\partial \tau_\xi} = \frac{1}{\tau_A} b \left(b + 2\tau_\xi \frac{\partial b}{\partial \tau_\xi}\right).$$

By applying the Implicit Function Theorem again, we obtain

$$\frac{\partial b}{\partial \tau_\xi} = \frac{\left(\tau_A + \frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right)\tau_\xi^{-1} b + \frac{\tau_\xi^{-1}\tau_s}{1-\phi(1-\eta)}}{3b^2 + \left(\tau_A + \frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right)\tau_\xi^{-1}}\tau_\xi^{-1} = \frac{-b^3}{3b^2 + \left(\tau_A + \frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right)\tau_\xi^{-1}}\tau_\xi^{-1} > 0.$$

By substituting this into the above expression for $\frac{\partial \pi}{\partial \tau_\xi}$, we find that

$$\frac{\partial \pi}{\partial \tau_\xi} = \frac{b^2}{\tau_A}\left(\frac{b^2 + \left(\tau_A + \frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right)\tau_\xi^{-1}}{3b^2 + \left(\tau_A + \frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right)\tau_\xi^{-1}}\right) > 0.$$

Therefore, π is monotonically increasing in τ_ξ.

Proof of Proposition 2.6:

Based on l_P and h_ξ given in equations (A.16) and (25), $l_P > 0$ is equivalent to $b^2 > \frac{k^{-1}\tau_\xi^{-1}\tau_s}{1-\phi(1-\eta)}$, which, as $l_s = -b > 0$, is in turn equivalent to $l_s > l_s^* = \sqrt{\frac{k^{-1}\tau_\xi^{-1}\tau_s}{1-\phi(1-\eta)}}$. In words, this condition states that the commodity price has to be sufficiently informative. As b is the unique real and negative root of equation (A.18), this condition is equivalent to the following condition on the LHS of equation (A.18): $LHS(-l_s^*) > 0$. By substituting l_s^* into the LHS, we obtain the following condition:

$$-\frac{k^{-3/2}\tau_\xi^{-1}\tau_s}{1-\phi\,(1-\eta)} - \left(\tau_A + \frac{1-\phi}{1-\phi\,(1-\eta)}\tau_s\right)\tau_\xi^{-1}k^{-1/2} + \sqrt{\frac{\tau_\xi^{-1}\tau_s}{1-\phi\,(1-\eta)}} > 0,$$

which, as $1-\phi(1-\eta)>0$ and defining $u=\sqrt{1-\phi(1-\eta)}$, can be rewritten as

$$u^2 - u\tau_A^{-1}\sqrt{k\tau_\xi\tau_s} + \left(1-\phi+k^{-1}\right)\tau_A^{-1}\tau_s < 0.$$

Note that the LHS of this condition $LHS(u)$ is a quadratic form of u, which has its minimum at $u^* = \frac{1}{2\tau_A}\sqrt{k\tau_\xi\tau_s}$. Thus, this condition is satisfied if and only if the following occurs. First, $LHS(u^*)<0$, which is equivalent to $\tau_\xi/\tau_A > 4k^{-1}\left(1-\phi+k^{-1}\right)$, the first condition given in Proposition 6. Second,

$$LHS(u) = (u-u^*)^2 - \left[(u^*)^2 - \left(1-\phi+k^{-1}\right)\tau_A^{-1}\tau_s\right] < 0,$$

which is equivalent to

$$\frac{u - \frac{1}{2\tau_A}\sqrt{k\tau_\xi\tau_s}}{\frac{1}{2}\tau_s^{1/2}\tau_A^{-1}\sqrt{k\tau_\xi - 4\left(1-\phi+k^{-1}\right)\tau_A}} \in (-1,1).$$

This leads to the second condition given in Proposition 6.

Proof of Proposition 2.7:

We begin by evaluating the first component of the social welfare from the island households' goods consumption. We denote this component by

$$W^C = E\left[\int_0^1 \left(\frac{C(i)}{1-\eta}\right)^{1-\eta}\left(\frac{C^*(i)}{\eta}\right)^{\eta} di\right].$$

In the perfect-information benchmark, by substituting the symmetric consumption of all island households, the expected social welfare from consumption is

$$\log W^C_{bench} = \log E\left[\int_0^1 AX_i^\phi di\right] = \log E\left[AX_i^\phi\right].$$

Given $\log X_i$ derived in Proposition 3, we have

$$\begin{aligned}\log W^C_{bench} &= \frac{\phi k}{1+k(1-\phi)}\log\phi + \frac{1+k}{1+k(1-\phi)}\bar{a} + \frac{1}{2}\left(\frac{1+k}{1+k(1-\phi)}\right)^2\tau_A^{-1} \\ &\quad + \frac{\phi}{1+k(1-\phi)}\bar{\xi} + \frac{1}{2}\left(\frac{\phi}{1+k(1-\phi)}\right)^2\tau_\xi^{-1}.\end{aligned}$$

Note that the total goods output in this economy is given by

$$E\left[Y^{aggr}_{bench}\right] = E\left[\int_0^1 Y_i di\right] = E\left[\int_0^1 AX_i^\phi di\right] = W^C_{bench},$$

which indicates that in this symmetric equilibrium with perfect information, the expected social welfare from consumption is equal to the expected aggregate goods output.

In the presence of informational frictions, by using Proposition 1, the expected social welfare from consumption is given by

$$\log W^C = \log E\left[A\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} X\left(s_i,P_X\right)^{\phi(1-\eta)} X\left(s_j,P_X\right)^{\phi\eta}\, d\Phi\left(\varepsilon_i\right) d\Phi\left(\varepsilon_j\right)\right],$$

where in the second line, an integral over ε_j, that is, noise in the signal of the goods producer of island j, is taken to compute expectation over uncertainty in ε_j. By substituting

$$\log X\left(\varepsilon_i\right) = l_0 + l_P \log P_X + l_s s_i = l_0 + l_P \log P_X + l_s\left(\log A + \varepsilon_i\right)$$

and $\log P_X = h_0 + h_A \log A + h_\xi \xi$ with our expressions for l_s, l_P, h_ξ, h_A, and b^3 from equation (A.18), we obtain

$$\begin{aligned}\log W^C \;=\;& \frac{\phi k}{1+k(1-\phi)}\log\phi + \frac{1+k}{1+k(1-\phi)}\bar{a} + \frac{1}{2}\left(1+\phi k h_A\right)^2\tau_A^{-1} + \frac{\phi}{1+k(1-\phi)}\bar{\xi}\\ &+\frac{1}{2}\phi^2\left(1+kb^{-1}h_A\right)^2\tau_\xi^{-1} + \frac{1}{2}\left(1-\frac{1}{\phi}-2\eta(1-\eta)\right)\phi^2 b^2\tau_s^{-1}\\ &-\frac{1}{2}\frac{\phi k\left(1-\phi(1-\eta)\right)}{1+k(1-\phi)}\left(1-\left(\frac{1-\phi+\phi^2\eta^2}{1-\phi(1-\eta)}+\phi\eta\right)b\right)b\tau_s^{-1}.\end{aligned}$$

The logarithm of the expected total output in this economy is given by

$$\begin{aligned}\log E\left[Y^{aggr}\right] \;=\;& \log E\left[\int_{-\infty}^{\infty} Y_i di\right] = \log E\left[A\int_{-\infty}^{\infty} X\left(\varepsilon_i\right)^{\phi} d\Phi\left(\varepsilon_i\right)\right]\\ =\;& \log E\left[e^{\phi l_0+\phi l_P\log P_X+(1+\phi l_s)\log A}\int_{-\infty}^{\infty} e^{\phi l_s\varepsilon_i} d\Phi\left(\varepsilon_i\right)\right]\\ =\;& \phi l_0 + \frac{1}{2}\phi^2 l_s^2\tau_s^{-1} + \phi l_P h_0 + \phi l_P h_\xi\bar{\xi} + \frac{1}{2}\phi^2 l_P^2 h_\xi^2\tau_\xi^{-1} + \left(1+\phi l_s+\phi l_P h_A\right)\bar{a}\\ &+\frac{1}{2}\left(1+\phi l_s+\phi l_P h_A\right)^2\tau_A^{-1}.\end{aligned}$$

Again by substituting the expressions for l_s, l_P, h_ξ, and h_A, we have

$$\begin{aligned}\log E\left[Y^{aggr}\right] \;=\;& \frac{\phi k}{1+k(1-\phi)}\log\phi + \left(\frac{1+k}{1+k(1-\phi)}\right)\bar{a} + \frac{1}{2}\left(1+\phi k h_A\right)^2\tau_A^{-1}\\ &+\phi\left(\frac{1}{1+k(1-\phi)}\right)\bar{\xi} + \frac{1}{2}\phi^2\left(1+kb^{-1}h_A\right)^2\tau_\xi^{-1} - \frac{1}{2}\phi(1-\phi)b^2\tau_s^{-1}\\ &-\frac{1}{2}\frac{\phi k\left(1-\phi(1-\eta)\right)}{1+k(1-\phi)}\left(1-\left(\frac{1-\phi+\phi^2\eta^2}{1-\phi(1-\eta)}+\phi\eta\right)b\right)b\tau_s^{-1}.\end{aligned}$$

It is then easy to compute

$$\log E\left[Y^{aggr}\right] - \log W = 2\phi^2\eta(1-\eta)b^2\tau_s^{-1} > 0.$$

We now compare expected aggregate goods output with and without informational frictions:

$$\begin{aligned}
& \log E\left[Y^{aggr}\right] - \log E\left[Y^{aggr}_{bench}\right] \\
= \; & \frac{1}{2}\left(\left(1+\phi k h_A\right)^2 - \left(1+\frac{\phi k}{1+k\left(1-\phi\right)}\right)^2\right)\tau_A^{-1} \\
& + \frac{1}{2}\phi^2\left(\left(1+kb^{-1}h_A\right)^2 - \left(1-k\frac{1-\phi}{1+k\left(1-\phi\right)}\right)^2\right)\tau_\xi^{-1} - \frac{1}{2}\phi\left(1-\phi\right)b^2\tau_s^{-1} \\
& - \frac{1}{2}\frac{\phi k\left(1-\phi\left(1-\eta\right)\right)}{1+k\left(1-\phi\right)}\left(1-\left(1-\phi+\phi\eta\left(1+\phi\eta\right)\right)b\right)b\tau_s^{-1}.
\end{aligned}$$

Substituting with equations (24) and (A.18), we arrive at

$$\begin{aligned}
& \log E\left[Y^{aggr}\right] - \log E\left[Y^{aggr}_{bench}\right] \\
= \; & \frac{1}{2}\frac{\phi k\left(1-\phi\left(1-\eta\right)\right)\tau_s^{-1}b}{\left(1+k\left(1-\phi\right)\right)^2}\left(1+k-\phi b\right) - \frac{1}{2}\phi\left(1-\phi\right)b^2\tau_s^{-1} \\
& + \frac{1}{2}\frac{\phi k\left(1-\phi\left(1-\eta\right)\right)\tau_s^{-1}b^2}{1+k\left(1-\phi\right)}\left(\frac{1-\phi}{1-\phi\left(1-\eta\right)} + \phi^2\eta^2 + \left(\frac{\phi^2\eta^2}{1-\phi\left(1-\eta\right)} - 1\right)\phi\left(1-\eta\right)\right).
\end{aligned}$$

Notice that $\frac{\phi^2\eta^2}{1-\phi(1-\eta)} < 1$ and the first term is negative since $b<0$. Further note that

$$\begin{aligned}
& \frac{1}{2}\frac{\phi k\left(1-\phi\left(1-\eta\right)\right)\tau_s^{-1}b^2}{1+k\left(1-\phi\right)}\frac{1-\phi}{1-\phi\left(1-\eta\right)} - \frac{1}{2}\phi\left(1-\phi\right)b^2\tau_s^{-1} \\
= \; & -\left(\frac{1-k\phi}{1+k\left(1-\phi\right)}\right)\frac{1}{2}\phi\left(1-\phi\right)b^2\tau_s^{-1} < 0,
\end{aligned}$$

and

$$\frac{1}{2}\frac{\phi k\left(1-\phi\left(1-\eta\right)\right)\tau_s^{-1}b^2}{1+k\left(1-\phi\right)}\left(\phi^2\eta^2 - \frac{\phi-\left(1+k\right)b^{-1}}{1+k\left(1-\phi\right)}\right) < 0,$$

because $\phi^2\eta^2 < \phi$ and, since $\phi < 1$ and $0 > b > -\frac{1}{1-\phi}$, $k\left(1-\phi\right)\phi^2\eta^2 + \left(1+k\right)b^{-1} < 0$. To see that $b > -\frac{1}{1-\phi}$, we rewrite equation (A.18) as

$$b^3 + \tau_A\tau_\xi^{-1}b + \left(\left(1-\phi\right)b+1\right)\frac{\tau_s\tau_\xi^{-1}}{1-\phi\left(1-\eta\right)} = 0,$$

from which it follows that $b > -\frac{1}{1-\phi}$. Therefore, we see that

$$\log E\left[Y^{aggr}\right] - \log E\left[Y^{aggr}_{bench}\right] < 0.$$

Given that the expected social welfare from consumption W^C_{bench} is equal to the expected aggregate output $E\left[Y^{aggr}_{bench}\right]$ in the perfect-information benchmark, and in the presence of informational frictions the expected social welfare from consumption W^C is strictly less than the expected aggregate goods output $E\left[Y^{aggr}\right]$, the expected social welfare from consumption is lower in the presence of information frictions than in the perfect-information benchmark.

Now we return to the second part of the expected social welfare from commodity suppliers' disutility of labor. We denote this part by

$$W^L = E\left[\frac{k}{1+k}e^{-\xi/k}X_S^{\frac{1+k}{k}}\right].$$

In the perfect-information benchmark, by using $\log X_S$ derived in Proposition 3, we have

$$\begin{aligned}\log W_{bench}^L &= \log\frac{k}{1+k} + \frac{1+k}{1+k(1-\phi)}\log\phi + \frac{1+k}{1+k(1-\phi)}\bar{a} + \frac{\phi}{1+k(1-\phi)}\bar{\xi}\\ &\quad + \frac{1}{2}\left(\frac{1+k}{1+k(1-\phi)}\right)^2\tau_A^{-1} + \frac{1}{2}\left(\frac{\phi}{1+k(1-\phi)}\right)^2\tau_\xi^{-1}.\end{aligned}$$

In the presence of informational frictions, aggregate demand X_S is given by

$$\log X_S = k\log P_X + \xi = kh_A\log A + (kh_\xi + 1)\,\xi + kh_0,$$

and therefore the suppliers' disutility of labor reduces to

$$\begin{aligned}\log W^L &= \log\frac{k}{1+k} + (1+k)\,h_0 + (1+k)\,h_A\bar{a} + (1+(1+k)\,h_\xi)\,\bar{\xi}\\ &\quad + \frac{1}{2}(1+k)^2 h_A^2\tau_A^{-1} + \frac{1}{2}(1+(1+k)\,h_\xi)^2\tau_\xi^{-1}.\end{aligned}$$

We now analyze the overall social welfare $W = W^C - W^L$. We can express the relative welfare in the two economies as

$$\frac{W}{W_{bench}} = \frac{W^C - W^L}{W_{bench}^C - W_{bench}^L} = \frac{W^C}{W_{bench}^C}\frac{1-W^L/W^C}{1-W_{bench}^L/W_{bench}^C} < \frac{1-W^L/W^C}{1-W_{bench}^L/W_{bench}^C},$$

where the last inequality follows from $W^C < W_{bench}^C$, as proved above.

Note that in the perfect-information benchmark,

$$\log W_{bench}^L - \log W_{bench}^C = \log\frac{\phi k}{1+k}.$$

Thus, $1 - W_{bench}^L/W_{bench}^C = 1 - \frac{\phi k}{1+k} > 0$. Therefore, it is sufficient to show that $W^L/W^C \geq W_{bench}^L/W_{bench}^C$ to establish that $\frac{W}{W_{bench}} < 1$.

With some manipulation of our expressions for $\log W^L$ and $\log W^C$, and by substituting our expressions for h_A and h_ξ and making use of equation (A.18), we arrive at

$$\begin{aligned}&\log\left(W^L/W^C\right)\\ =\;& \log\frac{\phi k}{1+k} + \frac{1}{2}b\tau_s^{-1}\left(\left(1-\phi^2+\phi^2\eta^2+\phi\eta+\phi^2\eta\,(1-\eta)\right)b - (1-\phi\,(1-\eta))\right)\\ &- (1-\phi\,(1-\eta))\,\tau_s^{-1}b^2 - \frac{1}{2}(1-\phi\,(1-\eta))^2\,\tau_s^{-2}b^2\left(\tau_A + \tau_\xi b^2\right).\end{aligned}$$

Finally, by invoking equation (A.18) to rewrite the last term, we find that

$$\log\left(W^L/W^C\right) = \log \frac{\phi k}{1+k}.$$

Thus, $\log\left(W^L/W^C\right) = \log\left(W^L_{bench}/W^C_{bench}\right)$, which in turn establishes the proposition.

2.B: Model Extension

2.B.1 Model Setting

We introduce a new date $t = 0$ before the dates $t = 1$ and 2 in the baseline model, and a centralized futures market at $t = 0$ for delivery of the commodity at $t = 1$. All agents can take positions in the futures market at $t = 0$, and can choose to revise or unwind their positions before delivery at $t = 1$. The ability to unwind positions before delivery is an advantage that makes futures market trading appealing in practice.

We keep all of the agents in the baseline model: island households, goods producers, and commodity suppliers and add a group of financial traders. These traders invest in the commodity by taking a long position in the futures market at $t = 0$ and then unwinding this position at $t = 1$ without taking delivery.

To focus on information aggregation through trading in the futures market, we assume that there is no spot market trading at $t = 0$. At $t = 1$, a spot market naturally emerges through commodity delivery for the futures market. Commodity suppliers take a short position in the futures market at $t = 0$ and then make delivery at $t = 1$. Suppliers' marginal cost of supplying the commodity determines the spot price. When a trader chooses to unwind a futures position at $t = 1$, his gain/loss is determined by this spot price.

Table 2B.1

Timeline of the Extended Model

	t=0	t=1	t=2
	Futures Market	Spot Market	Goods Market
Households			Trade/Consume Goods
Producers	Observe Signals Long Futures	Take Delivery Produce Goods	
Suppliers	Short Futures	Observe Supply Shock Deliver Commodity	
Fin Traders	Long/Short Futures	Unwind Position	

Table 2B.1 specifies the timeline of the extended model. We keep the same specification for the island households, who trade and consume both home and away goods at $t = 2$ as described in Section I.A of the main paper. We modify some of the specifications for goods producers and commodity suppliers and describe our specifications for financial traders below.

2.B.1.1 Goods Producers

As in the main model, we allow goods producers to have the same production technology and receive their private signals at $t = 0$. Each producer takes a long position in the futures market at $t = 0$ and then commodity delivery at $t = 1$. The timing of the producer's information flow is key to our analysis. At $t = 0$, producer i's information set $I_i^0 = \{s_i, F\}$ includes its private signal s_i and the traded futures price F. At $t = 1$, its information set $I_i^1 = \{s_i, F, P_X\}$ includes the updated spot price P_X.

We allow the producer to use its updated information set at $t = 1$ to revise its futures position for commodity delivery. That is, its production decision is based on not only its private signal and the futures price but also the updated spot price. Thus, it is not obvious that noise in the futures market can affect the producer's production decision and commodity demand. We examine this key issue with our extended model.

At $t = 1$, the producer optimizes its production decision X_i (i.e., commodity demand) based on its updated information set I_i^1:

$$\max_{X_i} \; E\left[P_i Y_i \middle| \, \mathcal{I}_i^1 \right] - P_X X_i + (P_X - F)\tilde{X}_i.$$

The first two terms above represent the producer's expected profit from goods production and the last term is the gain/loss from its futures position. The producer's optimal production decision is then

$$X_i = \left\{ \phi E\left[A X_j^{\phi\eta} \middle| \mathcal{I}_i^1 \right] \Big/ P_X \right\}^{1/(1-\phi(1-\eta))}. \tag{2.B.1}$$

When deciding its futures position at $t = 0$, the producer faces a nuanced issue in that, because it does not need to commit its later production decision to the initial futures position, it may engage in dynamic trading. In other words, it could choose a futures position to maximize its expected trading profit at $t = 0$. This trading motive is not essential for our focus on analyzing the aggregation of the producers' information but significantly complicates derivation of the futures market equilibrium. To avoid this complication, we make a simplifying assumption that the producers are myopic at $t = 0$. That is, at $t = 0$, each producer chooses a futures position as if it commits to taking full delivery and using the good for production, even though the producer can revise its production decision based on the updated information at $t = 1$. While this simplifying assumption affects each producer's trading profit, it is innocuous for our analysis of how the futures price feeds back to the producers' later production decisions because each producer still makes good use of its information and the futures price is informative by aggregating each producer's information.

Specifically, at $t = 0$ the producer chooses a futures position $\tilde{X}_i$ to maximize the following expected production profit based on its information set I_i^0:

$$\max_{\tilde{X}_i} E\left[P_i Y_i \middle| \mathcal{I}_i^0\right] - F\tilde{X}_i,$$

where it treats $\tilde{X}_i$ as its production input at $t = 1$. Throughout the rest of this appendix, we use a tilde to denote variables and coefficients associated with the futures market at $t = 0$; we maintain the same notation without the tilde for variables related to the spot market at $t = 1$. The producer's futures position is then

$$\tilde{X}_i = \left\{ \phi E\left[A\tilde{X}_j^{\phi\eta} \middle| \mathcal{I}_i^0 \right] \Big/ F \right\}^{1/(1-\phi(1-\eta))}. \tag{2.B.2}$$

2.B.1.2 Financial Traders

We introduce a group of financial traders, who trade in the futures market at $t = 0$ and unwind their position at $t = 1$ before delivery. For simplicity, we assume that the aggregate position of financial traders and goods producers is given by the aggregate position of producers multiplied by a factor $e^{\kappa \log A + \theta}$:

$$e^{\kappa \log A + \theta} \int_{-\infty}^{\infty} \tilde{X}_i\left(s_i, F\right) d\Phi\left(\varepsilon_i\right),$$

where the factor $e^{\kappa \log A + \theta}$ represents the contribution of financial traders. This multiplicative specification is useful for ensuring the tractable log-linear equilibrium of our model.[56]

[56] From an economic perspective, this specification implies that the position of financial traders tends to expand and contract with producers' futures position, which is broadly consistent with the expansion and contraction of the aggregate commodity futures positions of portfolio investors and hedge funds in the recent commodity price boom-and-bust cycle (e.g., Cheng, Kirilenko, and Xiong (2012)). Also note that $e^{\kappa \log A + \theta}$ can be less than one. This implies that financial traders may take a net short position at some point, which is consistent with short positions taken by hedge funds in practice.

We allow the contribution of financial traders $e^{\kappa \log A+\theta}$ to contain a component $\kappa \log A$, where $\kappa > 0$, to capture the possibility that the trading of financial traders is partially driven by their knowledge of the global fundamental $\log A$.

The trading of financial traders also contains a random component θ, which is unobservable by other market participants. This assumption is realistic in two respects. First, in practice, the trading of financial traders is often driven by portfolio diversification and risk-control purposes unrelated to fundamentals of commodity markets. Second, market participants cannot directly observe others' positions.[57] Specifically, we assume that θ has a normal distribution independent of other sources of uncertainty in the model,

$$\theta \backsim \mathcal{N}\left(\overline{\theta}, \tau_\theta^{-1}\right),$$

with mean $\overline{\theta}$ and variance τ_θ^{-1}.

The presence of financial traders introduces an additional source of uncertainty to the futures market, as both goods producers and commodity suppliers cannot observe θ at $t = 0$. At $t = 1$, financial traders unwind their positions, and commodity suppliers make delivery only to goods producers.

2.B.1.3 Commodity Suppliers

Commodity suppliers take a short position of $\tilde{X}_S$ in the futures market at $t = 0$ and then make delivery of X_S units of the commodity at $t = 1$. We maintain the same convex cost function for the suppliers: $\frac{k}{1+k} e^{-\xi/k} \left(X_S\right)^{\frac{1+k}{k}}$, where the supply shock ξ has a Gaussian distribution $N\left(\bar{\xi}, \tau_\xi^{-1}\right)$.

We assume that the suppliers observe their supply shock ξ only at $t = 1$, which implies that the supply shock does not affect the futures price at $t = 0$ and instead hits the spot market at $t = 1$. Due to this timing, the supply shock provides a camouflage for the unwinding of financial traders' aggregate futures position at $t = 1$. That is, even after financial

[57] Despite the fact that large traders need to report their futures positions to the Commodities Future Trading Commission (CFTC) on a daily basis, ambiguity in trader classification and netting of positions taken by traders who are involved in different lines of business nevertheless make the aggregate positions provided by the CFTC's weekly Commitment of Traders Report to the public imprecise. See Cheng, Kirilenko, and Xiong (2012) for a more detailed discussion of the trader classification and netting problems in the CFTC's Large Trader Reporting System and a summary of positions taken by commodity index traders and hedge funds.

traders unwind their position, the commodity spot price does not reveal their position.[58]

In summary, the suppliers' information set at $t = 0$ is $I_S^0 = \{F\}$, and at $t = 1$ is $I_S^1 = \{F, P_X, \xi\}$. At $t = 1$, the suppliers face the following optimization problem:

$$\max_{X_S} P_X X_S - \frac{k}{1+k} e^{-\xi/k} X_S^{\frac{1+k}{k}} + (F - P_X)\, \tilde{X}_S,$$

where they choose X_S—the quantity of commodity delivery—to maximize the profit from delivery in the first two terms. The last term is the gain/loss from their initial futures position. The suppliers' optimal supply curve is then given by $X_S = e^{\xi} P_X^k$, which is identical to their supply curve in the baseline model.

At $t = 0$, like the goods producers, the suppliers also face a nuanced issue related to dynamic trading. As their initial futures position does not necessarily equal their later commodity delivery, they may also choose to maximize the trading profit from $t = 0$ to $t = 1$. To be consistent with our earlier assumption about the myopic behavior of goods producers, we assume that at $t = 0$ the suppliers believe that goods producers will take full delivery of their futures positions and that the suppliers choose their initial short position to myopically maximize the profit from making delivery of $e^{-(\kappa \log A + \theta)} \tilde{X}_S$ units of the commodity to goods producers:

$$\max_{\tilde{X}_S} E\left[F e^{-(\kappa \log A + \theta)} \tilde{X}_S \middle| \; \mathcal{I}_S^0 \right] - E\left[\frac{k}{1+k} e^{-\xi/k} \left(e^{-(\kappa \log A + \theta)} \tilde{X}_S \right)^{\frac{1+k}{k}} \middle| \; \mathcal{I}_S^0 \right].$$

Since ξ is independent of θ and $\log A$, it is easy to derive

$$\tilde{X}_S = e^{\bar{\xi} - \sigma_\xi^2/2k} \left\{ E\left[e^{-(\kappa \log A + \theta)} \middle| \; \mathcal{I}_S^0 \right] / E\left[e^{-\frac{1+k}{k}(\kappa \log A + \theta)} \middle| \; \mathcal{I}_S^0 \right] \right\}^k F^k, \tag{2.B.3}$$

which is a function of the futures price F.

2.B.1.4 Joint Equilibrium of Different Markets

We analyze the joint equilibrium of a number of markets: the goods markets between each pair of matched islands at $t = 2$, the spot market for the commodity at $t = 1$, and the futures market at $t = 0$. Equilibrium requires clearing of each of these markets:

[58] This timing may appear special in our static setting with only one round of futures market trading followed by physical commodity delivery, as there is no particular reason to argue whether letting the suppliers observe the supply shock at $t = 0$ or $t = 1$ is more natural. However, if we view this setting as one module of a more realistic setting with many recurrent periods and a supply shock arriving in each period, then there is always a supply shock hitting the market when financial traders unwind their futures position.

- At $t = 2$, for each pair of randomly matched islands $\{i, j\}$, the households of these islands trade their produced goods and clear the market of each good:

$$\begin{aligned} C_i + C_j^* &= AX_i^{\phi}, \\ C_i^* + C_j &= AX_j^{\phi}. \end{aligned}$$

- At $t = 1$, the commodity supply equals the goods producers' aggregate demand:

$$\int_{-\infty}^{\infty} X(s_i, F, P_X)\, d\Phi(\varepsilon_i) = X_S(P_X, \xi).$$

- At $t = 0$, the futures market clears:

$$e^{\kappa \log A + \theta} \int_{-\infty}^{\infty} \tilde{X}_i(s_i, F)\, d\Phi(\varepsilon_i) = \tilde{X}_S(F).$$

2.B.2 The Equilibrium

The goods market equilibrium at $t = 2$ remains identical to that derived in Proposition 1 for the main model. The futures market equilibrium at $t = 0$ and the spot market equilibrium at $t = 1$ also remain log-linear and can be derived following a similar procedure as the derivation of Proposition 2. The following proposition summarizes the key features of the equilibrium with explicit expressions for all coefficients given in Section 2.B.4.

PROPOSITION 2.B1: At $t = 0$, the futures market has a unique log-linear equilibrium: the futures price is a log-linear function of $\log A$ and θ,

$$\log F = \tilde{h}_A \log A + \tilde{h}_\theta \theta + \tilde{h}_0, \tag{2.B.4}$$

with the coefficients $\tilde{h}_A > 0$ and $\tilde{h}_\theta > 0$, while the long position taken by goods producer i is a log-linear function of its private signal s_i and $\log F$,

$$\log \tilde{X}_i = \tilde{l}_s s_i + \tilde{l}_F \log F + \tilde{l}_0, \tag{2.B.5}$$

with the coefficient $\tilde{l}_s > 0$.

At $t = 1$, the spot market also has a unique log-linear equilibrium: the spot price of the commodity is a log-linear function of $\log A$, $\log F$, and ξ,

$$\log P_X = h_A \log A + h_F \log F + h_\xi \xi + h_0, \tag{2.B.6}$$

with the coefficients $h_A > 0$, $h_F > 0$, and $h_\xi < 0$, while the commodity consumed by producer i is a log-linear function of s_i, $\log F$, and $\log P_X$,

$$\log X_i = l_s s_i + l_F \log F + l_P \log P_X + l_0, \tag{2.B.7}$$

with the coefficients $l_s > 0$ and $l_F > 0$, and the sign of l_P undetermined.

There are two rounds of information aggregation in the equilibrium. During the first round of trading in the futures market at $t = 0$, goods producers take long positions based on their private signals. The futures price $\log F$ aggregates producers' information, and reflects a linear combination of $\log A$ and θ, as given in (28). The futures price does not fully reveal $\log A$ due to the θ noise originated from the trading of financial traders. The spot price that emerges from the commodity delivery at $t = 1$ represents another round of information aggregation by pooling together the goods producers' demand for delivery. As a result of the arrival of the supply shock ξ, the spot price $\log P_X$ does not fully reveal either $\log A$ or θ, and instead reflects a linear combination of $\log A$ and ξ, as derived in (29).

Despite the updated information from the spot price at $t = 1$, the informational content of $\log F$ is not subsumed by the spot price, and still has an influence on goods producers' expectations of $\log A$. As a result of this informational role, equation (30) confirms that each goods producer's commodity demand at $t = 1$ is increasing with $\log F$, as $l_F > 0$, and equation (29) shows that the spot price is also increasing with $\log F$, as $h_F > 0$. This is the key feedback channel through which futures market trading affects commodity demand and the spot price despite the availability of information from the spot price.

The simplifying assumptions we make regarding the myopic trading of goods producers and commodity suppliers at $t = 0$ are innocuous to the informational role of the futures price at $t = 1$. As long as goods producers trade on their private signals, the futures price would aggregate the information, which in turn establishes the futures price as a useful price signal for the later round at $t = 1$. Our simplifying assumptions have quantitative consequences for goods producers' trading profits and the efficiency of the futures price signal, but should not critically affect the qualitative feedback channel of the futures price, which we characterize in the next subsection.[59]

Interestingly, Proposition 2.B1 also reveals that l_P can be either positive or negative,

[59] Note that despite the different information content of the futures price and the spot price, there is no arbitrage between the two prices because the two prices are traded at different points in time and the spot price is exposed to the supply shock realized later.

due to the offsetting cost effect and informational effect of the spot price, similar to our characterization of the main model.

2.B.3. Implications

2.B.3.1 Feedback on Commodity Demand

As financial traders do not take or make any physical delivery, their trading in the futures market does not have direct effect on commodity supply or demand. However, their trading affects the futures price, through which it can further impact commodity demand and spot prices. By substituting equation (28) into (29), we express the spot price $\log P_X$ as a linear combination of primitive variables $\log A$, θ, and ξ:

$$\log P_X = \left(h_A + h_F \tilde{h}_A\right) \log A + h_F \tilde{h}_\theta \theta + h_\xi \xi + h_F \tilde{h}_0 + h_0. \tag{2.B.8}$$

The θ term arises through the futures price. As $h_F > 0$ and $\tilde{h}_\theta > 0$, the noise from financial traders' trading in the futures market, θ, has a positive effect on the spot price.

Furthermore, by substituting the equation above and (28) into (30), we obtain an individual producer's commodity demand as

$$\begin{aligned} \log X_i &= l_s s_i + \left(l_F \tilde{h}_A + l_P \left(h_A + h_F \tilde{h}_A\right)\right) \log A + \left(l_F + l_P h_F\right) \tilde{h}_\theta \theta + l_P h_\xi \xi \\ &\quad + \left(l_F + l_P h_F\right) \tilde{h}_0 + l_P h_0 + l_0, \end{aligned}$$

and the producers' aggregate demand as

$$\begin{aligned} \log\left[\int_{-\infty}^{\infty} X\left(s_i, F, P_X\right) d\Phi\left(\varepsilon_i\right)\right] &= \left[l_s + l_P h_A + l_F \tilde{h}_A + l_P h_F \tilde{h}_A\right] \log A + \left(l_F + l_P h_F\right) \tilde{h}_\theta \theta + l_P h_\xi \xi \\ &\quad + \left(l_F + l_P h_F\right) \tilde{h}_0 + l_P h_0 + l_0 + \frac{1}{2} l_s^2 \tau_s^{-1}. \end{aligned}$$

By using equation (2.B.28) in the proof of Proposition 2.B1, the coefficient on θ in the aggregate commodity demand is

$$l_F + l_P h_F = k h_F > 0.$$

Thus, θ also has a positive effect on aggregate commodity demand.

The effects of θ on commodity demand and the spot price clarify the simple yet important conceptual point that traders in commodity futures markets, who never take or make physical commodity delivery, can nevertheless impact commodity markets through the informational feedback channel of commodity futures prices.

2.B.3.2 Market Transparency

Information frictions in the futures market, originating from the unobservability of the positions of different participants, are essential in order for the trading of financial traders to impact the demand for the commodity and spot prices. The following proposition confirms that as $\tau_\theta \to \infty$ (i.e., the position of financial traders becomes publicly observable), the spot market equilibrium converges to the perfect-information benchmark.

PROPOSITION 2.B2: As $\tau_\theta \to \infty$, the spot price and aggregate demand converge to the perfect-information benchmark.

Proposition 2.B2 shows that by improving transparency of the futures market, one can achieve the perfect-information benchmark because by making the position of financial traders publicly observable, the θ noise no longer interferes with the information aggregation in the futures market. As a result, the futures price fully reveals the global fundamental, which allows goods producers to achieve the same efficiency allowed by the perfect-information benchmark. This nice convergence result relies on the assumption that the supply noise ξ does not affect the futures market trading at $t = 0$ and hits the spot market only at $t = 1$. Nevertheless, this result highlights the importance of improving market transparency.[60]

Imposing position limits on speculators in commodity futures markets has occupied much of the post-2008 policy debate, while improving market transparency has received much less attention. By highlighting the feedback effect originating from information frictions as a key channel for noise in futures market trading to affect commodity prices and demand, our model suggests that imposing position limits may not address the central information frictions that confront participants in commodity markets and thus may not be effective in reducing potential distortion caused by speculative trading. Instead, increasing the transparency of trading positions might be more effective.

[60] While our analysis focuses on the noise effect of their trading, financial traders can also contribute to information aggregation. As κ increases, the futures position of financial traders builds more on the global economic fundamental $\log A$, in which case the futures price $\log F$ becomes more informative of $\log A$. This is because one can prove based on Proposition B1 that $\tilde{h}_A/\tilde{h}_\theta$, the ratio of the loadings of $\log F$ on $\log A$ and θ, increases with κ.

2.B.4 Technical Proofs

Proof of Proposition 2.B1

We follow the same procedure as in the proof of Proposition 2 in the main paper to derive the futures market equilibrium at $t=0$. We first conjecture the log-linear forms for the futures price and each island producer's long position in (28) and (2.B.5) with the coefficients $\tilde{h}_0$, $\tilde{h}_A$, $\tilde{h}_\theta$, $\tilde{l}_0$, $\tilde{l}_s$, and $\tilde{l}_F$ to be determined by equilibrium conditions.

Let z be a sufficient statistic of the information contained in F :

$$z \equiv \frac{\log F - \tilde{h}_0 - \tilde{h}_\theta \bar{\theta}}{\tilde{h}_A} = \log A + \frac{\tilde{h}_\theta}{\tilde{h}_A}\left(\theta - \bar{\theta}\right).$$

Then, conditional on observing s_i and F, producer i's expectation of $\log A$ is

$$\begin{aligned} E\left[\log A \mid s_i, \log F\right] &= E\left[\log A \mid s_i, z\right] = \frac{1}{\tau_A + \tau_s + \frac{\tilde{h}_A^2}{\tilde{h}_\theta^2}\tau_\theta}\left(\tau_A \bar{a} + \tau_s s_i + \frac{\tilde{h}_A^2}{\tilde{h}_\theta^2}\tau_\theta z\right) \\ &= c_0 + c_s s_i + c_F\left(\log F - \tilde{h}_0 - \tilde{h}_\theta \bar{\theta}\right), \end{aligned} \tag{2.B.10}$$

where

$$\begin{aligned} c_0 &= \left(\tau_A + \tau_s + \frac{\tilde{h}_A^2}{\tilde{h}_\theta^2}\tau_\theta\right)^{-1}\left(\tau_A \bar{a} - \frac{\tilde{h}_A^2}{\tilde{h}_\theta^2}\tau_\theta \frac{\tilde{h}_0 + \tilde{h}_\theta \bar{\theta}}{\tilde{h}_A}\right), \\ c_s &= \left(\tau_A + \tau_s + \frac{\tilde{h}_A^2}{\tilde{h}_\theta^2}\tau_\theta\right)^{-1}\tau_s, \\ c_F &= \left(\tau_A + \tau_s + \frac{\tilde{h}_A^2}{\tilde{h}_\theta^2}\tau_\theta\right)^{-1}\frac{\tilde{h}_A}{\tilde{h}_\theta^2}\tau_\theta. \end{aligned}$$

Producer i's conditional variance of $\log A$ is

$$\tilde{\tau}_{A,i} = Var\left[\log A \mid s_i, \log F\right] = \left(\tau_A + \tau_s + \frac{\tilde{h}_A^2}{\tilde{h}_\theta^2}\tau_\theta\right)^{-1}. \tag{2.B.11}$$

By substituting equation (2.B.5) into producer i's optimal production decision in equation (2.B.2), we obtain

$$\begin{aligned} \log \tilde{X}_i &= \frac{1}{1-\phi(1-\eta)}\log\phi + \frac{\phi\eta}{1-\phi(1-\eta)}\tilde{l}_0 + \frac{1}{1-\phi(1-\eta)}\left(\phi\eta\tilde{l}_F - 1\right)\log F \\ &\quad + \left(\frac{1+\phi\eta\tilde{l}_s}{1-\phi(1-\eta)}\right)\left(c_0 + c_s s_i + c_F \frac{\log F}{\tilde{h}_A}\right) + \frac{\left(1+\phi\eta\tilde{l}_s\right)^2}{2\left(1-\phi(1-\eta)\right)}\tilde{\tau}_{A,i} + \frac{\phi^2\eta^2\tilde{l}_s^2}{2\left(1-\phi(1-\eta)\right)}\tau_s^{-1}. \end{aligned}$$

For the above equation to match the conjectured equilibrium position in equation (2.B.5), the constant term and the coefficients on s_i and $\log F$ have to be identical:

$$\begin{aligned}\tilde{l}_0 &= \frac{\phi\eta}{1-\phi(1-\eta)}\tilde{l}_0 + \left(\frac{1+\phi\eta\tilde{l}_s}{1-\phi(1-\eta)}\right)c_0 + \frac{\left(1+\phi\eta\tilde{l}_s\right)^2}{2(1-\phi(1-\eta))}\tilde{\tau}_{A,i} \\ &\quad + \frac{\phi^2\eta^2\tilde{l}_s^2}{2(1-\phi(1-\eta))}\tau_s^{-1} + \frac{1}{1-\phi(1-\eta)}\log\phi, \qquad (2.B.12)\\ \tilde{l}_s &= \left(\frac{1+\phi\eta\tilde{l}_s}{1-\phi(1-\eta)}\right)c_s, \qquad (2.B.13)\\ \tilde{l}_F &= \frac{\phi\eta}{1-\phi(1-\eta)}\tilde{l}_F - \frac{1}{1-\phi(1-\eta)} + \left(\frac{1+\phi\eta\tilde{l}_s}{1-\phi(1-\eta)}\right)c_F. \qquad (2.B.14)\end{aligned}$$

By substituting equation (2.B.13) into (2.B.14), we have

$$\tilde{l}_s = \frac{1+(1-\phi)\tilde{l}_F}{1-\phi(1-\eta)}\frac{\tilde{h}_\theta^2}{\tilde{h}_A}\tau_s\tau_\theta^{-1}. \qquad (2.B.15)$$

By manipulating equation (2.B.13), we also have that

$$\tilde{l}_s = \left(\tau_A + \frac{1-\phi}{1-\phi(1-\eta)}\tau_s + \frac{\tilde{h}_A^2}{\tilde{h}_\theta^2}\tau_\theta\right)^{-1}\frac{\tau_s}{1-\phi(1-\eta)}. \qquad (2.B.16)$$

We now use market clearing of the futures market to determine three other equations for the coefficients. Aggregating equation (2.B.5) gives the producers' aggregate position,

$$\int_{-\infty}^{\infty}\tilde{X}_i(s_i,F)\,d\Phi(\varepsilon_i) = \exp\left[\left(\tilde{l}_s + \tilde{l}_F\tilde{h}_A\right)\log A + \tilde{l}_F\tilde{h}_\theta\theta + \tilde{l}_0 + \tilde{l}_F\tilde{h}_0 + \frac{1}{2}\tilde{l}_s^2\tau_s^{-1}\right]. \qquad (2.B.17)$$

Equation (2.B.3) gives $\tilde{X}_S$. Define

$$z_\theta \equiv \frac{\log F - \tilde{h}_0 - \tilde{h}_A\bar{a}}{\tilde{h}_\theta} = \frac{\tilde{h}_A}{\tilde{h}_\theta}(\log A - \bar{a}) + \theta.$$

The suppliers' conditional expection of θ is then

$$E[\theta \mid \log F] = E[\theta \mid z_\theta] = \left(\tau_\theta + \frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\right)^{-1}\left[\tau_\theta\bar{\theta} + \frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\left(\frac{\log F - \tilde{h}_0}{\tilde{h}_\theta} - \frac{\tilde{h}_A}{\tilde{h}_\theta}\bar{a}\right)\right],$$

with conditional variance $Var[\theta \mid \log F] = \left(\tau_\theta + \frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\right)^{-1}$. Their conditional expectation of $\log A$ is

$$E[\log A \mid \log F] = E[\log A \mid z_\theta] = \left(\tau_A + \frac{\tilde{h}_A^2}{\tilde{h}_\theta^2}\tau_\theta\right)^{-1}\left[\tau_A\bar{a} + \frac{\tilde{h}_A^2}{\tilde{h}_\theta^2}\tau_\theta\left(\frac{\log F - \tilde{h}_0 - \tilde{h}_\theta\bar{\theta}}{\tilde{h}_A}\right)\right],$$

with conditional variance $Var\left[\log A \mid \log F\right] = \left(\tau_A + \frac{\tilde{h}_A^2}{\tilde{h}_\theta^2}\tau_\theta\right)^{-1}$. Thus, we obtain an expression for $\log \tilde{X}_S$ that is linear in $\log A$ and θ.

Next, the market-clearing condition $\log\left[e^{\kappa \log A + \theta}\int_{-\infty}^{\infty}\tilde{X}_i\left(s_i, F\right)d\Phi\left(\varepsilon_i\right)\right] = \log \tilde{X}_S$ requires that the coefficients on $\log A$ and θ and the constant term be identical on both sides:

$$\begin{aligned}
\kappa + \tilde{l}_s + \tilde{l}_F\tilde{h}_A &= k\tilde{h}_A + \left(\tau_\theta + \frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\right)^{-1}\left(\frac{\tilde{h}_\theta}{\tilde{h}_A}\tau_A + \kappa\tau_\theta\right), && \text{(2.B.18)}\\
1 + \tilde{l}_F\tilde{h}_\theta &= k\tilde{h}_\theta + \left(\tau_\theta + \frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\right)^{-1}\frac{\tilde{h}_\theta}{\tilde{h}_A}\left(\frac{\tilde{h}_\theta}{\tilde{h}_A}\tau_A + \kappa\tau_\theta\right), && \text{(2.B.19)}\\
\tilde{l}_0 + \tilde{l}_F\tilde{h}_0 + \frac{1}{2}\tilde{l}_s^2\tau_s^{-1} &= k\tilde{h}_0 + \left(\tau_\theta + \frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\right)^{-1}\left(1 + \kappa\frac{\tilde{h}_\theta}{\tilde{h}_A}\right)\tau_\theta\bar{\theta} && \text{(2.B.20)}\\
&\quad - \left(\tau_\theta + \frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\right)^{-1}\frac{\tilde{h}_\theta}{\tilde{h}_A}\left(1 + \kappa\frac{\tilde{h}_\theta}{\tilde{h}_A}\right)\tau_A\bar{a} + \bar{\xi} - \sigma_\xi^2/2k\\
&\quad - \frac{\kappa^2}{2k}\left(1+2k\right)\left(\tau_A + \frac{\tilde{h}_A^2}{\tilde{h}_\theta^2}\tau_\theta\right)^{-1} - \frac{1}{2k}\left(1+2k\right)\left(\tau_\theta + \frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\right)^{-1}.
\end{aligned}$$

Equation (2.B.19) directly implies that

$$\tilde{l}_F = k + \left(\tau_\theta + \frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\right)^{-1}\left(\kappa\frac{\tilde{h}_\theta}{\tilde{h}_A} - 1\right)\tau_\theta\tilde{h}_\theta^{-1}. \tag{2.B.21}$$

Equations (2.B.18) and (2.B.19) together imply that

$$\tilde{l}_s = \tilde{h}_\theta^{-1}\tilde{h}_A - \kappa.$$

By combining this equation with (2.B.16), we arrive at

$$\tilde{l}_s^3 + 2\kappa\tilde{l}_s^2 + \left(\tau_\theta^{-1}\tau_A + \frac{1-\phi}{1-\phi\left(1-\eta\right)}\tau_\theta^{-1}\tau_s + \kappa^2\right)\tilde{l}_s - \frac{\tau_\theta^{-1}\tau_s}{1-\phi\left(1-\eta\right)} = 0. \tag{2.B.22}$$

By further making the convenient substitution $L_s = \tilde{l}_s + \frac{2}{3}\kappa$, called the Tschirnhaus transformation, we obtain the depressed cubic polynomial

$$L_s^3 + pL_s + q = 0,$$

where

$$\begin{aligned}
p &= \tau_\theta^{-1}\tau_A + \frac{1-\phi}{1-\phi\left(1-\eta\right)}\tau_\theta^{-1}\tau_s - \frac{1}{3}\kappa^2,\\
q &= -\frac{2}{3}\kappa\tau_\theta^{-1}\tau_A - \frac{2}{3}\kappa\frac{1-\phi}{1-\phi\left(1-\eta\right)}\tau_\theta^{-1}\tau_s - \frac{2}{27}\kappa^3 - \frac{\tau_\theta^{-1}\tau_s}{1-\phi\left(1-\eta\right)}.
\end{aligned}$$

It is easy to verify that $\frac{q^2}{4}+\frac{p^3}{27}>0$ and therefore L_s is a real root of this depressed cubic polynomial, which has one real and two complex roots. Following Cardano's method, the one real root of equation (2.B.22) is given by

$$\tilde{l}_s=\sqrt[3]{-\frac{q}{2}+\sqrt{\frac{q^2}{4}+\frac{p^3}{27}}}+\sqrt[3]{-\frac{q}{2}-\sqrt{\frac{q^2}{4}+\frac{p^3}{27}}}-\frac{2}{3}\kappa.$$

Since the coefficients of equation (2.B.22) change sign only once, by Descartes' Rule of Signs the real root must be positive.

Since $\tilde{l}_s=\tilde{h}_\theta^{-1}\tilde{h}_A-\kappa$, we have

$$\tilde{h}_\theta=\left(\tilde{l}_s+\kappa\right)^{-1}\tilde{h}_A,$$

which, together with our expression for $\tilde{l}_s$ and equations (2.B.15) and (2.B.21), implies that

$$\tilde{h}_\theta=\left(\left(1-\phi\left(1-\eta\right)\right)\tau_s^{-1}+\frac{1-\phi}{\tau_\theta\left(\tilde{l}_s+\kappa\right)^2+\tau_A}\right)\frac{\tau_\theta}{1+k\left(1-\phi\right)}\tilde{l}_s\left(\tilde{l}_s+\kappa\right) \tag{2.B.23}$$

and therefore

$$\tilde{h}_A=\left(\left(1-\phi\left(1-\eta\right)\right)\tau_s^{-1}+\frac{1-\phi}{\tau_\theta\left(\tilde{l}_s+\kappa\right)^2+\tau_A}\right)\frac{\tau_\theta}{1+k\left(1-\phi\right)}\tilde{l}_s\left(\tilde{l}_s+\kappa\right)^2. \tag{2.B.24}$$

Since by equation (2.B.22), $\tilde{l}_s>0$, $\tilde{h}_A$ and $\tilde{h}_\theta$ must have the same sign. With $\tilde{h}_A$ and $\tilde{h}_\theta$ determined, $\tilde{l}_F$ is then given by equation (2.B.21),

$$\tilde{l}_F=k+\left(\tau_\theta+\frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\right)^{-1}\left(\kappa\frac{\tilde{h}_\theta}{\tilde{h}_A}-1\right)\tau_\theta\tilde{h}_\theta^{-1},$$

$\tilde{h}_0$ by equation (2.B.12),

$$\begin{aligned}\tilde{h}_0 \quad = \quad & \left(k-\tilde{l}_F+\frac{1-\phi\left(1-\eta\right)}{1-\phi}\tilde{l}_s\tau_s^{-1}\frac{\tilde{h}_A}{\tilde{h}_\theta^2}\tau_\theta\right)^{-1}\cdot\\ & \left(\frac{1}{1-\phi}\log\phi-\bar{\xi}+\sigma_\xi^2/2k+\frac{1}{2k}\left(1+2k\right)\left(1+\kappa^2\frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\right)\left(\tau_\theta+\frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\right)^{-1}\right.\\ & +\frac{1}{2}\left(\tilde{l}_s+\frac{1-\phi\left(1-\eta\right)}{1-\phi}\left(1+\phi\eta\tilde{l}_s+\frac{\phi^2\eta^2\tilde{l}_s}{1-\phi\left(1-\eta\right)}\right)\right)\tilde{l}_s\tau_s^{-1}\\ & \left.+\left(\frac{1-\phi\left(1-\eta\right)}{1-\phi}\tilde{l}_s\tau_s^{-1}+\left(\tau_\theta\frac{\tilde{h}_A^2}{\tilde{h}_\theta^2}+\tau_A\right)^{-1}\left(\frac{\tilde{h}_A}{\tilde{h}_\theta}+\kappa\right)\right)\left(\tau_A\bar{a}-\left(\tilde{l}_s+\kappa\right)\tau_\theta\bar{\theta}\right)\right),\end{aligned}$$

and $\tilde{l}_0$ by equation (2.B.20),

$$\begin{aligned}\tilde{l}_0 &= \left(k-\tilde{l}_F\right)\tilde{h}_0+\bar{\xi}-\sigma_\xi^2/2k+\left(\tau_\theta+\frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\right)^{-1}\left(1+\kappa\frac{\tilde{h}_\theta}{\tilde{h}_A}\right)\tau_\theta\bar{\theta} \\ &\quad -\left(\tau_\theta+\frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\right)^{-1}\frac{\tilde{h}_\theta}{\tilde{h}_A}\left(1+\kappa\frac{\tilde{h}_\theta}{\tilde{h}_A}\right)\tau_A\bar{a}-\frac{1}{2}\tilde{l}_s^2\tau_s^{-1} \\ &\quad -\frac{1}{2k}(1+2k)\left(1+\kappa^2\frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\right)\left(\tau_\theta+\frac{\tilde{h}_\theta^2}{\tilde{h}_A^2}\tau_A\right)^{-1}.\end{aligned}$$

We now derive the spot market equilibrium at $t=1$. We again first conjecture that the spot price P_X and a goods producer's updated commodity demand take the log-linear forms given in equations (29) and (30) with the coefficients h_0, h_A, h_F, h_ξ, l_0, l_s, l_F, and l_P to be determined by equilibrium conditions.

The mean and variance of producer i's prior belief over $\log A$ carried from $t=0$ is derived in (2.B.10) and (2.B.11). Define

$$z_p=\frac{\log P_X-h_0-h_F\log F-h_\xi\bar{\xi}}{h_A}=\log A+\frac{h_\xi}{h_A}\left(\xi-\bar{\xi}\right).$$

Then, after observing the spot price P_X at $t=1$, the producer's expectation of $\log A$ is

$$E\left[\log A\mid s_i,\log F,\log P_X\right]=E\left[\log A\mid s_i,\log F,z_p\right]=\frac{\tilde{\tau}_{A,i}c_s s_i+\frac{h_A^2}{h_\xi^2}\left(\frac{\log P_X-h_0-h_F\log F-h_\xi\bar{\xi}}{h_A}\right)}{\tilde{\tau}_{A,i}+\frac{h_A^2}{h_\xi^2}\tau_\xi},$$

with conditional variance

$$\tau_{A,i}=Var\left[\log A\mid s_i,\log F,\log P_X\right]=\left(\tilde{\tau}_{A,i}+\frac{h_A^2}{h_\xi^2}\tau_\xi\right)^{-1}.$$

We use (2.B.1) to compute $\log X_i$, and obtain a linear expression of s_i, $\log F$, and P_X. By matching the coefficients of this expression with the conjectured form in (30), we obtain

$$\begin{aligned}l_0 &= \frac{1}{1-\phi}\log\phi+\frac{(1+\phi\eta l_s)^2}{2(1-\phi)}\tau_{A,i}+\frac{1}{2(1-\phi)}\phi^2\eta^2 l_s^2\tau_s^{-1}-\frac{1+\phi\eta l_s}{1-\phi}\tau_{A,i}\frac{h_A}{h_\xi^2}\left(h_0+h_\xi\bar{\xi}\right) \\ &\quad +\frac{1}{1-\phi}(1+\phi\eta l_s)\tau_{A,i}\tilde{\tau}_{A,i}\left(c_0-c_F\left(\tilde{h}_0+\tilde{h}_\theta\bar{\theta}\right)\right), \\ l_s &= \frac{\tilde{\tau}_{A,i}c_s}{(1-\phi(1-\eta))\tau_{A,i}^{-1}-\phi\eta\tilde{\tau}_{A,i}c_s}, \\ l_F &= \frac{1}{1-\phi}(1+\phi\eta l_s)\tau_{A,i}\left(\tilde{\tau}_{A,i}c_F-\frac{h_A}{h_\xi^2}h_F\right), &(2.B.25) \\ l_P &= \frac{1}{1-\phi}(1+\phi\eta l_s)\tau_{A,i}\frac{h_A}{h_\xi^2}-\frac{1}{1-\phi}. &(2.B.26)\end{aligned}$$

Market clearing of the spot market requires $\int_{-\infty}^{\infty} X_i d\Phi(\varepsilon_i) = X_S$, which implies

$$(k - l_P)\log P_X = l_0 + \frac{1}{2} l_s^2 \tau_s^{-1} + l_s \log A + l_F \log F - \xi.$$

By matching coefficients on both sides, we have

$$\begin{aligned} (k-l_P)\, h_0 &= l_0 + \frac{1}{2} l_s^2 \tau_s^{-1}, \\ (k-l_P)\, h_A &= l_s, \qquad (2.B.27) \\ (k-l_P)\, h_F &= l_F, \qquad (2.B.28) \\ (k-l_P)\, h_\xi &= -1. \qquad (2.B.29) \end{aligned}$$

From equations (2.B.27) and (2.B.29), we have that $l_s = -\frac{h_A}{h_\xi}$, and given our expression for l_0 and l_F above, we also see that

$$\begin{aligned} h_0 &= \left(k - l_P + \frac{1+\phi\eta l_s}{1-\phi}\tau_{A,i}\frac{h_A}{h_\xi^2}\right)^{-1} \cdot \left(\frac{1}{1-\phi}\log\phi + \frac{(1+\phi\eta l_s)^2}{2(1-\phi)}\tau_{A,i} + \frac{1}{2(1-\phi)}\phi^2\eta^2 l_s^2 \tau_s^{-1}\right. \\ &\quad \left. + \frac{1}{1-\phi}(1+\phi\eta l_s)\tau_{A,i}\tilde{\tau}_{A,i}\left(c_0 - c_F\left(\tilde{h}_0 + \tilde{h}_\theta\bar{\theta}\right)\right) - \frac{1+\phi\eta l_s}{1-\phi}\tau_{A,i}\frac{h_A}{h_\xi^2}h_\xi\bar{\xi} + \frac{1}{2}l_s^2\tau_s^{-1}\right), \\ h_F &= \left(\frac{1-\phi}{1+\phi\eta l_s}\tau_{A,i}^{-1}(k-l_P) + \frac{h_A}{h_\xi^2}\right)^{-1}\tilde{\tau}_{A,i}c_F. \qquad (2.B.30) \end{aligned}$$

From our expression for l_s above and $l_s = -h_A/h_\xi$, we have

$$l_s^3 + \tau_\xi^{-1}\left(\tilde{\tau}_{A,i} - \frac{\phi\eta\tilde{\tau}_{A,i}c_s}{1-\phi(1-\eta)}\right)l_s - \frac{\tau_\xi^{-1}\tilde{\tau}_{A,i}c_s}{1-\phi(1-\eta)} = 0. \qquad (2.B.31)$$

This is a depressed cubic polynomial whose unique real and positive root is given by

$$\begin{aligned} l_s &= \sqrt[3]{-\frac{1}{2}\frac{\tau_\xi^{-1}\tilde{\tau}_{A,i}a_s}{1-\phi(1-\eta)} + \sqrt{\frac{1}{4}\left(\frac{\tau_\xi^{-1}\tilde{\tau}_{A,i}a_s}{1-\phi(1-\eta)}\right)^2 + \frac{1}{27}\tau_\xi^{-3}\left(\left(\tilde{\tau}_{A,i} - \frac{\phi\eta\tilde{\tau}_{A,i}a_s}{1-\phi(1-\eta)}\right)\right)^3}} \\ &\quad + \sqrt[3]{-\frac{1}{2}\left(\frac{\tau_\xi^{-1}\tilde{\tau}_{A,i}a_s}{1-\phi(1-\eta)}\right) - \sqrt{\frac{1}{4}\left(\frac{\tau_\xi^{-1}\tilde{\tau}_{A,i}a_s}{1-\phi(1-\eta)}\right)^2 + \frac{1}{27}\tau_\xi^{-3}\left(\left(\tilde{\tau}_{A,i} - \frac{\phi\eta\tilde{\tau}_{A,i}a_s}{1-\phi(1-\eta)}\right)\right)^3}}. \end{aligned}$$

It follows that $l_s > 0$ and from equation (2.B.29) that

$$h_A = \frac{(1-\phi)\, l_s + (1+\phi\eta l_s)\left(\tilde{\tau}_{A,i} + l_s^2\tau_\xi\right)^{-1} l_s^2}{1+(1-\phi)\, k} > 0,$$

and, since $l_s = -h_A/h_\xi > 0$,

$$h_\xi = -\frac{1-\phi+(1+\phi\eta l_s)\left(\tilde{\tau}_{A,i}+l_s^2\tau_\xi\right)^{-1} l_s}{1+(1-\phi)\,k} < 0.$$

We now prove that $l_F > 0$. Given the expression for l_F in (2.B.25) and given $l_s > 0$, it is sufficient for $l_F > 0$ if

$$\tilde{\tau}_{A,i} c_F > \frac{h_A}{h_\xi^2} h_F.$$

Given the expression for h_F in (2.B.30), and recognizing that $\tilde{\tau}_{A,i} > 0$ and $c_F > 0$, the above condition can be rewritten as

$$1 > \frac{h_A}{h_\xi^2}\left(\frac{1-\phi}{1+\phi\eta l_s}\tau_{A,i}^{-1}\,(k-l_P)+\frac{h_A}{h_\xi^2}\right)^{-1}.$$

Furthermore, from the expressions for $\tau_{A,i}$ and l_P, this condition can be further expressed as

$$\frac{1}{1+\phi\eta l_s}\left(1+k\,(1-\phi)\right)\left(\tilde{\tau}_{A,i}+\frac{h_A^2}{h_\xi^2}\tau_\xi\right) > \frac{h_A}{h_\xi^2}.$$

Since $l_s = -\frac{h_A}{h_\xi}$, given our expression for $h_\xi < 0$, the condition reduces to

$$\frac{1-\phi}{1+\phi\eta l_s}\left(\tilde{\tau}_{A,i}+l_s^2\tau_\xi\right) > 0,$$

which is always satisfied. Therefore, $l_F > 0$. In addition, since $(k-l_P)\,h_A = l_s$ implies that $k > l_P$, we see from $(k-l_P)\,h_F = l_F$ that $h_F > 0$.

We now examine the sign of l_P. By substituting $l_s = -\frac{h_A}{h_\xi}$ and the expressions of $\tau_{A,i}$ and h_ξ into (2.B.26), we have

$$l_P = -\frac{1}{h_\xi\,(1+(1-\phi)\,k)}\left(\tilde{\tau}_{A,i}+l_s^2\tau_\xi\right)^{-1}\left(kl_s-(\tau_\xi-k\phi\eta)\,l_s^2-\tilde{\tau}_{A,i}\right).$$

Consequently, l_P can be positive or negative depending on the sign of $kl_s-(\tau_\xi-k\phi\eta)\,l_s^2-\tilde{\tau}_{A,i}$.

Proof of Proposition 2.B2

In (31), $\log P_X$ is a linear expression of $\log A$, θ, and ξ. We need to show that as $\tau_\theta \to \infty$, the coefficients on $\log A$ and ξ converge to their corresponding values in the perfect-information benchmark (Proposition 3 of the main paper), and the variance of θ

$$V_\theta = h_F^2\tilde{h}_\theta^2\tau_\theta^{-1} \to 0.$$

We rewrite equation (2.B.22) as

$$\left(\tilde{l}_s+\kappa\right)^2\tilde{l}_s+\tau_\theta^{-1}\left(\tau_A+\frac{1-\phi}{1-\phi\left(1-\eta\right)}\tau_s\right)\tilde{l}_s=\frac{\tau_\theta^{-1}\tau_s}{1-\phi\left(1-\eta\right)}.$$

As τ_θ becomes sufficiently large, the right-hand side converges to zero and therefore, since the cubic polynomial has a unique real solution, $\tilde{l}_s \to 0$. By substituting equation (2.B.22) into our expression for $\tilde{h}_A$, one can express $\tilde{h}_A$ as

$$\tilde{h}_A=\frac{1-\phi\left(1-\eta\right)}{1+k\left(1-\phi\right)}\tau_s^{-1}\left(1+\frac{\left(1-\phi\right)\tilde{l}_s}{1-\left(1-\phi\right)\tilde{l}_s}\right)\left(\frac{\tau_s}{1-\phi\left(1-\eta\right)}-\left(\tau_A+\frac{1-\phi}{1-\phi\left(1-\eta\right)}\tau_s\right)\tilde{l}_s\right).$$

As $\tau_\theta \to \infty$, $\tilde{l}_s \to 0$, and thus $\tilde{h}_A \to \frac{1}{1+k(1-\phi)}$. In addition, by substituting for c_s, we can rewrite (2.B.31) as

$$\tau_\xi l_s^3+\tilde{\tau}_{A,i}l_s=\left(1+\phi\eta l_s\right)\frac{\tilde{\tau}_{A,i}^2\tau_s}{1-\phi\left(1-\eta\right)}.$$

Since $\tau_\theta\left(l_s+\kappa\right)^2$ grows as τ_θ increases, $\tilde{\tau}_{A,i}=\left(\tau_s+\tau_A+\tau_\theta\left(l_s+\kappa\right)^2\right)^{-1}\to 0$ as $\tau_\theta\to\infty$. It then follows that $l_s \to 0$.

By substituting (2.B.31) and our expression for c_s into our expression for h_ξ, we have

$$h_\xi=-\frac{1-\phi}{1+\left(1-\phi\right)k}-\frac{1-\phi\left(1-\eta\right)}{1+\left(1-\phi\right)k}\tau_s^{-1}\left(\tilde{\tau}_{A,i}l_s\right)^2.$$

As $\tau_\theta\to\infty$, $\tilde{\tau}_{A,i}l_s=\left(1+\phi\eta l_s\right)\frac{\tilde{\tau}_{A,i}^2\tau_s}{1-\phi(1-\eta)}-\tau_\xi l_s^3\to 0$, and therefore $h_\xi\to-\frac{1-\phi}{1+k(1-\phi)}$. Given that $l_s=-\frac{h_A}{h_\xi}$ and given our expression for l_P, we have that as $\tau_\theta\to\infty$, the coefficient of ξ in (2.B.8) equals

$$l_P h_\xi=-\frac{1}{1-\phi}\left(1+\phi\eta l_s\right)\tau_{A,i}l_s-\frac{1}{1-\phi}h_\xi\to\frac{1}{1+k\left(1-\phi\right)},$$

which is its value in the perfect-information benchmark.

Since $l_s=-\frac{h_A}{h_\xi}$, and given that as $\tau_\theta\to\infty$, $l_s\to 0$ and $h_\xi\to-\frac{1-\phi}{1+k(1-\phi)}$, we have $h_A\to 0$. By substituting for $\tau_{A,i}$, c_F, l_P, and $\tilde{h}_A/\tilde{h}_\theta$, we can rewrite $h_F\tilde{h}_A$ as

$$\begin{aligned}h_F\tilde{h}_A &= \frac{1-\phi\left(1-\eta\right)}{1+k\left(1-\phi\right)}\tau_s^{-1}\tau_\theta\left(l_s+\kappa\right)^2 l_s\\ &= \frac{1-\phi\left(1-\eta\right)}{1+k\left(1-\phi\right)}\tau_s^{-1}\left(\left(1+\phi\eta l_s\right)\frac{\tau_s}{1-\phi\left(1-\eta\right)}-\tau_\xi\left(\tilde{\tau}_{A,i}^{-1}l_s^{3/2}\right)^2-\left(\tau_s+\tau_A\right)l_s\right),\end{aligned}$$

where we use substitution with equation (2.B.31). As $\tau_\theta\to\infty$, $l_s\to 0$, and $\left(\tilde{\tau}_{A,i}^{-1}l_s^{3/2}\right)^2\to 0$, the coefficient on $\log A$ in (2.B.8) $h_A+h_F\tilde{h}_A\to\frac{1}{1+k(1-\phi)}$, which is its value in the perfect-information benchmark.

By using the expressions of h_F, l_P, l_s, c_F, $\tau_{A,i}$, $\tilde{l}_s$, and $\tilde{h}_\theta$ in Proposition B1 and by manipulating terms, we have

$$h_F\tilde{h}_\theta = \frac{1-\phi(1-\eta)}{1+k(1-\phi)}\tau_s^{-1} l_s \left(\tilde{l}_s+\kappa\right)\tau_\theta.$$

Consequently, we can write V_θ as

$$V_\theta = \left(\frac{1-\phi(1-\eta)}{1+k(1-\phi)}\tau_s^{-1}\right)^2 l_s^2 \left(\tilde{l}_s+\kappa\right)^2 \tau_\theta.$$

We can rewrite equation (2.B.22) as

$$\tau_\theta\left(\tilde{l}_s+\kappa\right)^2 = \frac{\tau_s}{1-\phi(1-\eta)}\tilde{l}_s^{-1} - \left(\tau_A + \frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right).$$

By applying the Implicit Function Theorem to equation (2.B.22),

$$\frac{\partial \tilde{l}_s}{\partial \tau_\theta} = -\frac{\left(\tilde{l}_s+\kappa\right)^2 \tilde{l}_s^2}{2\tau_\theta\left(\tilde{l}_s+\kappa\right)\tilde{l}_s^2 + \frac{\tau_s}{1-\phi(1-\eta)}} < 0.$$

Consequently, $\tau_\theta\left(\tilde{l}_s+\kappa\right)^2$ is growing in τ_θ. Now we can rewrite equation (2.B.31) by substituting for $\tilde{\tau}_{A,i}$ and c_s as

$$\left(\tau_A+\tau_s+\left(\tilde{l}_s+\kappa\right)^2\tau_\theta\right)\sqrt{\frac{\tau_\xi l_s^3 + \left(\tau_A+\tau_s+\left(\tilde{l}_s+\kappa\right)^2\tau_\theta\right)^{-1} l_s}{1+\phi\eta l_s}} = \sqrt{\frac{\tau_s}{1-\phi(1-\eta)}}.$$

As $\tau_\theta \to \infty$, $l_s \to 0$. Thus, for this equation to hold, we must have $\tau_\theta\left(l_s+\kappa\right)^2 \to \infty$. The left-hand side (LHS) of the above equation can then be expressed as

$$\begin{aligned} &\left(\tau_A+\tau_s+\left(\tilde{l}_s+\kappa\right)^2\tau_\theta\right)\sqrt{\frac{\tau_\xi l_s^3 + \left(\tau_A+\tau_s+\left(\tilde{l}_s+\kappa\right)^2\tau_\theta\right)^{-1} l_s}{1+\phi\eta l_s}} \\ \approx\; & \tau_\theta\left(\tilde{l}_s+\kappa\right)^2 l_s^{3/2}\sqrt{\frac{\tau_\xi}{1+\phi\eta l_s}} + o\left(\tau_\theta^{-1}\left(\tilde{l}_s+\kappa\right)^{-2}\right). \end{aligned}$$

This suggests that $l_s^{3/2}$ must be shrinking at approximately the same rate as $\tau_\theta\left(\tilde{l}_s+\kappa\right)^2$ is growing for the LHS to remain finite. Therefore, l_s^2 must be shrinking at a faster rate and $V_\theta \to 0$ as $\tau_\theta \to \infty$.

In summary, we have shown that as $\tau_\theta \to \infty$, $\log P_X$ converges with its counterpart in the perfect-information benchmark. We can similarly prove that the producers' aggregate demand also converges.

Chapter 3

Supply Elasticity and Housing Cycles[61]

3.1. Introduction

Conventional wisdom posits that supply elasticity attenuates housing cycles. As a result, one expects housing prices to be more volatile in areas with more inelastic housing supplies. However, as noted by Glaeser (2013) and other commentators, during the recent U.S. housing cycle in the 2000s, some areas such as Las Vegas and Phoenix experienced more dramatic housing price booms and busts, despite their relatively elastic housing supply, compared to areas with more inelastic supply, such as New York and Los Angeles. Interestingly, by systematically examining the cross-section of the booms and busts experienced by different counties during this housing cycle, we find that the monotonically decreasing relationship between the magnitude of housing cycles and supply elasticity is more fragile than commonly perceived. If one simply sorts counties into three groups based on Saiz's (2010) widely used measure of supply elasticity, each with an equal number of counties, the average housing price increase in the boom period of 2004-2006 and drop in the bust period of 2007-2009 monotonically decreases across the inelastic, middle, and elastic groups. However, as the inelastic group holds more than half of the population, this coarse grouping may disguise non-monotonicity under present finer parsings. Indeed, when we sort the counties into ten elasticity groups, each with an equal number of counties, or into either three or ten elasticity groups each with an equal population, we uncover a non-monotonic relationship between the magnitudes of the housing price booms and busts experienced by different counties and their supply elasticity. The most dramatic boom and bust cycle occurs in an intermediate range of supply elasticity.

This hump-shaped relationship between housing cycle and supply elasticity, which we summarize in Section 2, is intriguing and cannot be explained by the usual supply-side

[61] This chapter is based on joint work with Zhenyu Gao at CUHK and Wei Xiong at Princeton University.

mechanisms. In this paper, we develop a theoretical model to highlight a novel mechanism for supply elasticity to affect housing demand through a learning channel. We emphasize that home buyers observe neither the economic strength of a neighborhood, which ultimately determines the demand for housing in the neighborhood, nor the supply of housing. In the presence of these pervasive informational frictions, local housing markets provide a useful platform for aggregating information. This fundamental aspect of housing markets, however, has received little attention in the academic literature. It is intuitive that traded housing prices reflect the net effect of demand and supply factors. Supply elasticity determines the weight of supply-side factors in determining housing prices and therefore by extension determines the informational noise faced by home buyers in using housing prices as signals for the strength of demand.

Our model integrates the standard framework of Grossman and Stiglitz (1980) and Hellwig (1980) for information aggregation in asset markets with a housing market in a local neighborhood. This setting allows us to extend the insights of market microstructure analysis to explore the real consequences of informational frictions in housing markets. In particular, our model allows us to analyze how agents form expectations in housing markets, how these expectations interact with characteristics endemic to a neighborhood, and how these expectations feed into housing prices.

We first present a baseline setting in Section 3 to highlight the basic information aggregation mechanism with each household purchasing homes for their own consumption and then extend the model in Section 4 to further incorporate purchases of investment homes. The baseline model features a continuum of households in a closed neighborhood, which can be viewed as a county. Each household in the neighborhood has a Cobb-Douglas utility function over its goods consumption and housing consumption, as well as housing consumption of its neighbors. This complementarity in the households' housing consumption motivates each household to learn about the unobservable economic strength of the neighborhood, which determines the common productivity of all households and thus their housing demand.

Despite each household's housing demand being non-linear, the Law of Large Numbers allows us to aggregate their housing demand in closed form and to derive a unique log-linear equilibrium for the housing market. Each household possesses a private signal regarding the neighborhood common productivity. By aggregating the households' housing demand, the housing price aggregates their private signals. However, the presence of unobservable supply shocks prevents the housing price from perfectly revealing the neighborhood strength and acts as a source of informational noise in the housing price.

Our model also builds in another important feature that households underestimate supply elasticity. By examining a series of historical episodes of real estate speculation in the U.S., Glaeser (2013) summarizes the tendency of speculators to underestimate the response of housing supply to rising prices as a key for understanding these historical experiences. In our model, underestimation of supply elasticity implies that households underestimate the amount of informational noise in the observed price signal, which in turn causes the households' expectations of the neighborhood strength and housing demand to overreact to the housing price. The amplification of housing price volatility induced by such overreaction depends on the uncertainty faced by households and the informational content of the price, both of which are endogenously linked to the neighborhood's supply elasticity.

It is useful to consider two polar cases. At one end with the supply being infinitely inelastic, housing prices are fully determined by the strength of the neighborhood and thus perfectly reveals it. At the other end, with housing supply being infinitely elastic, housing prices are fully determined by the supply shock, and households' uncertainty about the strength of the neighborhood does not interact with the housing price. In between these two polar cases, the households face uncertainty regarding neighborhood strength and the uncertainty matters to the housing price. Consequently, households' overreaction to the price signal has the most pronounced effect on their housing demand and housing price in an intermediate range of supply elasticity, causing the price volatility to have a hump-shaped relationship with supply elasticity. That is, housing price volatility is largest at an intermediate supply elasticity rather than when supply is infinitely inelastic. This insight helps explain the aforementioned empirical observation that during the recent U.S. housing cycle, counties with supply elasticities in an intermediate range experienced the most dramatic price booms and busts.

We further extend the baseline model in Section 4 to incorporate immigrants who are attracted to the neighborhood by its economic strength in a later period and the speculation of the current households in acquiring secondary homes in anticipation of selling them to immigrants. This model extension generates two additional predictions. First, the households' learning effects can induce another non-monotonic relationship between the variability of the share of secondary home purchases among total home purchases and supply elasticity.

The intuition is similar to before. As secondary home purchases are more sensitive than primary home purchases to the households' expectations of the neighborhood strength, informational frictions and the households' overreaction to the price signal make households' secondary home purchases most variable at an intermediate range of supply elasticity. This mechanism also leads to a second prediction regarding a positive relationship between the variability of the share of secondary home purchases and the volatility of housing prices across neighborhoods with different elasticities.

Interestingly, we are able to confirm these new model predictions in the data. First, we find that counties in an intermediate range of supply elasticity indeed had the largest change in the share of non-owner-occupied home (secondary home) purchases from the pre-boom period of 2001-2003 to the boom period of 2004-2006, as opposed to counties with either the most elastic or inelastic supplies. Second, counties with greater increases in the share of non-owner-occupied home purchases from 2001-2003 to 2004-2006 also experienced larger price increases in 2004-2006 and larger price decreases in the bust period of 2007-2009. These empirical findings provide evidence from a new dimension to support the important roles played by informational frictions and household learning in driving housing cycles.

3.2. Related Literature

The existing literature has emphasized the importance of accounting for home buyers' expectations (and in particular extrapolative expectations) in understanding dramatic housing boom and bust cycles, e.g., Case and Shiller (2003); Glaeser, Gyourko, and Saiz (2008); and Piazzesi and Schneider (2009). Much of the analyses and discussions, however, are made in the absence of a systematic framework that anchors home buyers' expectations to their information aggregation and learning process. In this paper, we fill this gap by developing a model for analyzing information aggregation and learning in housing markets. By doing so, we are able to uncover a novel effect of supply elasticity, beyond its role in driving housing supply, in determining the informational content of the housing price and households' learning from the price signal. This learning effect implies non-monotonic patterns in housing price volatility and the variability of share of secondary home purchases across neighborhoods with different supply elasticities. This learning mechanism also differentiates our model from Gao (2013), which shares a similar motivation as ours to explain the dramatic housing cycles in the 2000s experienced by areas with intermediate supply elasticities, and which emphasizes a joint effect of time-to-build and housing speculators' extrapolative expectations as an explanation.

In our model, households overreact to the housing price signal. Such overreaction is driven by their underestimation of supply elasticity. This overreaction mechanism, which depends on the informational frictions faced by households and the endogenous informational content of the housing price, is different from the commonly discussed mechanisms in the behavioral finance literature, such as overconfidence highlighted by Daniel, Hirshleifer, and Subrahmanyam (1998), slow information diffusion by Hong and Stein (1999), and extrapolation by Barberis, Shleifer, and Vishny (1998) and Barberis et al. (2014).

Our model also differs from Burnside, Eichenbaum, and Rebelo (2013), which offers a model of housing market booms and busts based on the epidemic spreading of optimistic or pessimistic beliefs among home buyers through their social interactions. Our learning-based mechanism is also different from Nathanson and Zwick (2014), which studies the hoarding of land by home builders in certain elastic areas as a mechanism to amplify price volatility in the recent U.S. housing cycle. Glaeser and Nathanson (2015) presents a model of biased learning in housing markets, building on current buyers not adjusting for the expectations of past buyers and instead assuming that past prices reflect only contemporaneous demand. This incorrect inference gives rise to correlated errors in housing demand forecasts over time, which in turn generate excess volatility, momentum, and mean-reversion in housing prices. In contrast to this model, informational frictions in our model anchor on the interaction between the demand and supply sides, and, in particular, on the elasticity of housing supply. This key feature is also different from the amplification to price volatility induced by dispersed information and short-sale constraints featured in Favara and Song (2014).

By focusing on information aggregation and learning of symmetrically informed households with dispersed private information, our study differs in emphasis from those that analyze the presence of information asymmetry between buyers and sellers of homes, such as Garmaise and Moskowitz (2004) and Kurlat and Stroebel (2014). Neither does our model emphasize the potential asymmetry between in-town and out-of-town home buyers, which is shown to be important by a recent study of Chinco and Mayer (2013).

There are extensive studies in the housing literature highlighting the roles played by both demand-side and supply-side factors in driving housing cycles. On the demand side, Himmelberg, Mayer, and Sinai (2006) focus on interest rates, Poterba (1991) on tax changes, and Mian and Sufi (2009) on credit expansion. On the supply side, Glaeser, Gyourko, Saiz (2008) emphasize supply as a key force in mitigating housing bubbles, Haughwout, Peach, Sporn and Tracy (2012) provide a detailed account of the housing supply side during the U.S. housing cycle in the 2000s, and Gyourko (2009) systematically reviews the literature on housing supply. By introducing informational frictions, our analysis shows that supply-side and demand-side factors are not mutually independent. In particular, supply shocks may affect housing demand by acting as informational noise in household learning and thus influencing households' expectations of the strength of the neighborhood.

3.3. Some Basic Facts

Before we present a model to analyze how supply elasticity affects learning in housing markets, we present some basic facts regarding the relationship between supply elasticity and the magnitudes of housing price booms and busts experienced by different counties during the recent U.S. housing cycle. Even though common wisdom holds that supply elasticity attenuates boom and bust cycles, the data do not support a robust, monotonic relationship between the magnitude of the housing cycle in a county and its supply elasticity. In fact, our analysis uncovers that counties with supply elasticities in an intermediate range had experienced more dramatic housing booms and busts than counties with the most inelastic supply.

Our focus is on the most recent housing cycle, which was a national cycle for the U.S. housing market. Many factors, such as the Clinton-era initiatives to broaden homeownership, the low interest rate environment of the late 1990s and early 2000s, the inflow of foreign capital, and the increase in securitization and sub-prime lending, contributed to the boom. While this was a well-known national phenomenon at the time, how these factors expressed themselves at the regional level was more idiosyncratic and uncertain. The magnitude of housing price cycles experienced by different regions reflect such idiosyncratic uncertainty, which is the focus of our empirical as well as theoretical analysis.[62]

Our county-level house price data come from the Case-Shiller home price indices, which are constructed from repeat home sales. There are 420 counties in 46 states with a large enough number of repeat home sales to compute the Case-Shiller home price indices. We

[62] The regional uncertainty introduced by this national phenomenon is absent from the local boom and bust episodes throughout the 1970s and 1980s. While there are other national housing cycles in history, data limitations restrict our attention to the most recent U.S. housing cycle.

use the Consumer Price Index (CPI) from the Bureau of Labor Statistics to deflate the Case-Shiller home price indices. In addition, we also use population data from the 2000 U.S. census.

For housing supply elasticity, we employ the commonly used elasticity measure constructed by Saiz (2010). This elasticity measure focuses on geographic constraints by defining undevelopable land for construction as terrain with a slope of 15 degrees or more and areas lost to bodies of water including seas, lakes, and wetlands. This measure provides an exogenous measure of supply elasticity, with a higher value if an area is more geographically restricted. Saiz's measure is available for 269 Metropolitan Statistical Areas (MSAs). By matching counties with MSAs, our sample includes 326 counties for which we have data on both house prices and supply elasticity available from 2000 to 2010. Though our sample covers only 11 percent of the counties in the U.S., they represent 53 percent of the U.S. population and 57 percent of the housing trading volume in 2000.

Figure 1 displays the real home price indices for the U.S. and three cities, New York, Las Vegas, and Charlotte, from 2000 to 2010. We normalize all indices to 100 in 2000. The national housing market experienced a significant boom and bust cycle in the 2000s with the national home price index increasing over 60 percent from 2000 to 2006 and then falling back to the 2000 level through 2010. Different cities in the U.S experienced largely synchronized price booms and busts during this period even though the magnitudes of the cycle varied across these cities. According to Saiz's measure, the elasticity measures for New York, Las Vegas, and Charlotte are 0.76, 1.39, and 3.09, respectively. New York, which has severe geographic constraints and building regulations, had a real housing price appreciation of more than 80 percent during the boom, and then declined by over 25 percent during the bust. Charlotte, with its vast developable land and few building restrictions, had an almost flat real housing price level throughout this decade. Sitting in between New York and Charlotte, Las Vegas, with its intermediate supply elasticity, experienced the most pronounced price expansion of over 120 percent during the boom, and the most dramatic price drop of over 50 percent during the bust. Many commentators, including Glaeser (2013), have pointed out that the dramatic boom and bust cycles experienced by Las Vegas and other cities such as Phoenix are peculiar given the relatively elastic supply in these areas.

Are Las Vegas and Phoenix unique in experiencing these dramatic housing cycles despite their relatively elastic housing supply? We now systematically examine this issue by sorting different counties in our sample into three groups, an inelastic, a middle, and an elastic group, based on Saiz's elasticity measure, each with the same number of counties. Figure 2 plots the average price expansion and contraction experienced by each group during the housing cycle (the top panel), together with the total population in each group (the bottom panel). We measure the price expansion in 2004-2006, the period that is often defined as the housing bubble period, and the price contraction in 2007-2009.[63]

The top panel of Figure 2 shows that the inelastic group had the largest house price expansion in 2004-2006 and the largest price contraction in 2007-2009, the middle group experienced a milder cycle, and the elastic group had the most modest cycle. This pattern appears to be consistent with the aforementioned common wisdom that supply elasticity attenuates housing cycles.

It seems natural to sort the counties into several groups each with an equal number of counties. In fact, this is a common practice used in the literature to demonstrate a monotonic relationship between housing cycles and supply elasticity. Interestingly, the bottom panel of Figure 2 shows that the population is unevenly distributed across the three groups, with the inelastic group having more than half of the total population. This is consistent with the fact that inelastic areas tend to be densely populated. As the inelastic group pools together a large fraction of the population, there might be substantial heterogeneity between counties within the inelastic group. Indeed, both New York and Las Vegas fall into this inelastic group. This consideration motivates us to examine alternative ways of grouping the counties.

In Figure 3, we sort the counties into ten groups from the most inelastic group to the most elastic group, still with each group holding an equal number of counties. The top panel shows that the housing price expansion and contraction experienced by these ten groups are no longer monotonic with elasticity. In particular, group 3, which has the third-most-inelastic supply, experienced the largest price expansion during the boom, and the largest price contraction during the bust. Interestingly, Las Vegas falls into group 3, while New York falls into group 1. The bottom panel again shows that the population tends to be concentrated in the more inelastic groups. Taken together, Figure 3 shows that the commonly perceived monotonic relationship between housing cycles and supply elasticity is not robust.

In Figure 4, we sort the counties into three groups based on supply elasticity in an

[63] We have also used an alternative boom period of 2001-2006 and obtained qualitatively similar results as defining the boom in 2004-2006.

alternative way. Instead of letting different groups have an equal number of counties, we let them have the same population. If the magnitude of the housing cycle monotonically decreases with supply elasticity, whether we group the counties by number or population should not affect the monotonically decreasing pattern across the groups. In contrast, the top panel of Figure 4 shows that the middle group has the most pronounced housing cycle, with its price expansion during the boom being substantially more pronounced than that of the inelastic group, and its price contraction during the bust slightly greater than that of the inelastic group. The bottom panel shows that the inelastic group has only 40 counties, the middle group slightly below 120 counties, and the elastic group over 160 counties. Under this grouping, while New York remains in the inelastic group, Las Vegas is now in the middle group.

In Figure 5, we further sort the counties into ten groups from the most inelastic group to the most elastic group, with each group having the same population. This figure shows a finer non-monotonicity with groups 3 and 5 experiencing the most pronounced price expansions and contractions.

To further examine whether the more pronounced housing cycles experienced by the intermediate elasticity groups are robust to controlling for other fundamental factors, such as changes of income and population and fraction of subprime households, we adopt a regression approach. Specifically, we separately regress the housing price expansion in 2004-2006 and contraction in 2007-2009 on two dummy variables that indicate whether a county is in the middle elasticity group or the most elastic group, which are constructed in Figure 4, together with a list of control variables. This regression implicitly uses the inelastic group as the benchmark for the middle and elastic groups. The control variables include the fraction of subprime households in the county in 2005, which is computed based on individual mortgage loan applications reported by the "Home Mortgage Disclosure Act" (HMDA) dataset, as well as the contemporaneous population change and annualized per capita income change.

Table 1 reports the regression results. Columns 1 and 2 report the regressions of the housing price expansion in 2004-2006 without and with the controls. Columns 3 and 4 report the regressions of the housing price contraction in 2007-2009 without and with the controls. Among the control variables, the fraction of subprime households is significantly correlated with both the price expansion during the boom and the price contraction during the bust. This result is consistent with Mian and Sufi (2009), which shows that credit expansion to subprime households before 2006 was a key factor in explaining the recent housing cycle. The changes in population and income are insignificant in explaining either the price expansion or the contraction across the cycle. More important, even after controlling for these fundamental factors, the middle group experienced a significantly more pronounced housing price expansion in 2004-2006 and a more pronounced price contraction in 2007-2009 relative to the inelastic group.

It is important to note that the findings shown in Figures 2-5 and Table 1 are not driven by a few areas such as Las Vegas and Phoenix. In unreported analysis, we dropped Las Vegas and Phoenix and found that results are similar both qualitatively and quantitatively. These robustness results are available upon request.

Taken together, Figures 2-5 and Table 1 show that the commonly perceived monotonic relationship between housing cycle and supply elasticity is not robust to finer groupings of counties. Finer groupings and an alternative method of grouping counties by population rather than number of counties, however, reveal a robust non-monotonic relationship in which the counties in a median elasticity range experienced more pronounced price booms and busts in the 2000s than counties with the most inelastic supply. This non-monotonic relationship is intriguing and cannot easily be explained by the usual role of elasticity in affecting the supply side of housing. In the next section, we present a simple model to illustrate a learning mechanism through which supply elasticity affects the informational role of housing prices and households' learning from housing prices.

3.4. A Baseline Model

In this section we develop a simple model with two dates $t = 1, 2$ to analyze the effects of informational frictions on the housing market equilibrium in a closed neighborhood. One can think of this neighborhood as a county or a township. A key feature of the model is that the housing market is not only a place for households to trade housing but also a platform to aggregate private information about the unobservable strength or quality of the neighborhood. In addition to its direct role in affecting housing supply in the neighborhood, supply elasticity also indirectly affects the informational noise in the housing prices.

3.4.1 Model Setting

There are two types of agents in the economy: households looking to buy homes in the neighborhood and home builders. Suppose that the neighborhood is new and all households purchase houses from home builders in a centralized market at $t = 1$ and consume both housing and consumption goods at $t = 2$.[64]

Each household cares about the strength of the neighborhood, as its utility depends on not only its own housing consumption but also the housing consumption of other households in their neighborhood. This assumption is motivated by the empirical findings of Ioannides and Zabel (2003) and leads to strategic complementarity in each household's housing demand.[65] The strength of this closed neighborhood is reflected by the aggregate productivity of its households. A strong aggregate productivity implies greater output by all households, and thus greater housing demand by them as well. In the presence of realistic informational frictions in gauging the strength of the neighborhood, the housing market provides an important platform for aggregating information about this aggregate productivity. As a consequence, the resulting housing price serves as a useful signal about the neighborhood's strength.

3.4.1.1 Demand Side

There is a continuum of households, indexed by $i \in [0, 1]$. Household i has a Cobb-Douglas utility function over its own housing H_i, consumption good C_i, and the housing consumption of all other households in the neighborhood $\{H_j\}_{j\in[0,1]}$:[66]

$$U\left(\{H_j\}_{j\in[0,1]}, C_i\right) = \left\{\frac{1}{1-\eta_H}\left(\frac{H_i}{1-\eta_c}\right)^{1-\eta_c}\left(\frac{\int_{[0,1]/i} H_j dj}{\eta_c}\right)^{\eta_c}\right\}^{1-\eta_H}\left(\frac{1}{\eta_H}C_i\right)^{\eta_H}. \quad (33)$$

The parameters $\eta_H \in (0, 1)$ and $\eta_c \in (0, 1)$ measure the weights of different consumption components in the utility function. A higher η_H means a stronger complementarity between

[64] For simplicity, we do not consider the endogenous decision of households choosing their neighborhood and instead take the pool of households in the neighborhood as given. See Van Nieuwerburgh and Weill (2010) for a systematic treatment of moving decisions by households across neighborhoods.

[65] There are other types of social interactions between households living in a neighborhood, which are explored, for instance, in Durlauf (2004) and Glaeser, Sacerdote, and Scheinkman (2003).

[66] Our modeling choice of non-separable preferences for housing and consumption is similar to the CES specification of Piazzesi, Schneider, and Tuzel (2007).

housing consumption and goods consumption, while a higher η_c means a stronger complementarity between the housing of household i and housing of the composite house $\int_{[0,1]/i} H_j dj$ purchased by the other households in the neighborhood.

The production function of household i is $e^{A_i} l_i$, where l_i is the household's labor choice and A_i is its productivity. A_i is comprised of a component A common to all households in the neighborhood and an idiosyncratic component ε_i:

$$A_i = A + \varepsilon_i,$$

where $A \sim \mathcal{N}\left(\bar{A}, \tau_A^{-1}\right)$ and $\varepsilon_i \sim \mathcal{N}\left(0, \tau_\varepsilon^{-1}\right)$ are both normally distributed. The common productivity A represents the strength of the neighborhood, as a higher A implies a more productive neighborhood. As A determines the households' aggregate demand for housing, it represents the demand-side fundamental.

As a result of realistic informational frictions, neither A nor A_i is observable to the households. Instead, each household observes a noisy private signal about A at $t = 1$. Specifically, household i observes

$$\theta_i = A + \nu_i,$$

where $\nu_i \sim \mathcal{N}\left(0, \tau_\theta^{-1}\right)$ is signal noise independent across households. The parameter τ_θ measures the precision of the private signal. As $\tau_\theta \to \infty$, the households' signals become infinitely precise and the informational frictions about A vanish.

While our model setting is static and focuses on a closed neighborhood, one can provide a broad interpretation of the uncertainly in the neighborhood strength A. In relating this setting to the recent housing cycle, we interpret the uncertainty in A as being induced by a nationwide shock to credit expansion for homeowners in the 2000's due to the factors mentioned in Section 3.3. While this shock affected the whole country, its potential impact on individual neighborhoods was different. Some neighborhoods might attract migrants with high productivity from other neighborhoods as a result of the national shock, while others might lose their high-quality residents who could now more easily reallocate to other locations. As a result, home buyers faced a realistic problem in inferring how this national shock might have influenced the strength of an individual neighborhood when buying a home.

Households care about the strength of the neighborhood A not only because it determines their own productivity, but also because of complementarity in their housing demand. Since households want to live in similar-sized houses to their neighbors, they need to learn about

A because it affects their neighbors' housing decisions. Consequently, while a household may have a fairly good understanding of its own productivity when moving into a neighborhood, complementarity in housing demand motivates it to pay attention to housing prices to learn about the average level A for the neighborhood.

We assume that each household experiences a disutility for labor $\frac{l_i^{1+\psi}}{1+\psi}$, and that it maximizes its expected utility at $t = 1$ by choosing its housing demand H_i and labor l_i:

$$\max_{\{H_i, l_i\}} E\left[U\left(\{H_j\}_{j\in[0,1]}, C_i \right) - \frac{l_i^{1+\psi}}{1+\psi} \middle| \mathcal{I}_i \right] \tag{34}$$

$$\text{such that} \quad C_i = e^{A_i} l_i - PH_i + \Pi_i.$$

We assume for simplicity that the home builder for household i is part of the household and that the builder brings home its profit $\Pi_i = PH_i$ to the household after construction has taken place. Furthermore, we normalize the interest rate from $t = 1$ to $t = 2$ to be zero. As a result, at $t = 2$, household i's budget constraint satisfies $C_i = e^{A_i} l_i$. The choices of labor and housing are made at $t = 1$ subject to each household's information set $\mathcal{I}_i = \{\theta_i, P\}$, which includes its private signal θ_i and the housing price P.[67]

3.4.1.2 Supply Side

Home builders face a convex labor cost

$$\frac{k}{1+k} e^{-\zeta} S_H^{\frac{1+k}{k}}$$

in supplying housing, where S_H is the quantity of housing supplied, $k \in (0, \infty)$ is a constant parameter, and ζ represents a shock to the building cost. We assume that ζ is observed by builders but not households,[68] and that from the perspective of households $\zeta \sim \mathcal{N}\left(\bar{\zeta}, 1\right)$, i.e., a normal distribution with $\bar{\zeta}$ as the mean and unit variance.

[67] We do not include the volume of housing transactions in the information set as a result of a realistic consideration that, in practice, people observe only delayed reports of total housing transactions at highly aggregated levels, such as national or metropolitan levels.

[68] Even though we assume that builders perfectly observe the supply shock, a more realistic setting would have builders each observing part of the supply and thus needing to aggregate their respective information in order to fully observe the supply-side shock. We have explored this more general setting, which entails an additional layer of information aggregation on the builder side of the housing market. Nevertheless, it gives qualitatively similar insights as our current setting.

Builders at $t = 1$ maximize their profit subject to their supply curve

$$\Pi\left(S_H\right) = \max_{S_H} \; PS_H - \frac{k}{1+k} e^{-\zeta} S_H^{\frac{1+k}{k}}. \tag{35}$$

It is easy to determine the builders' optimal supply curve:

$$S_H = P^k e^{\xi}, \tag{36}$$

where $\xi = k\zeta$ is interpretated as being a supply shock with normal distribution $\xi \sim \mathcal{N}\left(\bar{\xi}, k^2\right)$, where $\bar{\xi} = k\bar{\zeta}$. The parameter k measures the supply elasticity of the neighborhood. A more elastic neighborhood has a larger supply shock, i.e., the supply shock has greater mean and variance. In the housing market equilibrium, the supply shock ξ not only affects the supply side but also the demand side, as it acts as informational noise in the price signal when the households use the price to learn about the common productivity A.

We also incorporate a behavioral feature that households may underestimate the supply elasticity in the neighborhood and incorrectly believe it to be ϕk rather than k, where $\phi \leq 1$. This feature is motivated by the observation made by Glaeser (2013) that agents tend to underestimate supply shocks during various episodes of real-estate speculation observed in U.S. history. Specifically, Glaeser identifies the under-appreciation of the supply response by buyers as a systematic, cognitive limitation that helps explain historical boom and bust episodes of real-estate speculation. As a result of this behavioral feature, households may put too much weight on housing prices in their housing decisions because they overestimate the precision of prices as a signal about the neighborhood strength A.

This overweighing is reminiscent of extrapolative beliefs, which, as referenced in the introduction, have been recognized in the literature as important for understanding housing cycles. It is useful to note that while extrapolative beliefs amplify housing price volatility, they nevertheless imply price volatility monotonically decreases in supply elasticity. The overweighing of prices highlighted in our model anchors on a characteristic of neighborhoods that allows for it to explain the hump-shaped pattern in price volatility across supply elasticity.

3.4.2 The Equilibrium

Our model features a noisy rational expectations equilibrium, which requires clearing of the housing market that is consistent with the optimal behavior of both households and home builders:

- Household optimization: $\left\{ \{H_i\}_{i\in[0,1]}, l_i \right\}$ solves each household's maximization problem in (34).

- Builder optimization: S_H solves the builders' maximization problem in (35).

- At $t = 1$, the housing market clears:

$$\int_{-\infty}^{\infty} H_i(\theta_i, P)\, d\Phi(\nu_i) = P^k e^{\xi},$$

  where each household's housing demand $H_i(\theta_i, P)$ depends on its private signal θ_i and the housing price P. The demand from households is integrated over the idiosyncratic component of their private signals $\{\nu_i\}_{i\in[0,1]}$.

We first solve for the optimal labor and housing choices for a household given its utility function and budget constraint in (34), which are characterized in the following proposition.

PROPOSITION 3.1: Household i's optimal labor choice depends on its expected productivity:

$$l_i = \left\{ \eta_H E\left[\left\{ \frac{1}{1-\eta_H} \left(\frac{H_i}{1-\eta_c} \right)^{1-\eta_c} \left(\frac{\int_{[0,1]/i} H_j dj}{\eta_c} \right)^{\eta_c} \right\}^{1-\eta_H} \left(\frac{1}{\eta_H} e^{A_i} \right)^{\eta_H} \middle| \mathcal{I}_i \right] \right\}^{\frac{1}{1+\psi-\eta_H}},$$

and its demand for housing is

$$\begin{aligned} \log H_i &= \frac{2+\psi-\eta_H}{\psi+(1-\eta_H)\eta_c} \log E\left[\left(\int_{[0,1]} H_j dj \right)^{\eta_c(1-\eta_H)} e^{\eta_H A_i} \middle| \mathcal{I}_i \right] - \frac{1+\psi-\eta_H}{\psi+(1-\eta_H)\eta_c} \log P \\ &\quad - \frac{1+\psi-\eta_H}{\psi+(1-\eta_H)\eta_c} \log E\left[\left(\int_{[0,1]} H_j dj \right)^{\eta_c(1-\eta_H)} e^{(\eta_H-1)A_i} \middle| \mathcal{I}_i \right] \\ &\quad + \frac{\psi}{\psi+(1-\eta_H)\eta_c} \log\left((1-\eta_c) \left(\frac{1-\eta_H}{\eta_H} \right) \left(\frac{1-\eta_c}{\eta_c} \right)^{\frac{\eta_c(1-\eta_H)}{\psi}} \right). \end{aligned} \tag{37}$$

Proposition 1 demonstrates that the labor chosen by a household is determined by its expected productivity and that its housing demand is determined by not only its own productivity e^{A_i} but also the aggregate housing consumption of other households. This latter component arises from the complementarity in the utility function of the household.

By clearing the aggregate housing demand of the households with the supply from the builders, we derive the housing market equilibrium. Despite the nonlinearity in each household's demand and in the supply from builders, we obtain a tractable unique log-linear equilibrium. The following proposition summarizes the housing price and each household's housing demand in this equilibrium.

PROPOSITION 3.2: At $t = 1$, the housing market has a unique log-linear equilibrium:
1) The housing price is a log-linear function of A and ξ:

$$\log P = p_A A + p_\xi \xi + p_0, \tag{38}$$

with the coefficients p_A and p_ξ given by

$$p_A = \frac{1+\psi}{1+\psi(1+k)-\eta_H} - \frac{\psi+(1-\eta_H)\eta_c}{1+\psi(1+k)-\eta_H}\tau_\theta^{-1}\tau_A b > 0, \tag{39}$$

$$p_\xi = -\frac{\psi}{1+\psi(1+k)-\eta_H} - \frac{\psi+(1-\eta_H)\eta_c}{1+\psi(1+k)-\eta_H}\tau_\theta^{-1}\left(\frac{b}{\phi k}\right)^2 < 0, \tag{40}$$

where $b \in \left[0, \frac{1+\psi}{\psi}\frac{\tau_\theta}{\left(1+\frac{(1-\eta_H)\eta_c}{\psi}\right)\tau_A+\tau_\theta}\right]$ is the unique positive, real root of equation (61), and p_0 is given in equation (66).

2) The housing demand of household i is a log-linear function of its private signal θ_i and $\log P$:

$$\log H_i = h_\theta \theta_i + h_P \log P + h_0, \tag{41}$$

with the coefficients h_θ and h_P given by

$$h_\theta = b > 0, \tag{42}$$

$$h_P = -\frac{1+\psi-\eta_H}{\psi} + \frac{1+\psi+\eta_c(1-\eta_H)b}{\psi}\frac{\left(\frac{b}{\phi k}\right)^2}{\tau_A+\tau_\theta+\left(\frac{b}{\phi k}\right)^2}\frac{1}{p_A}, \tag{43}$$

and h_0 given by equation (55).

Proposition 3.2 establishes that the housing price P is a log-linear function of the neighborhood strength A and the housing supply shock ξ, and that each household's housing demand is a log-linear function of its private signal θ_i and the log housing price $\log P$. Similar to Hellwig (1980), the housing price aggregates the households' dispersed private information to partially reveal A. The price does not depend on the idiosyncratic noise in any individual household's signal because of the Law of Large Numbers. This last observation is key to the tractability of our model and ensures that the housing demand from the households retains a log-normal distribution after aggregation. This tractable log-linear equilibrium is different from the frameworks of Goldstein, Ozdenoren, and Yuan (2013) and Albagi, Hellwig, and Tsyvinski (2012, 2014), which employ risk-neutral agents, normally distributed asset fundamentals, and position limits to deliver tractable nonlinear equilibria.

In the presence of informational frictions, the housing supply shock ξ serves the same role as noise trading in standard models of asset market trading with dispersed information. This feature is new to the housing literature and highlights an important channel for supply shocks to affect the expectations of potential home buyers. Since households cannot perfectly disentangle changes in housing prices caused by supply shocks from those brought about by shocks to demand, they partially confuse a housing price change caused by a supply shock to be a signal about the strength of the neighborhood.

To facilitate our discussion of the impact of learning, it will be useful to introduce a perfect-information benchmark in which all households perfectly observe the strength of the neighborhood A. The following proposition characterizes this benchmark equilibrium.

PROPOSITION 3.3: Consider a benchmark setting, in which households perfectly observe A (i.e., $\theta_i = A$, $\forall i$.) There is also a log-linear equilibrium, in which the housing price is

$$\begin{aligned}\log P &= \frac{1+\psi}{1+\psi(1+k)-\eta_H}A - \frac{\psi}{1+\psi(1+k)-\eta_H}\xi + \frac{(1+\psi)\eta_H - (1-\eta_H)(1+\psi-\eta_H)}{2(1+\psi(1+k)-\eta_H)}\tau_\varepsilon^{-1} \\ &\quad + \frac{\psi}{1+\psi(1+k)-\eta_H}\log\left((1-\eta_c)\left(\frac{1-\eta_H}{\eta_H}\right)\left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}}\right).\end{aligned}$$

and all households have the same housing demand

$$\begin{aligned}\log H &= \frac{1+\psi}{\psi}A - \frac{1+\psi-\eta_H}{\psi}\log P + \frac{(1+\psi)\eta_H - (1-\eta_H)(1+\psi-\eta_H)}{2\psi}\tau_\varepsilon^{-1} \\ &\quad + \log\left((1-\eta_c)\left(\frac{1-\eta_H}{\eta_H}\right)\left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}}\right).\end{aligned}$$

Furthermore, the housing market equilibrium with informational frictions characterized in Proposition 2 converges to this benchmark equilibrium as $\tau_\theta \nearrow \infty$, and the variance of the housing price $Var[\log P]$ has a U-shaped relationship with the supply elasticity k.

It is reassuring that as the households' private information becomes infinitely precise, the housing market equilibrium converges to the perfect-information benchmark. In this perfect-information benchmark, the housing price is also a log-linear function of the demand-side fundamental A and the supply shock ξ, and each household's identical demand is a log-linear function of the perfectly observed A and the housing price $\log P$. Consistent with the standard intuition, a higher A increases both the housing price and aggregate housing demand, while a larger supply shock ξ reduces the housing price but increases aggregate housing demand. It is also easy to see that in this benchmark setting, as the supply elasticity k rises from zero to infinity, the weight of A (the demand-side fundamental) in the housing price decreases, while the weight of ξ (the supply-side shock) increases.

Furthermore, in the perfect-information benchmark, the housing price variance has a U-shaped relationship with the housing supply elasticity k. This is because, as k varies, it causes the housing price to assign different weights to the demand-side fundamental and the supply-side shock. The standard intuition from diversification implies that the price has the lowest variance when the weights of the two factors are balanced, i.e., the supply elasticity takes an intermediate value. This U-shaped price variance serves a benchmark to evaluate the housing price variance in the presence of informational frictions.

3.4.3 Impact of Learning

In the presence of informational frictions about the strength of the neighborhood A, each household needs to use its private signal θ_i and the publicly observed housing price $\log P$ to learn about A. As the housing price $\log P$ is a linear combination of the demand-side fundamental A and the housing supply shock ξ, the supply shock interferes with this learning process. A larger supply shock ξ, by depressing the housing price, will have an additional effect of reducing the households' expectations of A. This, in turn, reduces their housing demand and consequently further depresses the housing price. This learning effect thus causes the supply shock to have a larger negative effect on the equilibrium housing price than it would in the perfect-information benchmark. Similarly, this learning effect also causes the demand-side fundamental A to have a smaller positive effect on the price than in the perfect-information benchmark because informational frictions cause households to partially discount the value of A. The following proposition formally establishes this learning effect on the housing price.

PROPOSITION 3.4: In the presence of informational frictions, coefficients $p_A > 0$ and $p_\xi < 0$ derived in Proposition 3.2 are both lower than their corresponding values in the perfect-information benchmark.

The precision of the households' private information τ_θ determines the informational frictions they face. The next proposition establishes that an increase in τ_θ mitigates the informational frictions and brings the coefficient p_A closer to its value in the perfect-information benchmark. In fact, as τ_θ goes to infinity, the housing market equilibrium converges to the perfect-information benchmark (Proposition 3).

PROPOSITION 3.5: p_A increases with the precision τ_θ of each household's private signal and decreases with the degree of complementarity in households' housing consumption η_c.

Each household's housing demand also reveals how the households learn from the housing price. In the presence of informational frictions about A, housing price is not only the cost of acquiring shelter but also a signal about A. The housing demand of each household derived in (41) reflects both of these effects. Specifically, we can decompose the price elasticity of each household's housing demand h_P in equation (43) into two components. The first component $-\frac{1}{\eta_H}$ is negative and represents the standard cost effect (i.e., downward sloping demand curve), as in the perfect-information benchmark in Proposition 3.3, and the second component $\frac{1+\psi+\eta_c(1-\eta_H)b}{\psi} \frac{\left(\frac{b}{\phi k}\right)^2}{\tau_A+\tau_\theta+\left(\frac{b}{\phi k}\right)^2} \frac{1}{p_A}$ is positive and represents the learning effect.

A higher housing price raises the household's expectation of A and induces it to consume more housing through two related but distinct learning channels. First, a higher A implies a higher productivity for the household itself. Second, a higher A also implies that other households demand more housing, which in turn induces each household to demand more housing. As a reflection of this complementarity effect, the second component in the price elasticity of housing demand increases with η_c, the degree of complementarity in the household's utility of its own housing consumption and other households' housing consumption.

As a result of the presence of the complementarity channel, η_c also affects the impact of learning on the housing price. As η_c increases, each household puts a greater weight on the housing price in its learning of A and a smaller weight on its own private signal. This in turn makes the housing price less informative of A. In this way, a larger η_c exacerbates the informational frictions faced by households. Indeed, Proposition 5 shows that the loading of $\log P$ on A is decreasing with η_c.

Housing supply elasticity k plays an important role in determining the informational frictions faced by the households, in addition to its standard supply effect. To illustrate this learning effect of supply elasticity, we consider two limiting economies as k goes to 0 and ∞, which are characterized in the following proposition.

PROPOSITION 3.6: As $k \to \infty$, the housing price and each household's housing demand converge to

$$\log P = -\zeta,$$

and

$$\log H_i = \frac{1+\psi}{\psi} \frac{\tau_\theta}{\tau_A + \tau_\theta + \frac{(1-\eta_H)\eta_c}{\psi}\tau_A}\theta_i - \frac{1+\psi-\eta_H}{\psi}\log P + h_0.$$

As $k \to 0$, the housing price and each household's housing demand converge to

$$\begin{aligned}\log P &= \frac{1+\psi}{1+\psi-\eta_H}A + \frac{1}{2}\frac{(1+\psi)\,\eta_H - (1-\eta_H)\,(1+\psi-\eta_H)}{1+\psi-\eta_H}\tau_\varepsilon^{-1} \\ &\quad + \frac{\psi}{1+\psi-\eta_H}\log\left((1-\eta_c)\left(\frac{1-\eta_H}{\eta_H}\right)\left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}}\right),\end{aligned}$$

and $\log H_i = 0$.

At one end, as supply elasticity goes to zero, the housing price is completely driven by A and thus fully reveals it. In this case, each household precisely learns A from the price, and as a result, both the housing price and each household's housing demand coincide with their corresponding values in the prefect-information benchmark. At the other end, as supply elasticity goes to infinity, the housing price is completely driven by the supply shock ξ and contains no information about A. In this case, each household has to rely on its own private signal to infer A. But as the housing price is fully determined by the supply shock and independent of the demand-side fundamental, informational frictions about A do not matter for the housing price. Consequently, the housing price also coincides with that in the perfect-information benchmark, even though informational frictions still affect each household's housing demand. Taken together, when housing supply is either perfectly elastic or inelastic, housing price is not affected by informational frictions and coincides with that in the perfect-information benchmark.

The following proposition characterizes the housing price at an intermediate supply elasticity and, in particular, analyzes the role of the households' underestimation ϕ of supply elasticity.

PROPOSITION 3.7: Consider an intermediate level of supply elasticity $k \in (0, \infty)$. 1) In the presence of informational frictions, both p_A and $|p_\xi|$ monotonically decrease with ϕ. 2)

When $\phi = 1$, the housing price variance with informational frictions is lower than that of the perfect-information benchmark. 3) The variance of the housing price $\log P$ decreases with ϕ, and a sufficient condition $1 - \frac{\psi}{\psi+(1-\eta_H)\eta_c}\frac{\tau_\theta}{\tau_A} \leq \phi^2 \leq \frac{1}{2}$ ensures the price variance to be at least as large as its corresponding value in the perfect-information benchmark.

Proposition 3.7 shows that, at an intermediate supply elasticity, the households' underestimation of supply elasticity causes them to overinterpret the information contained in the price signal and thus overreact to the price signal. Consequently, the positive loading of the equilibrium housing price p_A on the demand-side fundamental A becomes larger and the negative loading p_ξ on the supply shock becomes more negative. That is, the housing price becomes more responsive to both demand and supply shocks.

Proposition 3.7 also shows that, in the absence of the households' underestimation of supply elasticity, the presence of informational frictions reduces the housing price variance. This is because informational frictions make households less responsive to demand shocks, causing the housing price to load less on demand shocks. When households underestimate the supply elasticity ($\phi < 1$), their overreaction to the price signal amplifies the price effects of both supply and demand shocks, and implies that the housing price variance monotonically decreases with ϕ. In fact, Proposition 3.7 shows that when ϕ is sufficiently small, the housing price variance is at least as large as its value in the perfect-information benchmark.

Interestingly, the volatility amplification induced by the households' overreaction to the housing price is most pronounced when the supply elasticity is in an intermediate range. This follows from our earlier discussion of the two limiting cases when the elasticity goes to either zero or infinity. At one end, when the supply is infinitely elastic, the households' learning about the demand side is irrelevant for the price. At the other end, when the supply is infinitely inelastic, the price fully reveals the demand-side fundamental and there is no room for the households to overreact. In between these two limiting cases, the demand-side fundamental plays a significant role in determining the housing price and at the same time households face substantial uncertainty about the demand-side fundamental, which leaves room for their overreaction to amplify the price volatility.

In Figure 6, we provide a numerical example to illustrate how informational frictions and households' overreaction jointly affect the housing price variance. The figure depicts the log-price variance $Var[\log P]$ against the supply elasticity under the following parameter values:

$$\tau_\theta = 0.1,\ \tau_A = 1, \phi = 0.1,\ \psi = 0.6, \eta_c = 0.5, \eta_H = 0.9.$$

For comparison, it also depicts the log-price variance in the perfect-information benchmark, which is obtained as $\tau_\theta \to \infty$. As the supply elasticity k rises from 0 to 1 (i.e., from infinitely inelastic to more elastic), the log-price variance decreases with the supply elasticity. In contrast, when the households face informational frictions with $\tau_\theta = 1$, Figure 6 shows that the log-price variance first increases with k, when k is lower than an intermediate level around 0.1, and then decreases with k.[69] The difference between this humped shape and the monotonically decreasing curve in the perfect-information benchmark illustrates the joint effect of informational frictions and the households' overreaction to the price signal.

The humped log-price variance illustrated in Figure 6 provides an explanation for the aforementioned, non-monotonic relationship between the housing boom and bust cycles experienced by different U.S. counties in the 2000s and supply elasticity.

3.5 Elasticity and Housing Speculation

In this section, we further explore the effects of household learning on housing speculation. We first extend the baseline model presented in the last section to incorporate secondary homes and, in particular, to show that the same learning effect discussed earlier leads to new predictions regarding the relationship between housing speculation and supply elasticity. Then, we examine these predictions in the data and provide some supportive evidence.

3.5.1 A Model Extension

We extend the model presented in the previous section to incorporate three types of agents in the economy: households, home builders, and immigrants looking to move into the neighborhood. The immigrants are the new addition to this extension. Suppose that these immigrants make their decision at $t = 1$ on whether to move into the neighborhood, based on the expected strength of the neighborhood. The immigrants arrive in the neighborhood at $t = 2$ and then buy the secondary homes initially owned by the households.

[69] Outside the range of k depicted in the figure, both of these two lines decrease and eventually converge to each other as $k \to \infty$, as derived in Proposition 3.6.

3.5.1.1 Households

At $t = 1$, households purchase two types of homes, one as their primary residence and the other as a secondary home to sell at $t = 2$ to the immigrants. Home builders build and sell these two types of homes in two separate housing markets. This separate treatment of primary and secondary homes is consistent with the fact that, in practice, primary homes tend to be single houses, while secondary homes tend to be apartments and condominiums. Another advantage of giving separate supply curves to primary and secondary homes is that it ensures a tractable log-linear equilibrium.

When making their decisions at $t = 1$, households again receive a private signal θ_i about the strength of the neighborhood. Like before, the demand of household i for a primary home is H_i, but now, in addition, the household has a demand for a secondary home, M_i. For simplicity, suppose that households have no initial wealth and must finance their purchases by borrowing debt D_i from home builders.[70] We also normalize the interest rate on the loans to be zero. Then, the budget constraint of household i at $t = 1$ is

$$PH_i + Q_1 M_i = D_i, \tag{44}$$

where P is the price of primary homes and Q_1 is the price of secondary homes.

At $t = 2$, households decide how much of their goods to produce, sell their secondary homes at a price Q_2 to the immigrants that have moved into the neighborhood, and repay their debt to home builders. Household i again employs its own labor l_i as an input to production with production function $e^{A_i} l_i$. As in the baseline model, household i earns income Π_{Hi} from the home builder of their primary home, who is part of the household. Thus, in equilibrium, $\Pi_{Hi} = PH_i$.

The budget constraint of household i at $t = 2$ is then

$$C_i = -D_i + e^{A_i} l_i + Q_2 M_i + \Pi_{Hi}, \tag{45}$$

where C_i is the goods consumption of household i. Households have the same Cobb-Douglas preferences as in the baseline model for consuming their primary housing and non-housing consumption at $t = 2$.

[70] This assumption is innocuous as our main interest is not to study the effects of the households' credit constraints.

At $t = 1$, each household maximizes

$$\max_{\{H_i, M_i, l_i\}} E\left[U\left(\{H_j\}_{j\in[0,1]}, C_i \right) - \frac{l_i^{1+\psi}}{1+\psi} \middle| \mathcal{I}_i \right] \tag{46}$$

under its information set $\mathcal{I}_i = \{\theta_i, P, Q_1\}$, which includes its private signal θ_i and the housing prices P and Q_1. The household's consumption C_i is determined by its budget constraint at $t = 2$ given in (45), which in turn depends on its budget constraint at $t = 1$ given in (44).

3.5.1.2 Home Builders

Home builders face separate production processes for building primary and secondary homes. Specifically, they face the following convex labor cost for building each type of home:

$$\frac{k}{1+k} e^{-\zeta_j} S_j^{\frac{1+k}{k}}$$

where $j \in \{H, M\}$ indicates the type of home with H representing primary homes and M representing secondary homes, S_j is the quantity of type-j homes supplied, and ζ_j represents a supply shock. We assume that ζ_j is observed by builders but not households. From the perspective of households, there are two components in the supply shock of type-j homes:

$$\zeta_j = \zeta + e_j, \ j \in \{H, M\}.$$

The first component ζ is common to the two types of homes. It has a normal distribution with $\bar{\zeta}$ as the mean and unit variance. The second component e_j is idiosyncratic to type-j homes. It has a normal distribution with zero mean and α as its standard deviation.

The builders' optimization determines their supply curves for both primary and secondary homes:

$$S_H = P^k e^{\xi_H} \text{ and } S_M = Q_1^k e^{\xi_M},$$

where, for $j \in \{H, M\}$, $\xi_j \sim N\left(\bar{\xi}, (1+\alpha^2) k^2\right)$ and $\bar{\xi} = k\bar{\zeta}$.

3.5.1.3 Immigrants

Immigrants decide whether they want to move into the neighborhood at $t = 1$, although they move into the neighborhood only at $t = 2$ and thus purchase secondary homes from the initial households at $t = 2$. This is realistic as it takes time for immigrants to move families from one area to another. As such, immigrants have to make their migration decisions based on their expectations of the strength of a neighborhood at $t = 1$, rather than the realized strength at $t = 2$. As immigrants are from outside of the neighborhood, it is reasonable to assume that they do not receive any private information and have to rely on $\mathcal{I}_c = \{P, Q_1\}$, which contains the publicly observable housing prices at $t = 1$, to form their expectations. Like the households, the immigrants also underestimate the housing supply elasticity by the same factor $\phi \leq 1$.

It is intuitive that when immigrants hold a higher expectation about the strength of a neighborhood, the neighborhood will attract a larger number of immigrants. This is because immigrants also enjoy living and working in a stronger neighborhood. Consequently, there will be a greater demand for secondary homes at $t = 2$. For simplicity, we adopt a reduced form to capture the immigrants' housing demand by assuming that the aggregate wealth W they bring to buy homes is proportional to their expected strength of the neighborhood:

$$W = E\left[e^A \middle| \mathcal{I}_c \right]. \tag{47}$$

This form allows us to maintain the tractable log-linear equilibrium.[71]

Anticipating the arrival of immigrants at $t = 2$, the initial households act as intermediaries by buying secondary homes at $t = 1$ and selling them to immigrants at $t = 2$. To the extent that the households can perfectly predict the future housing demand of the immigrants based on the public information available at $t = 1$, they do not bear any risk in speculating in secondary homes. This feature serves to simplify our analysis and to highlight the key insight that the households' demand for secondary homes is crucially influenced by housing prices P and Q_1 traded at $t = 1$ through their impact on the immigrants' expectation of the neighborhood strength.

3.5.1.4 Equilibrium

We derive the noisy rational expectations equilibrium as in the baseline model. The

[71] One may micro-found this form in different ways. One possibility is that the number of immigrants increases with their expectation of the strength of neighborhood and each immigrant arrives with a fixed amount of wealth to acquire housing. Another possibility, which we have explicitly worked out in an earlier draft, is to let the immigrants supply labor to the initial households based on their expectations of the strength of the neighborhood, which determines the productivity of the households.

equilibrium features the clearing of both primary and secondary homes at $t = 1$:

$$\begin{aligned}\int_{-\infty}^{\infty} H_i(\theta_i, P, Q_1)\, d\Phi(v_i) &= P^k e^{\xi_H},\\ \int_{-\infty}^{\infty} M_i(\theta_i, P, Q_1)\, d\Phi(v_i) &= Q_1^k e^{\xi_M},\end{aligned}$$

and the households' learning from the prices of both primary and secondary homes. We also impose clearing in the market for secondary homes at $t = 2$, which requires the immigrants to spend all their wealth W to purchase secondary homes:

$$Q_2 \int_{-\infty}^{\infty} M_i(\theta_i, P, Q_1)\, d\Phi(v_i) = W.$$

As the nature of the equilibrium and the key steps of deriving the equilibrium are similar to the baseline model, we leave the detailed description and derivation of the equilibrium to an Internet Appendix. Instead, we briefly summarize the key features of the equilibrium in the extended model here.

There is a unique log-linear equilibrium where the primary home price is a log-linear function of A, ξ_H, and $\log Q_1$:

$$\log P = p_A A + p_\xi \xi_H + p_Q \log Q_1 + p_0,$$

and the prices of secondary homes at $t = 1$ and $t = 2$ are identical and equal to a log-linear function of ξ_M and $\log P$:

$$\log Q_1 = \log Q_2 = q_\xi \xi_M + q_P \log P + q_0.$$

All coefficients are given in the Internet Appendix. As the immigrants' demand at $t = 2$ for secondary homes is determined by the public information available at $t = 1$, the households can fully anticipate the price of secondary homes Q_2 at $t = 2$. Competitive pressure ensures that they earn zero profit by buying secondary homes at $t = 1$ and then selling them at $t = 2$. As a result, the price of secondary homes Q_1 at $t = 1$ is equal to Q_2.

As a result of the separate supply shocks in the primary and secondary home markets, the prices of primary and secondary homes are not perfectly correlated. The price of primary homes P serves to aggregate the private information of households regarding the strength of the neighborhood A, while the price of secondary homes simply reflects P, together with another supply component $q_\xi \xi_M$, which in turn reveals the supply shock ξ_M. Each household, say household i, treats both prices P and Q_1 as useful signals, in addition to its private signal θ_i, in forming its expectation of A.

Like the baseline model, through this informational channel, informational frictions and households' overreaction to the price signals can jointly lead to a hump-shaped relationship between the log-price variance of both primary and secondary homes and the supply elasticity. To illustrate this relationship, we again use a numerical example based on the following parameter choices:

$$\tau_\theta = 0.1,\ \ \tau_A = 1, \phi = 0.1,\ \ \psi = 0.6, \eta_c = 0.5, \eta_H = 0.9, \alpha = 1, \tau_\varepsilon = 0.1. \tag{48}$$

The top two panels of Figure 7 depict the log-price variance of both primary and secondary homes against supply elasticity. It shows humped-shapes for both curves in the presence of informational frictions, consistent with that in Figure 6 for the baseline model.

The households' demand for secondary homes is ultimately driven by the immigrants' learning about the neighborhood strength through the housing prices. As a consequence, the learning effects are particularly important to the households' demand for secondary homes. Thus, in this extended model, the households' demand for secondary homes provides an additional dimension to examine learning effects. Specifically, the demand of household i for primary homes is a log-linear function of its private signal θ_i and housing prices $\log P$ and $\log Q_1$, while its demand for secondary homes is a log-linear function of $\log P$ and $\log Q_1$:

$$\begin{aligned} \log H_i &= h_\theta \theta_i + h_P \log P + h_Q \log Q_1 + h_0, \\ \log M_i &= \log M = m_P \log P + m_Q \log Q_1 + m_0, \end{aligned}$$

with all coefficients given in the Internet Appendix. As all households agree on the housing demand of immigrants at $t = 2$, they choose an identical demand schedule for secondary homes. We are particularly interested in the fraction of demand for secondary homes relative to the total housing demand $\frac{M_i}{H_i+M_i}$, as this ratio is directly measurable in the data.

3.5.1.5 Empirical Predictions

The bottom-left panel of Figure 7 depicts the variance of $\frac{M_i}{H_i+M_i}$, which measures the variability of the share of investment-driven housing demand in total housing demand with respect to supply elasticity in the presence and absence of informational frictions. In the absence of informational frictions, $Var\left[\frac{M_i}{H_i+M_i}\right]$ monotonically decreases with supply elasticity.

This pattern is intuitive and reflects how the cost effect of a higher housing price impacts primary home purchases more than secondary home purchases. As seen in the Internet Appendix, the loading of primary home demand $\log H$ (which is identical across all households) on the supply shock ζ_H is $\frac{k(1+\psi-\eta_H)}{1+k\eta_H}$, where $\eta_H \in (0,1)$ is the degree of housing-consumption complementarity, while for secondary home demand $\log M$ it is $\frac{k}{1+k}$. The primary home demand is, therefore, more variable than secondary home demand for a given supply shock, and this difference increases with the supply elasticity of the neighborhood. Consequently, in areas with more elastic supply, the fraction of secondary home purchases is less variable.

Interestingly, in the presence of informational frictions, Figure 7 shows a hump-shaped pattern of $Var\left[\frac{M_i}{H_i+M_i}\right]$ with respect to supply elasticity. This humped shape highlights the learning effects on the households' demand for secondary homes. Building on the same insight from our earlier discussion, the demand for secondary homes is most variable in an intermediate range of supply elasticity because the joint effects of informational frictions and the overreaction of the households and immigrants to the price signals are most influential in affecting their expectations of the neighborhood strength.

The non-monotonic relationship between the variability of the fraction of secondary-home demand relative to total demand and housing supply elasticity is again in sharp contrast to the monotonic relationship in the perfect-information benchmark. This non-monotonic relationship provides a new prediction for us to explore in the data.

Figure 7 also highlights a second salient feature in the presence of informational frictions. The variance of housing prices and the variance of the share of secondary-home demand exhibit similar hump-shaped patterns across supply elasticity, which suggests a positive correlation between them. The bottom-right panel of Figure 7, by displaying a scatter plot of these two variances across supply elasticities, illustrates this positive correlation, which provides a second new prediction for us to test in the data.

3.5.2 Empirical Evidence

In this subsection, we examine the two empirical predictions provided by the model extension using data from the recent U.S. housing boom: 1) whether during the boom period of 2004-2006 the share of non-owner-occupied home purchases in the total home purchases had the greatest increases relative to the pre-boom period of 2001-2003 in counties with intermediate supply elasticities, and 2) whether counties with greater increases in the share of non-owner-occupied home purchases during the boom also experienced larger price increases in 2004-2006 and larger price decreases during the bust period of 2007-2009.

We construct the share of non-owner-occupied home purchases at the county level from the "Home Mortgage Disclosure Act" (HMDA) dataset. The HMDA has comprehensive coverage for mortgage applications and originations in the U.S. We use mortgages originated for home purchases. HMDA data identify owner occupancy for each individual mortgage. We then aggregate the HMDA data to the county level and calculate the fraction of mortgage origination for non-owner-occupied homes in the total mortgage origination as our measure of the share of secondary home purchases.

Figure 8 depicts the share of non-owner-occupied home purchases for the U.S. and for three cities, New York, Las Vegas, and Charlotte. At a national level, the share of non-owner-occupied home purchases rose steadily from a modest level of 7% in 2000 to peak at a level above 15% in 2005. It then fell gradually to less than 10% in 2010. The peak of the share of non-owner-occupied home purchases in 2005 was slightly in advance of the peak of the national home price index in 2006, as shown in Figure 1. Nevertheless, the rise and fall of the share of non-owner-occupied home purchases were roughly in sync with the boom and bust of home prices.

Across the three cities, it is interesting to note that Las Vegas had the most dramatic rise and fall in the share of non-owner-occupied home purchases, followed by Charlotte, with New York having the most modest rise and fall. The most variable share of non-owner-occupied home purchases experienced by Las Vegas is particularly interesting as Las Vegas also had the most dramatic price cycle among these cities.

We now systematically examine the share of non-owner-occupied home purchases across counties with different housing supply elasticities. We focus on the change in the fraction of non-owner-occupied home purchases from the pre-boom period of 2001-2003 to the boom period of 2004-2006.

In Figure 9, we sort the counties in our sample into three groups in the top panel and ten groups in the bottom panel using the Saiz's supply elasticity measure, with each group having the same number of counties. The top panel shows that the change in the share of non-owner-occupied home purchases is almost the same between the inelastic and the middle groups, and smaller in the elastic group. As we discussed before, this coarse grouping might hide finer non-monotonicity. Indeed, the bottom panel shows that the change in the fraction of non-owner-occupied home purchases displays a non-monotonic pattern across ten elasticity groups with the largest share of non-owner-occupied home purchases in groups 3 and 4. This non-monotonic pattern is consistent with the first prediction of the extended model.

In Figure 10, we sort the counties into elasticity groups each with an equal population rather than number of counties. Across either the three groups shown in the top panel or the ten groups shown in the bottom panel, there is a non-monotonic pattern in the change of the share of non-owner-occupied home purchases across the elasticity groups, with the change peaking in the middle of the groups.

To further examine whether the largest change in the share of non-owner-occupied home purchases in the middle elasticity groups is robust to controlling for other fundamental factors, we also adopt a regression approach in Table 2. Similar to the regressions reported in Table 1, we regress the change in the share of non-owner-occupied home purchases from 2001-2003 to 2004-2006 on two dummy variables that indicate whether a county is in the middle elastic group or the elastic group, which are constructed in the top panel of Figure 10, together with a list of control variables. This regression implicitly uses the inelastic group as the benchmark for the middle and elastic groups. The control variables include the fraction of subprime households in the county in 2005, the population change, and annualized per capita income change. Columns 1 and 2 of Table 2 report the regressions without and with the controls. In either regression specification, we observe the middle group has a significantly larger change in the share of non-owner-occupied home purchases than the other groups. Furthermore, none of the control variables is significant except the annualized per capita income change in 2004-2006.

Taken together, Figures 9-10 and Table 2 confirm the first prediction of the extended model that there is a non-monotonic relationship between the variability of the share of non-owner-occupied home purchases and housing supply elasticity.

Figure 11 and Table 3 provide evidence for the second prediction. The regressions in Table 3 show that change in the share of non-owner occupied home purchases from 2001-2003 to 2004-2006 is positively correlated with the size of the housing price boom in 2004-2006, and negatively correlated with the size of the housing price bust in 2007-2009. These results are robust to the inclusion of the control variables that are included in the test of the first prediction. The two panels of Figure 11 graphically illustrate these correlations by providing scatter plots of the size of the housing price boom in 2004-2006 and the housing price bust in 2007-2009 against the change in the share of non-owner-occupied home purchases from 2001-2003 to 2004-2006. These plots show that these correlations are a broad feature of the data rather than driven by a few outlying counties. Our empirical analysis thus confirms the second prediction of the extended model that there is a positive correlation between the volatility of housing prices and the variability of the share of non-owner occupied home purchases during the recent U.S. housing cycle.

3.6 Conclusion

This paper highlights a non-monotonic relationship between the magnitude of housing cycles and housing supply elasticity in the cross-section of U.S. county data during the U.S. housing cycle of the 2000s. We develop a model of information aggregation and learning in housing markets to explain this phenomenon. In the presence of pervasive informational frictions regarding economic strength and housing supply of a neighborhood, households face a realistic problem in learning about these fundamental variables with housing prices serving as important signals. Our model highlights how the households' learning interacts with characteristics endemic to local housing supply and demand to impact housing price dynamics. In particular, supply elasticity affects not only housing supply but also the informational noise in the price signal for the households' learning of the neighborhood strength. Our model predicts that housing price and share of investment home purchases are both most variable in areas with intermediate supply elasticities and that these variances are positively correlated. Our empirical analysis also provides evidence that supports these predictions.

3.A: Proofs of Propositions

Proof of Proposition 3.1:

The first order conditions for household i's choices of H_i and l_i at an interior point are

$$H_i \quad : \quad \frac{(1-\eta_c)(1-\eta_H)}{H_i} E\left[U\left(\{H_j\}_{j\in[0,1]}, C_i \right) \middle| \mathcal{I}_i \right] = PE\left[\frac{\eta_H}{C_i} U\left(\{H_j\}_{j\in[0,1]}, C_i \right) \middle| \mathcal{I}_i \right] \tag{49}$$

$$l_i \quad : \quad l_i^{\psi} = E\left[\frac{\eta_H}{C_i} U\left(\{H_j\}_{j\in[0,1]}, C_i \right) e^{A_i} \middle| \mathcal{I}_i \right]. \tag{50}$$

Imposing $C_i = e^{A_i} l_i$ in equation (49), one arrives at

$$PH_i = \frac{(1-\eta_c)(1-\eta_H)}{\eta_H} \frac{E\left[\left(\int_{[0,1]/i} H_j dj \right)^{\eta_c(1-\eta_H)} e^{\eta_H A_i} \middle| \mathcal{I}_i \right]}{E\left[\left(\int_{[0,1]/i} H_j dj \right)^{\eta_c(1-\eta_H)} e^{(\eta_H - 1) A_i} \middle| \mathcal{I}_i \right]} l_i.$$

From equation (50), it follows that

$$l_i = \left\{ \eta_H E\left[\left\{ \frac{1}{1-\eta_H} \left(\frac{H_i}{1-\eta_c} \right)^{1-\eta_c} \left(\frac{\int_{[0,1]/i} H_j dj}{\eta_c} \right)^{\eta_c} \right\}^{1-\eta_H} \left(\frac{1}{\eta_H} e^{A_i} \right)^{\eta_H} \middle| \mathcal{I}_i \right] \right\}^{\frac{1}{1+\psi-\eta_H}},$$

from which we see that

$$\begin{aligned} \log H_i \quad = \quad & \frac{2+\psi-\eta_H}{\psi + (1-\eta_H)\eta_c} \log E\left[\left(\int_{[0,1]/i} H_j dj \right)^{\eta_c(1-\eta_H)} e^{\eta_H A_i} \middle| \mathcal{I}_i \right] - \frac{1+\psi-\eta_H}{\psi + (1-\eta_H)\eta_c} \log P \\ & - \frac{1+\psi-\eta_H}{\psi + (1-\eta_H)\eta_c} \log E\left[\left(\int_{[0,1]/i} H_j dj \right)^{\eta_c(1-\eta_H)} e^{(\eta_H - 1) A_i} \middle| \mathcal{I}_i \right] \\ & + \frac{\psi}{\psi + (1-\eta_H)\eta_c} \log \left((1-\eta_c) \left(\frac{1-\eta_H}{\eta_H} \right) \left(\frac{1-\eta_c}{\eta_c} \right)^{\frac{\eta_c(1-\eta_H)}{\psi}} \right). \end{aligned}$$

Note that integrating over the continuum of other households' housing choices does not change when sets of measure zero are substracted from it. We then obtain equation (37).

Proof of Proposition 3.2:

We first conjecture that each household's housing purchasing and the housing price take the following log-linear forms:

$$\log H_i = h_P \log P + h_\theta \theta_i + h_0, \tag{51}$$

$$\log P = p_A A + p_\xi \xi + p_0, \tag{52}$$

where the coefficients h_0, h_P, h_θ, p_0, p_A, and p_ξ will be determined by equilibrium conditions.

Given the conjectured functional form for H_i, we can expand equation (37). It follows that

$$\begin{aligned} & E\left[\left(\int_{[0,1]} H_j dj\right)^{\eta_c(1-\eta_H)} e^{\eta_H A_i}\middle| \mathcal{I}_i\right] \\ = \quad & e^{\eta_c(1-\eta_H)\left(h_0+h_P\log P+\frac{1}{2}h_\theta^2\tau_\theta^{-1}\right)+\frac{1}{2}\eta_H^2\tau_\varepsilon^{-1}} E\left[e^{(\eta_H+\eta_c(1-\eta_H)h_\theta)A}\middle| \mathcal{I}_i\right], \end{aligned}$$

where we use the fact that A is independent of ε_j and exploit the Law of Large Number for the continuum when integrating over households, which still holds if we subtract sets of measure 0 from the integral. A similar expression obtains for $E\left[\left(\int_{[0,1]} H_j dj\right)^{\eta_c(1-\eta_H)} e^{(\eta_H-1)A_i}\middle| \mathcal{I}_i\right]$

Define

$$q \equiv \frac{\log P - p_0 - p_\xi\bar{\xi}}{p_A} = A + \frac{p_\xi}{p_A}\left(\xi - \bar{\xi}\right),$$

which is a sufficient statistic of information contained in P. Then, conditional on observing its own signal θ_i and the housing price P, household i's expectation of A is

$$E\left[A \mid \theta_i, \log P\right] = E\left[A \mid \theta_i, q\right] = \frac{1}{\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2}\frac{1}{(\phi k)^2}}\left(\tau_A\bar{A} + \tau_\theta\theta_i + \frac{p_A^2}{p_\xi^2}\frac{1}{(\phi k)^2}q\right),$$

and its conditional variance of A is

$$Var\left[A \mid \theta_i, \log P\right] = \left(\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2}\frac{1}{(\phi k)^2}\right)^{-1}.$$

Therefore,

$$\begin{aligned} & \log E\left[e^{(\eta_H+\eta_c(1-\eta_H)h_\theta)A}\middle| \mathcal{I}_i\right] \\ = \quad & \left(\eta_H + \eta_c\left(1-\eta_H\right)h_\theta\right)\left(\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2}\frac{1}{(\phi k)^2}\right)^{-1}\left(\tau_A\bar{A} + \tau_\theta\theta_i + \frac{p_A^2}{p_\xi^2}\frac{1}{(\phi k)^2}q\right) \\ & + \frac{1}{2}\left(\eta_H + \eta_c\left(1-\eta_H\right)h_\theta\right)^2\left(\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2}\frac{1}{(\phi k)^2}\right)^{-1}. \end{aligned}$$

Then,

$$
\begin{aligned}
& \log E\left[\left.\left(\int_{[0,1]} H_j dj\right)^{\eta_c(1-\eta_H)} e^{\eta_H A_i}\right| \mathcal{I}_i\right] \\
= \; & (\eta_H + \eta_c (1-\eta_H) h_\theta) \left(\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2} \frac{1}{(\phi k)^2}\right)^{-1} \\
& \cdot \left(\tau_A \bar{A} + \tau_\theta \theta_i + \frac{p_A}{p_\xi^2} \frac{1}{(\phi k)^2} \left(\log P - p_0 - p_\xi \bar{\xi}\right)\right) \\
& + \eta_c (1-\eta_H) \left(h_0 + h_P \log P + \frac{1}{2} h_\theta^2 \tau_\theta^{-1}\right) + \frac{1}{2} \eta_H^2 \tau_\varepsilon^{-1} \\
& + \frac{1}{2} (\eta_H + \eta_c (1-\eta_H) h_\theta)^2 \left(\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2} \frac{1}{(\phi k)^2}\right)^{-1}.
\end{aligned}
$$

Substituting this expression into equation (37) and matching coefficients with the conjectured log-linear form in (51), it follows that

$$
h_\theta = \frac{1+\psi+\eta_c(1-\eta_H) h_\theta}{\psi + (1-\eta_H)\eta_c} \left(\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2} \frac{1}{(\phi k)^2}\right)^{-1} \tau_\theta, \tag{53}
$$

$$
h_P = \frac{1+\psi+\eta_c(1-\eta_H) h_\theta}{\psi} \left(\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2} \frac{1}{(\phi k)^2}\right)^{-1} \frac{p_A}{p_\xi^2} \frac{1}{(\phi k)^2} - \frac{1+\psi-\eta_H}{\psi}, \tag{54}
$$

$$
\begin{aligned}
h_0 = \; & \frac{1+\psi+\eta_c(1-\eta_H) h_\theta}{\psi + (1-\eta_H)\eta_c} \left(\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2} \frac{1}{(\phi k)^2}\right)^{-1} \left(\tau_A \bar{A} - \frac{p_A}{p_\xi^2} \frac{1}{(\phi k)^2} \left(p_0 + p_\xi \bar{\xi}\right)\right) \\
& + \frac{\eta_c(1-\eta_H)}{\psi + (1-\eta_H)\eta_c} \left(h_0 + \frac{1}{2} h_\theta^2 \tau_\theta^{-1}\right) + \frac{1}{2} \frac{(1+\psi)\eta_H - (1+\psi-\eta_H)(1-\eta_H)}{\psi + (1-\eta_H)\eta_c} \tau_\varepsilon^{-1} \\
& + \frac{1}{2} \frac{1+\psi+\eta_c(1-\eta_H) h_\theta}{\psi + (1-\eta_H)\eta_c} (\eta_H + \eta_c(1-\eta_H) h_\theta) \left(\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2} \frac{1}{(\phi k)^2}\right)^{-1} \\
& + \frac{1}{2} \frac{1+\psi-\eta_H}{\psi + (1-\eta_H)\eta_c} (\eta_H - 1 + \eta_c(1-\eta_H) h_\theta) \left(\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2} \frac{1}{(\phi k)^2}\right)^{-1} \\
& + \frac{\psi}{\psi + (1-\eta_H)\eta_c} \log\left((1-\eta_c) \left(\frac{1-\eta_H}{\eta_H}\right) \left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}}\right).
\end{aligned} \tag{55}
$$

By aggregating households' housing demand and the builders' supply and imposing market clearing in the housing market, we have

$$
h_0 + h_P (p_0 + p_A A + p_\xi \xi) + h_\theta A + \frac{1}{2} h_\theta^2 \tau_\theta^{-1} = \xi + k (p_0 + p_A A + p_\xi \xi).
$$

Matching coefficients of the two sides of the equation leads to the following three conditions:

$$h_0 + h_P p_0 + \frac{1}{2} h_\theta^2 \tau_\theta^{-1} = k p_0, \tag{56}$$

$$h_P p_A + h_\theta = k p_A, \tag{57}$$

$$h_P p_\xi = 1 + k p_\xi. \tag{58}$$

It follows from equation (58) that

$$p_\xi = -\frac{1}{k - h_P}, \tag{59}$$

and further from equation (57) that

$$p_A = \frac{h_\theta}{k - h_P}. \tag{60}$$

Thus, by taking the ratio of equations (60) and (59), we arrive at

$$\frac{p_A}{p_\xi} = -h_\theta.$$

Substituting $\frac{p_A}{p_\xi} = -h_\theta$ into equation (53), and defining $b = -\frac{p_A}{p_\xi}$, we arrive at

$$\frac{1}{(\phi k)^2} b^3 + \left(\tau_A + \frac{\psi}{\psi + (1 - \eta_H)\, \eta_c} \tau_\theta \right) b - \frac{1 + \psi}{\psi + (1 - \eta_H)\, \eta_c} \tau_\theta = 0. \tag{61}$$

We see from equation (A.19) that b has at most one positive root since the above third order polynomial has only one sign change, by Descartes' Rule of Signs. By setting $b \to -b$, we see that there is no sign change, and therefore b has no negative root. Furthermore, by the Fundamental Theorem of Algebra, the roots of the polynomial (A.19) exist. Thus, it follows that equation (A.19) has only one real, nonnegative root $b \geq 0$ and 2 complex roots.[72]

Furthermore, by dropping the cubic term from equation (A.19), one arrives at an upper bound for b :

$$b \leq \frac{1 + \psi}{\psi} \frac{\tau_\theta}{\left(1 + \frac{(1 - \eta_H)\eta_c}{\psi} \right) \tau_A + \tau_\theta}.$$

Since $h_\theta = -\frac{p_A}{p_\xi} = b$, we can recover $h_\theta = b > 0$ and $p_\xi = -\frac{1}{b} p_A < 0$. From equation (54) and $b = -\frac{p_A}{p_\xi}$, it follows that

$$h_P = -\frac{1 + \psi - \eta_H}{\psi} + \frac{1 + \psi + \eta_c\,(1 - \eta_H)\, b}{\psi} \frac{\left(\frac{b}{\phi k} \right)^2}{\tau_A + \tau_\theta + \left(\frac{b}{\phi k} \right)^2} \frac{1}{p_A}. \tag{62}$$

[72] The uniqueness of the positive, real root also follows from the fact that the LHS of the polynomial equation monotonically increases in b.

From equation (58), one also has that $h_P = k + p_\xi^{-1}$. Since $p_\xi \leq 0$, it follows that $h_P < k$ whenever $k > 0$.

From $h_\theta = b$ and equations (60) and (62), we arrive at

$$p_A = \frac{\psi}{1+\psi(1+k)-\eta_H}\left(b + \frac{1+\psi+\eta_c(1-\eta_H)b}{\psi}\frac{\left(\frac{b}{\phi k}\right)^2}{\tau_A+\tau_\theta+\left(\frac{b}{\phi k}\right)^2}\right) > 0. \tag{63}$$

One arrives at p_ξ from recognizing that $p_\xi = -\frac{1}{b}p_A$. Manipulating equation (61), we recognize that

$$\frac{1+\psi+\eta_c(1-\eta_H)b}{\psi} = \frac{\psi+(1-\eta_H)\eta_c}{\psi}\left(\tau_A+\tau_\theta+\left(\frac{b}{\phi k}\right)^2\right)b\tau_\theta^{-1}. \tag{64}$$

Substituting equation (64) into equation (63), and invoking equation (61) to replace $\frac{1}{(\phi k)^2}b^3$, one arrives at

$$p_A = \frac{1+\psi}{1+\psi(1+k)-\eta_H} - \frac{\psi+(1-\eta_H)\eta_c}{1+\psi(1+k)-\eta_H}\tau_\theta^{-1}\tau_A b. \tag{65}$$

and from equation (63), equation (64). and $p_\xi = -\frac{1}{b}p_A$, one also has that

$$p_\xi = -\frac{\psi}{1+\psi(1+k)-\eta_H} - \frac{\psi+(1-\eta_H)\eta_c}{1+\psi(1+k)-\eta_H}\tau_\theta^{-1}\left(\frac{b}{\phi k}\right)^2 < 0.$$

From $h_\theta = b$, $b = -\frac{p_A}{p_\xi}$, and equations (62), (56) and (55), one also finds that

$$\begin{aligned} p_0 \quad = \quad & \frac{1+\psi+\eta_c(1-\eta_H)b}{1+\psi(1+k)-\eta_H}\frac{\tau_A\bar{A}+\frac{1}{(\phi k)^2}b\bar{\xi}}{\tau_A+\tau_\theta+\left(\frac{b}{\phi k}\right)^2} + \frac{1}{2}\frac{\psi+\eta_c(1-\eta_H)}{1+\psi(1+k)-\eta_H}b^2\tau_\theta^{-1} \\ & + \frac{1}{2}\frac{(1+\psi)\eta_H-(1-\eta_H)(1+\psi-\eta_H)}{1+\psi(1+k)-\eta_H}\tau_\varepsilon^{-1} \\ & + \frac{1}{2}\frac{1+\psi+\eta_c(1-\eta_H)b}{1+\psi(1+k)-\eta_H}(\eta_H+\eta_c(1-\eta_H)b)\left(\tau_A+\tau_\theta+\left(\frac{b}{\phi k}\right)^2\right)^{-1} \\ & + \frac{1}{2}\frac{1+\psi-\eta_H}{1+\psi(1+k)-\eta_H}(\eta_H-1+\eta_c(1-\eta_H)b)\left(\tau_A+\tau_\theta+\left(\frac{b}{\phi k}\right)^2\right)^{-1} \\ & + \frac{\psi}{1+\psi(1+k)-\eta_H}\log\left((1-\eta_c)\left(\frac{1-\eta_H}{\eta_H}\right)\left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}}\right). \end{aligned} \tag{66}$$

Given p_0, p_A, and $b = -\frac{p_A}{p_\xi}$, we can recover h_0 from equation (55).

Since we have explicit expressions for all other equilibrium objects as functions of b, and b exists and is unique, it follows that an equilibrium in the economy exists and is unique.

Proof of Proposition 3.3:

When all households observe A directly, there are no longer information frictions in the economy. Since the households' idiosyncratic productivity components are unobservable, they are now symmetric. Then, it follows that $H_j = H_i = H$. Imposing this symmetry in equation (37), we see that each household's housing demand is then given by

$$\begin{aligned} \log H &= \frac{1+\psi}{\psi} A - \frac{1+\psi-\eta_H}{\psi} \log P + \frac{(1+\psi)\,\eta_H - (1-\eta_H)(1+\psi-\eta_H)}{2\psi} \tau_\varepsilon^{-1} \\ &\quad + \log\left((1-\eta_c) \left(\frac{1-\eta_H}{\eta_H}\right) \left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}} \right). \end{aligned}$$

By market clearing, $\log H = \xi + k \log P$, it follows that

$$\begin{aligned} \log P &= \frac{1+\psi}{1+\psi(1+k)-\eta_H} A - \frac{\psi}{1+\psi(1+k)-\eta_H} \xi + \frac{(1+\psi)\,\eta_H - (1-\eta_H)(1+\psi-\eta_H)}{2\,(1+\psi(1+k)-\eta_H)} \tau_\varepsilon^{-1} \\ &\quad + \frac{\psi}{1+\psi(1+k)-\eta_H} \log\left((1-\eta_c) \left(\frac{1-\eta_H}{\eta_H}\right) \left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}} \right). \end{aligned}$$

This characterizes the economy in the limit as information frictions dissipate.

To see that the economy with information frictions (finite τ_θ) converges to this perfect-information limit, we consider a sequence of τ_θ that converges to ∞. From equation (A.19), it follows that, as $\tau_\theta \nearrow \infty$, $b \to \frac{1+\psi}{\psi}$. Since $h_\theta = b$, it follows that

$$h_\theta \to \frac{1+\psi}{\psi}.$$

Taking the limit $\tau_\theta \nearrow \infty$ in equation (63), recognizing that $h_\theta = b$ remains finite in the limit, we see that

$$p_A \to \frac{1+\psi}{1+\psi(1+k)-\eta_H}.$$

Since $p_\xi = -\frac{1}{b} p_A$, it follows that

$$p_\xi \to -\frac{\psi}{1+\psi(1+k)-\eta_H}.$$

In addition, from equation (62), we find that as $\tau_\theta \nearrow \infty$,

$$h_P \to -\frac{1+\psi-\eta_H}{\psi}$$

Finally, from equations (66) and (55), it follows that

$$
\begin{aligned}
p_0 \quad \to \quad & \frac{\psi}{1+\psi\left(1+k\right)-\eta_H}\log\left(\left(1-\eta_c\right)\left(\frac{1-\eta_H}{\eta_H}\right)\left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}}\right) \\
& +\frac{1}{2}\frac{\left(1+\psi\right)\eta_H-\left(1-\eta_H\right)\left(1+\psi-\eta_H\right)}{1+\psi\left(1+k\right)-\eta_H}\tau_\varepsilon^{-1}, \\
h_0 \quad \to \quad & \log\left(\left(1-\eta_c\right)\left(\frac{1-\eta_H}{\eta_H}\right)\left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}}\right)+\frac{\left(1+\psi\right)\eta_H-\left(1-\eta_H\right)\left(1+\psi-\eta_H\right)}{2\psi}\tau_\varepsilon^{-1}.
\end{aligned}
$$

Thus, we see that the economy with information frictions converges to the perfect-information benchmark as $\tau_\theta \nearrow \infty$.

Furthermore, the variance of the log housing price is given by

$$
Var\left[\log P\right]=\left(\frac{\psi k}{1+\psi\left(1+k\right)-\eta_H}\right)^2\left(1+\left(\frac{1+\psi}{\psi k}\right)^2\tau_A^{-1}\right),
$$

from which follows that

$$
\frac{\partial Var\left[\log P\right]}{\partial k}=\frac{2\psi^2 k}{\left(1+\psi\left(1+k\right)-\eta_H\right)^3}\left\{1+\psi-\eta_H-\frac{\left(1+\psi\right)^2}{\psi k}\tau_A^{-1}\right\}.
$$

For $k<\frac{1+\psi}{\psi}\frac{1+\psi}{1+\psi-\eta_H}\tau_A^{-1}$, $\frac{\partial Var[\log P]}{\partial k}<0$. For $k>\frac{1+\psi}{\psi}\frac{1+\psi}{1+\psi-\eta_H}\tau_A^{-1}$, $\frac{\partial Var[\log P]}{\partial k}>0$. Thus it follows that the log housing price is U-shaped in k.

Proof of Proposition 3.4:

From equation (65), it is clear that

$$
p_A=\frac{1+\psi}{1+\psi\left(1+k\right)-\eta_H}-\frac{\psi+\left(1-\eta_H\right)\eta_c}{1+\psi\left(1+k\right)-\eta_H}\tau_\theta^{-1}\tau_A b<\frac{1+\psi}{1+\psi\left(1+k\right)-\eta_H}.
$$

Thus, it follows that p_A is always lower than its corresponding value in the perfect-information benchmark.

Similarly, since $p_\xi=-\frac{1}{b}p_A$, it follows from equation (63) that we can express p_ξ as

$$
p_\xi=-\frac{\psi}{1+\psi\left(1+k\right)-\eta_H}-\frac{\psi+\left(1-\eta_H\right)\eta_c}{1+\psi\left(1+k\right)-\eta_H}\tau_\theta^{-1}\left(\frac{b}{\phi k}\right)^2<-\frac{\psi}{1+\psi\left(1+k\right)-\eta_H},
$$

which is the corresponding value of p_ξ in the perfect-information benchmark.

Proof of Proposition 3.5:

Note that b is determined by the polynomial equation (61). We define the LHS of the equation as $G(b)$. By using the Implicit Function Theorem and invoking equation (61), we have

$$\frac{\partial b}{\partial \eta_c} = -\frac{\partial G/\partial \eta_c}{\partial G/\partial b} = -\frac{(1-\eta_H)\,\tau_\theta}{3\frac{1}{(\phi k)^2}b^2 + \tau_A + \frac{\psi}{\psi+(1-\eta_H)\eta_c}\tau_\theta}\frac{1+\psi-\psi b}{(\psi+(1-\eta_H)\,\eta_c)^2}.$$

Since, from Proposition 3.2, $0 \le b \le \frac{1+\psi}{\psi}\frac{\tau_\theta}{\left(1+\frac{(1-\eta_H)\eta_c}{\psi}\right)\tau_A+\tau_\theta} \le \frac{1+\psi}{\psi}$, it follows that

$$\frac{1+\psi}{\psi} - b \ge 0$$

Thus $\frac{\partial b}{\partial \eta_c} < 0$. Similarly,

$$\frac{\partial b}{\partial \tau_\theta} = -\frac{\partial G/\partial \tau_\theta}{\partial G/\partial b} = \frac{1}{3\frac{1}{(\phi k)^2}b^2 + \tau_A + \frac{\psi}{\psi+(1-\eta_H)\eta_c}\tau_\theta}\frac{1+\psi-\psi b}{\psi+(1-\eta_H)\,\eta_c} > 0.$$

From the expression for p_A in Proposition 2,

$$\frac{\partial p_A}{\partial \eta_c} = -\frac{1-\eta_H}{1+\psi\,(1+k)-\eta_H}\tau_\theta^{-1}\tau_A b - \frac{\psi+(1-\eta_H)\,\eta_c}{1+\psi\,(1+k)-\eta_H}\tau_\theta^{-1}\tau_A\frac{\partial b}{\partial \eta_c}.$$

Then, it follows, subtituting with equation (61), that

$$\frac{\partial p_A}{\partial \eta_c} = -\frac{(1-\eta_H)\,\tau_\theta^{-1}\tau_A b}{1+\psi\,(1+k)-\eta_H}\frac{2\frac{1}{(\phi k)^2}b^3 + \frac{\psi b}{\psi+(1-\eta_H)\eta_c}\tau_\theta}{2\frac{1}{(\phi k)^2}b^3 + \frac{1+\psi}{\psi+(1-\eta_H)\eta_c}\tau_\theta} < 0.$$

Similarly, with respect to τ_θ, we have

$$\begin{aligned}\frac{\partial p_A}{\partial \tau_\theta} &= \frac{\psi+(1-\eta_H)\,\eta_c}{1+\psi\,(1+k)-\eta_H}\tau_\theta^{-2}\tau_A b\left(1-\tau_\theta\frac{1}{b}\frac{\partial b}{\partial \tau_\theta}\right)\\ &= \frac{\psi+(1-\eta_H)\,\eta_c}{1+\psi\,(1+k)-\eta_H}\tau_\theta^{-2}\tau_A b\frac{2\frac{1}{(\phi k)^2}b^3 + \frac{\psi b}{\psi+(1-\eta_H)\eta_c}\tau_\theta}{2\frac{1}{(\phi k)^2}b^3 + \frac{1+\psi}{\psi+(1-\eta_H)\eta_c}\tau_\theta} > 0.\end{aligned}$$

Proof of Proposition 3.6:

We first consider the limiting case for the economy as $k \to \infty$. Rewrite equation (61) as

$$\left(\frac{b}{\phi k}\right)^3 + \left(\tau_A + \frac{\psi}{\psi+(1-\eta_H)\,\eta_c}\tau_\theta\right)\frac{b}{\phi k} - \frac{1+\psi}{\psi+(1-\eta_H)\,\eta_c}\frac{1}{\phi k}\tau_\theta = 0. \tag{67}$$

Then it is apparent from equation (67) that, as $k \to \infty$, that either $\frac{b}{\phi k} = 0$ or $\frac{b}{\phi k} = \pm i\sqrt{\tau_A + \frac{\psi}{\psi+(1-\eta_H)\eta_c}\tau_\theta}$. Thus, as $k \to \infty$, one has that $\frac{b}{\phi k} \to 0$, and therefore $\frac{b}{k} \to 0$.

Consequently, from equation (65), $p_A \to 0$ and the housing price is completely driven by the supply shock ξ. From Proposition 2, one has that

$$p_\xi k = -\frac{\psi k}{1+\psi(1+k)-\eta_H} - \frac{\psi+(1-\eta_H)\eta_c}{1+\psi(1+k)-\eta_H}\tau_\theta^{-1}\frac{1}{\phi}\left(\frac{b}{\phi k}\right)b \to -1,$$

since b is bounded from above by $\frac{1+\psi}{\psi}$. Thus, $\log P = -\zeta$.

In addition, from equation (54), then, since $\frac{b}{k} \to 0$ and b is bounded from above by $\frac{1+\psi}{\psi}$, and from below by 0, one has that

$$h_P = -\frac{1+\psi+\eta_c(1-\eta_H)b}{\phi\psi}\left(\tau_A+\tau_\theta+\left(\frac{b}{\phi k}\right)^2\right)^{-1}\frac{b}{\phi k}\frac{1}{p_\xi k} - \frac{1+\psi-\eta_H}{\psi} \to -\frac{1+\psi-\eta_H}{\psi}.$$

From equation (53), it is straightforward to see that, as $k \to \infty$,

$$h_\theta = b \to \frac{1+\psi}{\psi}\frac{\tau_\theta}{\tau_A+\tau_\theta+\frac{(1-\eta_H)\eta_c}{\psi}\tau_A}.$$

Since h_θ remains bounded in the limit, it is easy to see from equation (66) that $p_0 \to 0$ as $k \to \infty$. It further follows from equation (55) that in the limit

$$\begin{aligned}
h_0 \;=\;& \frac{1+\psi}{\psi}\left(1+\frac{\eta_c(1-\eta_H)\tau_\theta}{\psi(\tau_A+\tau_\theta)+(1-\eta_H)\eta_c\tau_A}\right)\frac{\tau_A}{\tau_A+\tau_\theta}\bar{A} \\
&+\frac{\eta_c(1-\eta_H)}{2\psi}\left(\frac{\frac{1+\psi}{\psi}}{\tau_A+\tau_\theta+\frac{(1-\eta_H)\eta_c}{\psi}\tau_A}\right)^2\tau_\theta + \frac{1}{2}\frac{(1+\psi)\eta_H-(1+\psi-\eta_H)(1-\eta_H)}{\psi}\tau_\varepsilon^{-1} \\
&+\frac{1+\psi}{2\psi}\left(1+\frac{\eta_c(1-\eta_H)\tau_\theta}{\psi(\tau_A+\tau_\theta)+(1-\eta_H)\eta_c\tau_A}\right)\left(\eta_H+\frac{(1+\psi)\eta_c(1-\eta_H)\tau_\theta}{\psi(\tau_A+\tau_\theta)+(1-\eta_H)\eta_c\tau_A}\right)(\tau_A+\tau_\theta)^{-1} \\
&+\frac{1+\psi-\eta_H}{2\psi}\left(\eta_H-1+\frac{1+\psi}{\psi}\frac{\eta_c(1-\eta_H)\tau_\theta}{\tau_A+\tau_\theta+\frac{(1-\eta_H)\eta_c}{\psi}\tau_A}\right)(\tau_A+\tau_\theta)^{-1} \\
&+\log\left((1-\eta_c)\left(\frac{1-\eta_H}{\eta_H}\right)\left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}}\right).
\end{aligned} \tag{68}$$

In the case $k \to 0$, it follows from equation (67) that $b \to 0$ and $\frac{b}{k} \to \infty$. From equation (53), it follows that as $k \to 0$ one has that $h_\theta = b \to 0$. Furthermore, from equation (65), one has that

$$p_A \to \frac{1+\psi}{1+\psi-\eta_H}.$$

Since $h_\theta \to 0$, and p_A remain bounded as $k \to 0$, we also see from equation (57), subsituting for the limiting p_A, that $h_P \to 0$. Substituting for p_A in $p_\xi = -\frac{1}{b}p_A$ with equation (63), it follows that

$$p_\xi k = -\frac{\psi k}{1+\psi(1+k)-\eta_H} - \frac{1+\psi+\eta_c(1-\eta_H)b}{1+\psi(1+k)-\eta_H}\frac{\frac{1}{\phi^2}\frac{b}{k}}{\tau_A+\tau_\theta+\left(\frac{b}{\phi k}\right)^2} \to 0,$$

and the demand shock A completely drives the housing price.

Since h_θ remains bounded in the limit, it is easy to see from equation (66) that as $k \to 0$,

$$\begin{aligned} p_0 &= \frac{1}{2}\frac{(1+\psi)\eta_H - (1-\eta_H)(1+\psi-\eta_H)}{1+\psi-\eta_H}\tau_\varepsilon^{-1} \\ &\quad + \frac{\psi}{1+\psi-\eta_H}\log\left((1-\eta_c)\left(\frac{1-\eta_H}{\eta_H}\right)\left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}}\right). \end{aligned} \tag{69}$$

It further follows from equation (56) that in the limit $h_0 \to 0$.

Proof of Proposition 3.7:

We first prove that p_A decreases with ϕ and $p_\xi < 0$ increases with ϕ. Note that b is determined by the polynomial equation (A.19). We define the LHS of the equation as $G(b)$. Comparative statics of b with respect to ϕ reveal, by the Implicit Function Theorem and invoking equation (A.19), that

$$\begin{aligned} \frac{\partial b}{\partial \phi} &= -\frac{\partial G/\partial \phi}{\partial G/\partial b} = \frac{2\frac{1}{(\phi k)^2}b^3}{3\frac{1}{(\phi k)^2}b^2 + \tau_A + \frac{\psi}{\psi+(1-\eta_H)\eta_c}\tau_\theta}\frac{1}{\phi} \\ &= \frac{2\frac{1}{(\phi k)^2}b^4}{2\frac{1}{(\phi k)^2}b^3 + \frac{1+\psi}{\psi+(1-\eta_H)\eta_c}\tau_\theta}\frac{1}{\phi} > 0. \end{aligned}$$

From the expression for p_A in Proposition 3.2,

$$\frac{\partial p_A}{\partial \phi} = -\frac{\psi+(1-\eta_H)\eta_c}{1+\psi(1+k)-\eta_H}\tau_\theta^{-1}\tau_A\frac{\partial b}{\partial \phi} < 0.$$

Furthermore, by the Implicit Function Theorem, it follows that

$$\begin{aligned} \frac{\partial p_\xi}{\partial \phi} &= -\frac{2}{\phi}\frac{\psi+(1-\eta_H)\eta_c}{1+\psi(1+k)-\eta_H}\tau_\theta^{-1}\left(\frac{b}{\phi k}\right)^2\left(\frac{\phi}{b}\frac{\partial b}{\partial \phi} - 1\right) \\ &= \frac{2}{\phi}\frac{\psi+(1-\eta_H)\eta_c}{1+\psi(1+k)-\eta_H}\tau_\theta^{-1}\left(\frac{b}{\phi k}\right)^2\frac{\frac{1+\psi}{\psi+(1-\eta_H)\eta_c}\tau_\theta}{2\frac{1}{(\phi k)^2}b^3 + \frac{1+\psi}{\psi+(1-\eta_H)\eta_c}\tau_\theta}. \end{aligned}$$

Since $\phi \in [0,1]$, it follows that $\frac{\partial p_\xi}{\partial \phi} > 0$.

The variance of the housing price $Var[\log P]$ is given by

$$Var[\log P] = p_A^2\tau_A^{-1} + p_\xi^2 k^2,$$

from which follows that

$$\frac{\partial Var[\log P]}{\partial \phi} = 2p_A\tau_A^{-1}\frac{\partial p_A}{\partial \phi} + 2p_\xi k^2\frac{\partial p_\xi}{\partial \phi} < 0,$$

since $p_A \frac{\partial p_A}{\partial \phi} < 0$ and $p_\xi \frac{\partial p_\xi}{\partial \phi} < 0$.

From Proposition 3.3, the variance of the housing price in the perfect-information benchmark is

$$Var\left[\log P^{perf}\right] = \left(\frac{\psi k}{1+\psi(1+k)-\eta_H}\right)^2 \left(1+\left(\frac{1+\psi}{\psi k}\right)^2 \tau_A^{-1}\right).$$

It then follows, substituting for p_A and p_ξ with Proposition 3.2, that

$$\begin{aligned}
& Var\left[\log P\right] - Var\left[\log P^{perf}\right] \\
= & \left(p_A^2 - \left(\frac{1+\psi}{1+\psi(1+k)-\eta_H}\right)^2\right)\tau_A^{-1} + \left(p_\xi^2 - \left(\frac{\psi}{1+\psi(1+k)-\eta_H}\right)^2\right) k^2 \\
= & \left(\frac{\psi+(1-\eta_H)\eta_c}{1+\psi(1+k)-\eta_H}\right)^2 \tau_\theta^{-1} b\left(\tau_\theta^{-1}\tau_A b - \frac{2(1+\psi)}{\psi+(1-\eta_H)\eta_c}\right) \\
& + \left(\frac{\psi+(1-\eta_H)\eta_c}{1+\psi(1+k)-\eta_H}\right)^2 \tau_\theta^{-1}\left(\frac{b}{\phi}\right)^2\left(2\frac{\psi}{\psi+(1-\eta_H)\eta_c} + \tau_\theta^{-1}\left(\frac{b}{\phi k}\right)^2\right),
\end{aligned}$$

from which follows, substituting with equation (61), that $Var\left[\log P\right] - Var\left[\log P^{perf}\right] \geq 0$ whenever

$$b \geq \left(\left(\phi^2-1\right)\tau_\theta^{-1}\tau_A + \frac{\psi}{\psi+(1-\eta_H)\eta_c}\right)^{-1}\left(2\phi^2-1\right)\frac{1+\psi}{\psi+(1-\eta_H)\eta_c}. \tag{70}$$

Since $b \geq 0$, it is thus sufficient for $1 - \frac{\psi}{\psi+(1-\eta_H)\eta_c}\frac{\tau_\theta}{\tau_A} \leq \phi^2 \leq \frac{1}{2}$ for the condition in (70) to be satisfied.

When $\phi = 1$, then the condition in (70) becomes $b \geq \frac{1+\psi}{\psi}$. Since

$$0 \leq b \leq \frac{1+\psi}{\psi}\frac{\tau_\theta}{\left(1+\frac{(1-\eta_H)\eta_c}{\psi}\right)\tau_A + \tau_\theta} \leq \frac{1+\psi}{\psi}$$

from Proposition **??**, this condition can be satisfied only when $b = \frac{1+\psi}{\psi}$, which is the value of b in the perfect-information benchmark, in which case $Var\left[\log P\right] = Var\left[\log P^{perf}\right]$. Thus, when $\phi = 1$, variance with informational frictions is always less than that of the perfect-information benchmark.

3.B: Model Extension

3.B.1 The Equilibrium

Our model features a noisy rational expectations equilibrium, which requires clearing of the two housing markets that are consistent with the optimal behavior of both households and home builders:

- Household optimization: $\left\{\{H_j\}_{j\in[0,1]}, \{M_j\}_{j\in[0,1]}, l_i\right\}$ solves each household's maximization problem.

- Builder optimization: $\{S_j\}_{j\in\{H,M\}}$ solves the builders' maximization problem.

- At $t = 2$, the market for secondary homes clears:

$$Q_2 \int_{-\infty}^{\infty} M_i(\theta_i, P, Q_1)\, d\Phi(v_i) = W.$$

- At $t = 1$, the markets for both primary and secondary homes clear:

$$\begin{aligned} \int_{-\infty}^{\infty} H_i(\theta_i, P, Q_1)\, d\Phi(v_i) &= P^k e^{\xi_H}, \\ \int_{-\infty}^{\infty} M_i(\theta_i, P, Q_1)\, d\Phi(v_i) &= Q_1^k e^{\xi_M}, \end{aligned}$$

  where each household's housing demands $H_i(\theta_i, P, Q_1)$ and $M_i(\theta_i, P, Q_1)$ depend on its private signal θ_i and the housing prices P and Q_1. The two types of demands from the households are integrated over the idiosyncratic component of their private signals $\{\nu_i\}_{i\in[0,1]}$.

We first solve for the optimal labor and housing choices for a household given its utility function and budget constraint in (14). Similar to the baseline model, household $i's$ optimal primary housing H_i and labor demand l_i are given by Proposition 1. The first order condition for secondary home demand reveals that

$$E\left[\frac{\eta_H}{C_i} U\left(\{H_j\}_{j\in[0,1]}, C_i\right) Q_2 \middle| \mathcal{I}_i\right] = E\left[\frac{\eta_H}{C_i} U\left(\{H_j\}_{j\in[0,1]}, C_i\right) \middle| \mathcal{I}_i\right] Q_1. \tag{3.B.1}$$

Combining the market clearing conditions in the market for secondary homes at $t = 1$ and $t = 2$, and substituting $W = E\left[e^A \middle| \mathcal{I}_c\right]$, one also has that

$$Q_2 e^{\xi_M} = Q_1^{-k} E\left[e^A \middle| \mathcal{I}_c\right]. \tag{3.B.2}$$

The right-hand side, which is public information, is a function of the primary and secondary housing prices, P and Q_1. Thus, it follows that $Q_2 e^{\xi_M}$ must also be measurable with respect to the public information $\mathcal{I}_c$, and consequently $Q_2 e^{\xi_M} = f(P, Q_1)$ for some function $f : \mathbb{R}_+^2 \to \mathbb{R}_+$.

Substituting equation (3.B.2) into (3.B.1), it follows that Q_1 satisfies

$$Q_1^{1+k} = \frac{E\left[\left.\frac{\eta_H}{C_i} U\left(\{H_j\}_{j\in[0,1]}, C_i\right) e^{-\xi_M} \right| \mathcal{I}_i\right]}{E\left[\left.\frac{\eta_H}{C_i} U\left(\{H_j\}_{j\in[0,1]}, C_i\right) \right| \mathcal{I}_i\right]} E\left[\left. e^A \right| \mathcal{I}_c\right].$$

As the left-hand side of this equation is public information, the right-hand side must be public information as well. This is satisfied if $E\left[\left. e^{-\xi_M} \right| \mathcal{I}_i\right]$ is public information, i.e., $E\left[\left. e^{-\xi_M} \right| \mathcal{I}_i\right] = E\left[\left. e^{-\xi_M} \right| \mathcal{I}_c\right]$, in which case ξ_M must be public information. This happens because Q_1 is fully revealing about the supply shock ξ_M. Since ξ_M is public information, it follows by equation (3.B.2) that Q_2 must also be public information.[73] Then, equation (3.B.1) implies that

$$Q_2 = Q_1.$$

This means that households choose their demand for secondary homes so that they earn zero profit in equilibrium. They, consequently, serve as intermediaries between the immigrants and home builders in purchasing their homes, and there is effectively only one secondary housing price in the economy.

From the market clearing condition in the secondary housing market at $t = 1$, for this to be the case secondary housing demand M_i cannot depend on the private signals of households, and consequently it must be identical across households, $M_i = M$. By the market clearing condition in the secondary housing market at $t = 2$, then, it follows that

$$M = \frac{1}{Q_1} E\left[\left. e^A \right| \mathcal{I}_c\right]. \tag{3.B.3}$$

The secondary home demand, has two components: a cost component in the denominator Q_1 and a component in the numerator, which reflects migration driven by the immigrants' expectation based solely on housing prices $E\left[\left. e^A \right| \mathcal{I}_c\right]$. The second term is the additional piece that

[73] It is easy to see that there cannot be a log-linear equilibrium in which ξ_M is not public information. Suppose ξ_M is not public information and a log-linear equilibrium exists. Then Q_1 takes the log-linear form

$$\log Q_1 = \frac{1}{1+k}\left(E\left[A | \mathcal{I}_c\right] - E\left[\xi_M | \mathcal{I}_i\right]\right) + Q_0,$$

for some constant Q_0. Since $E\left[A | \mathcal{I}_c\right]$ is public information, it follows that $E\left[\xi_M | \mathcal{I}_i\right]$ must also be public information, which is a contradiction.

distinguishes primary from secondary homes. Buyers of secondary homes—immigrants—are less informed than the households in the neighborhood, and consequenly have to heavily rely on the prices to determine their housing demand.

By clearing the aggregate housing demand of primary and secondary homes from the households with the corresponding supplies from the builders, we derive the housing market equilibrium. Despite the nonlinearity in each household's demands and in the supplies, we obtain a tractable, unique log-linear equilibrium. The following proposition summarizes the housing price and each household's housing demand in this equilibrium.

PROPOSITION 3.B1: At $t = 1$, the primary and secondary housing markets have a unique log-linear equilibrium: 1) The primary and secondary housing demands of household i are log-linear function of its private signal θ_i, $\log P$, and $\log Q_1$:

$$\begin{aligned} \log H_i &= h_\theta \theta_i + h_P \log P + h_Q \log Q_1 + h_0, \\ \log M_i &= m_P \log P + m_Q \log Q_1 + m_0. \end{aligned}$$

with $h_\theta = b$, where $b \geq 0$ is a positive, real root of equation (3.B.27), and all other coefficients given in the proof.
2) The primary housing price is a log-linear function of A, ξ_H, and $\log Q_1$:

$$\log P = p_A A + p_\xi \xi_H + p_Q \log Q_1 + p_0,$$

and the secondary housing price is a log-linear function of ξ_M and $\log P$:

$$\log Q_1 = q_\xi \xi_M + q_P \log P + q_0,$$

and all coefficients given in the proof.

Proposition 3.B1 establishes that the housing prices P and Q_1 are log-linear functions of the common productivity of households A and the housing supply shocks ξ_H and ξ_M, respectively, and that each household's housing demand is a log-linear function of its private signal θ_i and the log housing prices $\log P$ and $\log Q_1$. The informational role of housing prices introduces a rich interaction between the two in housing demand. The conditional expectation of A

$$E\left[A \mid \mathcal{I}_i\right] = \frac{\tau_A \bar{A} + \tau_\theta \theta_i + \frac{p_A^2}{p_\xi^2} \frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2} x_H - \frac{p_A}{p_\xi} \frac{1}{(2+\alpha^2)\alpha^2(\phi k)^2} \frac{1}{q_\xi} x_Q}{\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2} \frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}}$$

now captures an informational hedging effect that having two markets to aggregate households' private information introduces because the supply shocks in both housing markets are correlated. The negative terms in the loadings on the two price signals capture that a high price in both housing markets can also be a sign of a low common supply shock ξ. This introduces a hedging effect that reduces the loadings on the two price signals because of this correlation.

Given that the expressions for the two housing prices in Proposition B1 form a linear system, we can solve it to express primary and secondary housing prices in terms of the fundamental shocks to the economy

$$\begin{aligned}\log P &= \frac{p_A}{1-p_Q q_P}A + \frac{p_\xi}{1-p_Q q_P}\xi_H + \frac{p_Q q_\xi}{1-p_Q q_P}\xi_M + \frac{p_0 + p_Q q_0}{1-p_Q q_P},\\ \log Q_1 &= q_P\frac{p_A}{1-p_Q q_P}A + q_P\frac{p_\xi}{1-p_Q q_P}\xi_H + \left(q_\xi + q_P\frac{p_Q q_\xi}{1-p_Q q_P}\right)\xi_M + q_0 + q_P\frac{p_0 + p_Q q_0}{1-p_Q q_P}.\end{aligned}$$

We characterize the perfect-information benchmark in the following proposition.

PROPOSITION 3.B2: As informational frictions dissipate, the demand for primary homes is

$$\begin{aligned}\log H &= \frac{1+\psi}{\psi}A - \frac{1+\psi-\eta_H}{\psi}\log P + \frac{(1+\psi)\eta_H - (1-\eta_H)(1+\psi-\eta_H)}{2\psi}\tau_\varepsilon^{-1}\\ &\quad + \log\left((1-\eta_c)\left(\frac{1-\eta_H}{\eta_H}\right)\left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}}\right).\end{aligned}$$

and the primary housing price is

$$\begin{aligned}\log P &= \frac{1+\psi}{1+\psi(1+k)-\eta_H}A - \frac{\psi}{1+\psi(1+k)-\eta_H}\xi_H + \frac{(1+\psi)\eta_H - (1-\eta_H)(1+\psi-\eta_H)}{2(1+\psi(1+k)-\eta_H)}\tau_\varepsilon^{-1}\\ &\quad + \frac{\psi}{1+\psi(1+k)-\eta_H}\log\left((1-\eta_c)\left(\frac{1-\eta_H}{\eta_H}\right)\left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}}\right).\end{aligned}$$

The demand for secondary homes is

$$\log M = A - \log Q_1,$$

and the secondary housing price is

$$\log Q_1 = \frac{1}{1+k}A - \frac{1}{1+k}\xi_M.$$

3.B.2 Technical Proofs

Proof of Proposition 3.B1

We first conjecture that each household's housing purchasing and the housing price take the following log-linear forms:

$$\log H_i = h_P \log P + h_Q \log Q_1 + h_\theta \theta_i + h_0, \tag{3.B.6}$$

$$\log P = p_A A + p_Q \log Q_1 + p_\xi \xi_H + p_0, \tag{3.B.7}$$

$$\log M_i = m_P \log P + m_Q \log Q_1 + m_0,$$

$$\log Q_1 = q_P \log P + q_\xi \xi_M + q_0,$$

where the coefficients h_0, h_P, h_Q, h_θ, p_0, p_A, p_Q, p_ξ, m_0, m_P, m_Q, q_0, q_P, and q_ξ will be determined by equilibrium conditions.

Given the conjectured functional form for H_i, it follows that

$$E\left[\left(\int_{[0,1]} H_j dj\right)^{\eta_c(1-\eta_H)} e^{\eta_H A_i}\middle| \mathcal{I}_i\right]$$
$$= e^{\eta_c(1-\eta_H)\left(h_0+h_P\log P+h_Q\log Q_1+\frac{1}{2}h_\theta^2\tau_\theta^{-1}\right)+\frac{1}{2}\eta_H^2\tau_\varepsilon^{-1}} E\left[e^{(\eta_H+\eta_c(1-\eta_H)h_\theta)A}\middle|\mathcal{I}_i\right],$$

which uses the fact that A is independent of ε_j and exploits the Law of Large Number for the continuum when integrating over households, which still holds if we subtract sets of measure 0 from the integral.

Define

$$x_H \equiv \frac{\log P - p_0 - p_Q \log Q_1 - p_\xi \bar{\xi}}{p_A} = A + \frac{p_\xi}{p_A}\left(\xi_H - \bar{\xi}\right),$$
$$x_Q \equiv \log Q_1 - q_0 - q_P \log P - q_\xi \bar{\xi} = q_\xi\left(\xi_M - \bar{\xi}\right),$$

which are sufficient statistics for information contained in P and Q_1. Then, household i's information set is

$$X_i = \begin{bmatrix} \theta_i \\ x_H \\ x_Q \end{bmatrix} = \begin{bmatrix} A+\nu_i \\ A+\frac{p_\xi}{p_A}\left(\xi-\bar{\xi}+\phi k e_H\right) \\ q_\xi\left(\xi-\bar{\xi}+\phi k e_M\right) \end{bmatrix}.$$

We can normalize the noise variables in the signals to standard normal distribution:

$$X_i = \begin{bmatrix} A+\sqrt{b}u_1 \\ A+\sqrt{c}\left(\sqrt{e}u_2+\sqrt{f}u_3\right) \\ \sqrt{d}\left(\sqrt{e}u_2+\sqrt{f}u_4\right) \end{bmatrix}$$

where $b=\tau_\theta^{-1},\ c=\frac{p_\xi^2}{p_A^2},\ d=q_\xi^2,\ e=(\phi k)^2$, and $f=(\alpha\phi k)^2$, and $[u_1,u_2,u_3,u_4]'\sim\mathcal{N}(0,I_4)$. Each household has a Gaussian prior over A:

$$A\sim\mathcal{N}\left(\bar{A},a\right),$$

where $a=\tau_A^{-1}$. Thus,

$$E[X_i]=\left[\begin{array}{ccc}\bar{A} & \bar{A} & 0\end{array}\right]'.$$

Then, conditional on observing the vector of signals X_i, household i arrives to its conditional belief $E[A\mid\mathcal{I}_i]$:

$$E[A\mid\mathcal{I}_i]=\bar{A}+Cov[A,X_i]'Var[X_i]^{-1}(X_i-E[X_i]),$$

where

$$Var[X_i]=\left[\begin{array}{ccc}a+b & a & 0\\ a & a+ce+cf & \sqrt{cde}\\ 0 & \sqrt{cde} & de+df\end{array}\right],$$

and

$$Cov[A,X_i]'=\left[\begin{array}{cc}a\mathbf{1}_{1\times2} & 0\end{array}\right].$$

The inverse of this matrix can be found using co-factors and the determinant simplified through a series of row manipulations. Following this approach, one can arrive at

$$Cov[A,X_i]'Var[X_i]^{-1}=a\frac{\left[\begin{array}{ccc}cfd(2e+f) & bd(e+f) & -b\sqrt{cde}\end{array}\right]}{(a+b)\,cdf\,(2e+f)+bda(e+f)}=\frac{\left[\begin{array}{ccc}\frac{1}{b} & \frac{e+f}{(2e+f)cf} & \frac{-e}{(2e+f)f\sqrt{cd}}\end{array}\right]}{\frac{1}{a}+\frac{1}{b}+\frac{e+f}{cf(2e+f)}}$$

Substituting this expression into $E[A\mid\mathcal{I}_i]$, we obtain

$$E[A\mid\mathcal{I}_i]=\frac{1}{\frac{1}{a}+\frac{1}{b}+\frac{e+f}{cf(2e+f)}}\left(\frac{1}{a}\bar{A}+\left[\begin{array}{ccc}\frac{1}{b} & \frac{e+f}{(2e+f)cf} & \frac{-e}{(2e+f)f\sqrt{cd}}\end{array}\right]X_i\right).$$

Similarly, the conditional variance is given by

$$Var[A\mid\mathcal{I}_i]=\frac{1}{\frac{1}{a}+\frac{1}{b}+\frac{e+f}{cf(2e+f)}}.$$

Substituting for a,b,c,d,e, and f, we arrive at

$$E[A\mid\mathcal{I}_i]=\frac{\tau_A\bar{A}+\tau_\theta\theta_i+\frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}x_H-\frac{p_A}{p_\xi}\frac{1}{(2+\alpha^2)\alpha^2(\phi k)^2}\frac{1}{q_\xi}x_Q}{\tau_A+\tau_\theta+\frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}},$$

and

$$Var\left[A \mid \mathcal{I}_i\right] = \frac{1}{\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}}.$$

It is straightforward to see that, under $\mathcal{I}_c$, the conditional expectation and variance of A take the form

$$E\left[A \mid \mathcal{I}_c\right] = \frac{\tau_A \bar{A} + \frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2} x_H - \frac{p_A}{p_\xi}\frac{1}{(2+\alpha^2)\alpha^2(\phi k)^2}\frac{1}{q_\xi} x_Q}{\tau_A + \frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}},$$

and

$$Var\left[A \mid \mathcal{I}_c\right] = \frac{1}{\tau_A + \frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}}.$$

Therefore,

$$\begin{aligned}
& \log E\left[\left(\int_{[0,1]} H_j dj\right)^{\eta_c(1-\eta_H)} e^{\eta_H A_i} \middle| \mathcal{I}_i\right] \\
= \; & \left(\eta_H + \eta_c\left(1-\eta_H\right) h_\theta\right) \frac{\tau_A \bar{A} + \tau_\theta \theta_i + \frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2} x_H - \frac{p_A}{p_\xi}\frac{1}{(2+\alpha^2)\alpha^2(\phi k)^2}\frac{1}{q_\xi} x_Q}{\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}} \\
& + \eta_c\left(1-\eta_H\right)\left(h_0 + h_P \log P + h_Q \log Q_1 + \frac{1}{2} h_\theta^2 \tau_\theta^{-1}\right) + \frac{1}{2}\eta_H^2 \tau_\varepsilon^{-1} \\
& + \frac{1}{2}\left(\left(\eta_H + \eta_c\left(1-\eta_H\right) h_\theta\right)^2\right)\left(\tau_A + \tau_\theta + \frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2}\right)^{-1}.
\end{aligned}$$

Substituting this expression into equation (5) in the main paper and matching coefficients with the conjectured log-linear form in (3.B.6), it follows that

$$h_\theta = \frac{1+\psi+\eta_c(1-\eta_H)h_\theta}{\psi+(1-\eta_H)\eta_c}\frac{\tau_\theta}{\tau_A+\tau_\theta+\frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}}, \tag{3.B.8}$$

$$h_P = \frac{1+\psi+\eta_c(1-\eta_H)h_\theta}{\psi}\frac{\frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}\frac{1}{p_A}+\frac{p_A}{p_\xi}\frac{1}{(2+\alpha^2)\alpha^2(\phi k)^2}\frac{q_P}{q_\xi}}{\tau_A+\tau_\theta+\frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}} - \frac{1+\psi-\eta_H}{\psi}, \tag{3.B.9}$$

$$h_Q = -\frac{1+\psi+\eta_c(1-\eta_H)h_\theta}{\psi}\frac{\frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}\frac{p_Q}{p_A}+\frac{p_A}{p_\xi}\frac{1}{(2+\alpha^2)\alpha^2(\phi k)^2}\frac{1}{q_\xi}}{\tau_A+\tau_\theta+\frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}}, \tag{3.B.10}$$

$$\begin{aligned} h_0 = {} & \frac{1+\psi+\eta_c(1-\eta_H)h_\theta}{\psi}\frac{\tau_A\bar{A}-\frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}\frac{p_0+p_\xi\bar{\xi}}{p_A}+\frac{p_A}{p_\xi}\frac{1}{(2+\alpha^2)\alpha^2(\phi k)^2}\frac{q_0+q_\xi\bar{\xi}}{q_\xi}}{\tau_A+\tau_\theta+\frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}} \\ & +\frac{1+\psi+\eta_c(1-\eta_H)h_\theta}{2\psi}\left(\eta_H+\eta_c(1-\eta_H)h_\theta\right)\left(\tau_A+\tau_\theta+\frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}\right)^{-1} \\ & +\frac{1+\psi-\eta_H}{2\psi}\left(\eta_H-1+\eta_c(1-\eta_H)h_\theta\right)\left(\tau_A+\tau_\theta+\frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{(2+\alpha^2)\alpha^2(\phi k)^2}\right)^{-1} \\ & +\log\left((1-\eta_c)\left(\frac{1-\eta_H}{\eta_H}\right)\left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}}\right) \\ & +\frac{1}{2\psi}\left(\eta_c(1-\eta_H)h_\theta^2\tau_\theta^{-1}+\eta_H^2\tau_\varepsilon^{-1}\right). \end{aligned} \tag{3.B.11}$$

By aggregating households' housing demand and the builders' supply and imposing market clearing in the housing market, we have

$$\begin{aligned} & h_0+h_P(p_0+p_Q\log Q_1+p_A A+p_\xi\xi_H)+h_Q\log Q_1+h_\theta A+\frac{1}{2}h_\theta^2\tau_\theta^{-1} \\ = {} & \xi_H+k(p_0+p_Q\log Q_1+p_A A+p_\xi\xi_H). \end{aligned}$$

Matching coefficients of the two sides of the equation leads to the following four conditions

$$h_0+h_P p_0+\frac{1}{2}h_\theta^2\tau_\theta^{-1} = kp_0, \tag{3.B.12}$$

$$h_P p_A+h_\theta = kp_A, \tag{3.B.13}$$

$$h_P p_\xi = 1+kp_\xi, \tag{3.B.14}$$

$$h_P p_Q+h_Q = kp_Q. \tag{3.B.15}$$

It follows from equation (3.B.14) that

$$p_\xi = -\frac{1}{k-h_P}, \tag{3.B.16}$$

and further from equation (3.B.13) that

$$p_A = \frac{h_\theta}{k - h_P}. \tag{3.B.17}$$

Thus, by taking the ratio of equations (3.B.17) and (3.B.16), we arrive at

$$\frac{p_A}{p_\xi} = -h_\theta.$$

From equation (3.B.8), $h_\theta > 0$, and therefore p_A and p_ξ must have opposite signs. Since the primary housing price cannot react negatively to an increase in neighborhood strength A, it must be the case that $p_A \geq 0$ and $p_\xi \leq 0$.

Finally, from equation (3.B.15), one has that

$$p_Q = \frac{h_Q}{k - h_P}. \tag{3.B.18}$$

From equation (3.B.3) and the functional form for the conditional expectations and variances for A given $\mathcal{I}_i$ and $\mathcal{I}_c$, M_i can be written as

$$\begin{aligned}\log M_i &= -\log Q_1 + \frac{\tau_A \bar{A}\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2 + \frac{p_A^2}{p_\xi^2}\left(1+\alpha^2\right)x_H - \frac{p_A}{p_\xi}\frac{1}{q_\xi}x_Q}{\tau_A\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2 + \frac{p_A^2}{p_\xi^2}\left(1+\alpha^2\right)} \\ &\quad + \frac{1}{2}\left(\tau_A + \frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2}\right)^{-1}.\end{aligned}$$

Given the common knowledge expectation about A and household $i's$ expectation about A, and matching coefficients of the resulting expression with the conjectured form for $\log M_i$, one arrives at the following restrictions

$$m_P = \frac{\frac{p_A^2}{p_\xi^2}\left(1+\alpha^2\right)\frac{1}{p_A} + \frac{p_A}{p_\xi}\frac{q_P}{q_\xi}}{\tau_A\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2 + \frac{p_A^2}{p_\xi^2}\left(1+\alpha^2\right)}, \tag{3.B.19}$$

$$m_Q = -1 - \frac{\frac{p_A^2}{p_\xi^2}\left(1+\alpha^2\right)\frac{p_Q}{p_A} + \frac{p_A}{p_\xi}\frac{1}{q_\xi}}{\tau_A\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2 + \frac{p_A^2}{p_\xi^2}\left(1+\alpha^2\right)}, \tag{3.B.20}$$

$$\begin{aligned}m_0 &= \frac{\tau_A \bar{A}\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2 - \frac{p_A^2}{p_\xi^2}\left(1+\alpha^2\right)\frac{p_0+p_\xi\bar{\xi}}{p_A} + \frac{p_A}{p_\xi}\frac{q_0+q_\xi\bar{\xi}}{q_\xi}}{\tau_A\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2 + \frac{p_A^2}{p_\xi^2}\left(1+\alpha^2\right)} \\ &\quad + \frac{1}{2}\left(\tau_A + \frac{p_A^2}{p_\xi^2}\frac{1+\alpha^2}{\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2}\right)^{-1}.\end{aligned} \tag{3.B.21}$$

Aggregating secondary housing demand across households and imposing market clearing

$$m_P \log P + m_Q \left(q_P \log P + q_\xi \xi_M + q_0\right) + m_0 = \xi_M + k \left(q_P \log P + q_\xi \xi_M + q_0\right).$$

Matching coefficients of the two sides of the equation leads to the following four conditions

$$\begin{aligned} m_0 + m_Q q_0 &= k q_0, & (3.B.22)\\ m_Q q_\xi &= 1 + k q_\xi, & (3.B.23)\\ m_Q q_P + m_P &= k q_P. & (3.B.24) \end{aligned}$$

It follows from equation (3.B.23) that

$$q_\xi = -\frac{1}{k - m_Q}, \tag{3.B.25}$$

Finally, from equation (3.B.24), one has that

$$q_P = \frac{m_P}{k - m_Q}. \tag{3.B.26}$$

Substituting $\frac{p_A}{p_\xi} = -h_\theta$ into equation (3.B.8), and defining $b = -\frac{p_A}{p_\xi}$, we arrive at

$$\frac{1+\alpha^2}{\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2} b^3 + \left(\tau_A + \frac{\psi}{\psi + \left(1-\eta_H\right)\eta_c}\tau_\theta\right) b - \frac{1+\psi}{\psi + \left(1-\eta_H\right)\eta_c}\tau_\theta = 0, \tag{3.B.27}$$

which is a cubic polynomial that identifies b, and therefore h_θ since $h_\theta = b$. Furthermore, by the Fundamental Theorem of Algebra, the cubic equation (3.B.27) has one unique, existent root.

Substituting equations (3.B.9) and (3.B.19) into (3.B.10) and (3.B.20), as well as $h_\theta = b = -\frac{p_A}{p_\xi}$, and equations (3.B.17), (3.B.18), and (3.B.26), one also has that

$$\begin{aligned} m_P &= \frac{\left(1+\alpha^2\right) b}{\tau_A \left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2 + \left(1+\alpha^2\right) b^2 - b}\left(k - h_P\right),\\ m_Q &= -\frac{\tau_A \left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2 + \left(1+\alpha^2\right) b^2 + kb + \left(1+\alpha^2\right) b h_Q}{\tau_A \left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2 + \left(1+\alpha^2\right) b^2 - b}, \end{aligned}$$

which can be rewritten as

$$\begin{aligned} m_P &= A\left(b\right)\left(k - h_P\right), & (3.B.28)\\ m_Q &= B\left(b\right) + C\left(b\right) h_Q. & (3.B.29) \end{aligned}$$

Substituting $h_\theta = b = -\frac{p_A}{p_\xi}$ and equations (3.B.17), (3.B.18), and (3.B.26) into (3.B.9) and (3.B.10), one arrives at

$$\begin{aligned} h_P &= \frac{\frac{1+\psi+\eta_c(1-\eta_H)b}{\psi}\left(\left(1+\alpha^2\right)k+m_P\right)b-\left(1+\psi-\eta_H\right)\frac{(\tau_A+\tau_\theta)\left(2+\alpha^2\right)\alpha^2(\phi k)^2+\left(1+\alpha^2\right)b^2}{\psi}}{\left(\tau_A+\tau_\theta\right)\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2+\left(1+\alpha^2\right)b^2+\frac{1+\psi+\eta_c(1-\eta_H)b}{\psi}\left(1+\alpha^2\right)b}, \\ h_Q &= -\frac{\frac{1+\psi+\eta_c(1-\eta_H)b}{\psi}\left(k-m_Q\right)}{\left(\tau_A+\tau_\theta\right)\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2+\left(1+\alpha^2\right)b^2+\frac{1+\psi+\eta_c(1-\eta_H)b}{\psi}\left(1+\alpha^2\right)b}b, \end{aligned}$$

which can be rewritten as

$$h_P = E(b) + F(b)\, m_P, \tag{3.B.30}$$

$$h_Q = G(b)\left(k - m_Q\right). \tag{3.B.31}$$

From equations (3.B.28), (3.B.29), (3.B.30), and (3.B.31), it then follows that

$$\begin{aligned} h_P &= \frac{E(b)+A(b)F(b)k}{1+A(b)F(b)} \\ h_Q &= \frac{G(b)}{1+C(b)G(b)}\left(k-B(b)\right). \end{aligned}$$

Finally, from equations (3.B.11), (3.B.12), (3.B.21), and (3.B.22), one can express p_0 and q_0 with the system of linear equations

$$\left(k-m_Q\right)q_0 = \frac{\tau_A\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)\bar{A}+\alpha^2 b\bar{\xi}-\left(1+\alpha^2\right)b^2\frac{p_0}{p_A}-b\frac{q_0}{q_\xi}}{\tau_A\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)+\left(1+\alpha^2\right)b^2}+\frac{1}{2}\left(\tau_A+b^2\frac{1+\alpha^2}{\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2}\right)^{-1}$$

and

$$\begin{aligned} \left(k-h_P\right)p_0 = {} & \frac{1+\psi+\eta_c\left(1-\eta_H\right)b}{\psi}\frac{\tau_A\bar{A}\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2+\alpha^2 b\bar{\xi}-\left(1+\alpha^2\right)b^2\frac{p_0}{p_A}-b\frac{q_0}{q_\xi}}{\left(\tau_A+\tau_\theta\right)\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2+\left(1+\alpha^2\right)b^2}+\frac{1}{2}b^2\tau_\theta^{-1} \\ & +\frac{1+\psi+\eta_c\left(1-\eta_H\right)b}{2\psi}\left(\eta_H+\eta_c\left(1-\eta_H\right)b\right)\left(\tau_A+\tau_\theta+b^2\frac{1+\alpha^2}{\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2}\right)^{-1} \\ & +\frac{1+\psi-\eta_H}{2\psi}\left(\eta_H-1+\eta_c\left(1-\eta_H\right)b\right)\left(\tau_A+\tau_\theta+b^2\frac{1+\alpha^2}{\left(2+\alpha^2\right)\alpha^2\left(\phi k\right)^2}\right)^{-1} \\ & +\frac{1}{2\psi}\left(\eta_c\left(1-\eta_H\right)b^2\tau_\theta^{-1}+\eta_H^2\tau_\varepsilon^{-1}\right)+\log\left(\left(1-\eta_c\right)\left(\frac{1-\eta_H}{\eta_H}\right)\left(\frac{1-\eta_c}{\eta_c}\right)^{\frac{\eta_c(1-\eta_H)}{\psi}}\right). \end{aligned}$$

From equation (3.B.14), one also has that $h_P = k\left(1+p_\xi^{-1}\right)$. Since $p_\xi \leq 0$, it follows that $h_P < k$ whenever $k > 0$. Similarly, $m_Q < k$ since $q_\xi \leq 0$.

Since we have explicit expressions for all other equilibrium objects as functions of b, and b exists and is unique, it follows that a log-linear equilibrium in the economy exists and is unique.

Proof of Proposition 3.B2

When all households observe A directly, there are no longer information frictions in the economy. Since the households' idiosyncratic productivity components are unobservable, they are now symmetric. Then, it follows that $H_j = H_i = H$. Imposing this symmetry in equation (5) in the main paper, we see that each household's housing demand and the housing price are the same as given in Proposition 3.

In addition, since $H_j = H_i = H$, $M_i = M_j = M$, one has that

$$\log M = A - \log Q_1.$$

By market clearing in the secondary housing market, $\log H = \xi_M + k \log Q_1$, it follows that

$$\log Q_1 = \frac{1}{1+k} A - \frac{1}{1+k} \xi_M.$$

This characterizes the economy in the limit as information frictions dissipate.

3.C: Figures and Tables

Figure 3.1: Case-Shiller Home Price Index

This figure plots the Case-Shiller home price index for the U.S. and three cities, New York, Las Vegas, and Charlotte. The price index is deflated by the CPI and normalized to 100 in 2000.

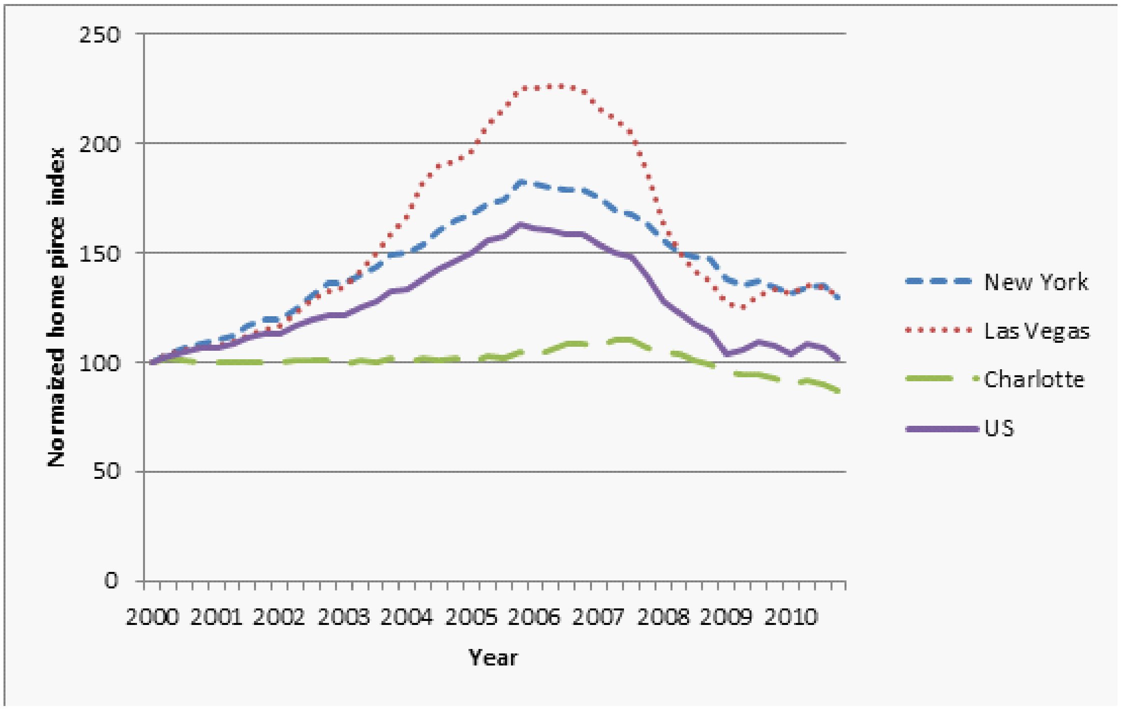

Figure 3.2: Housing Cycle across Three Elasticity Groups with an Equal Number of Counties

This figure is constructed from sorting the counties in the U.S. into three groups based on Saiz's (2010) housing supply elasticity measure, with each group holding an equal number of counties. The top panel depicts the average housing price expansion during the boom years of 2003 to 2006 and the average housing price contraction during the bust years of 2006 to 2009 in each of the groups. The bottom panel depicts the population in each group in the 2000 U.S. census.

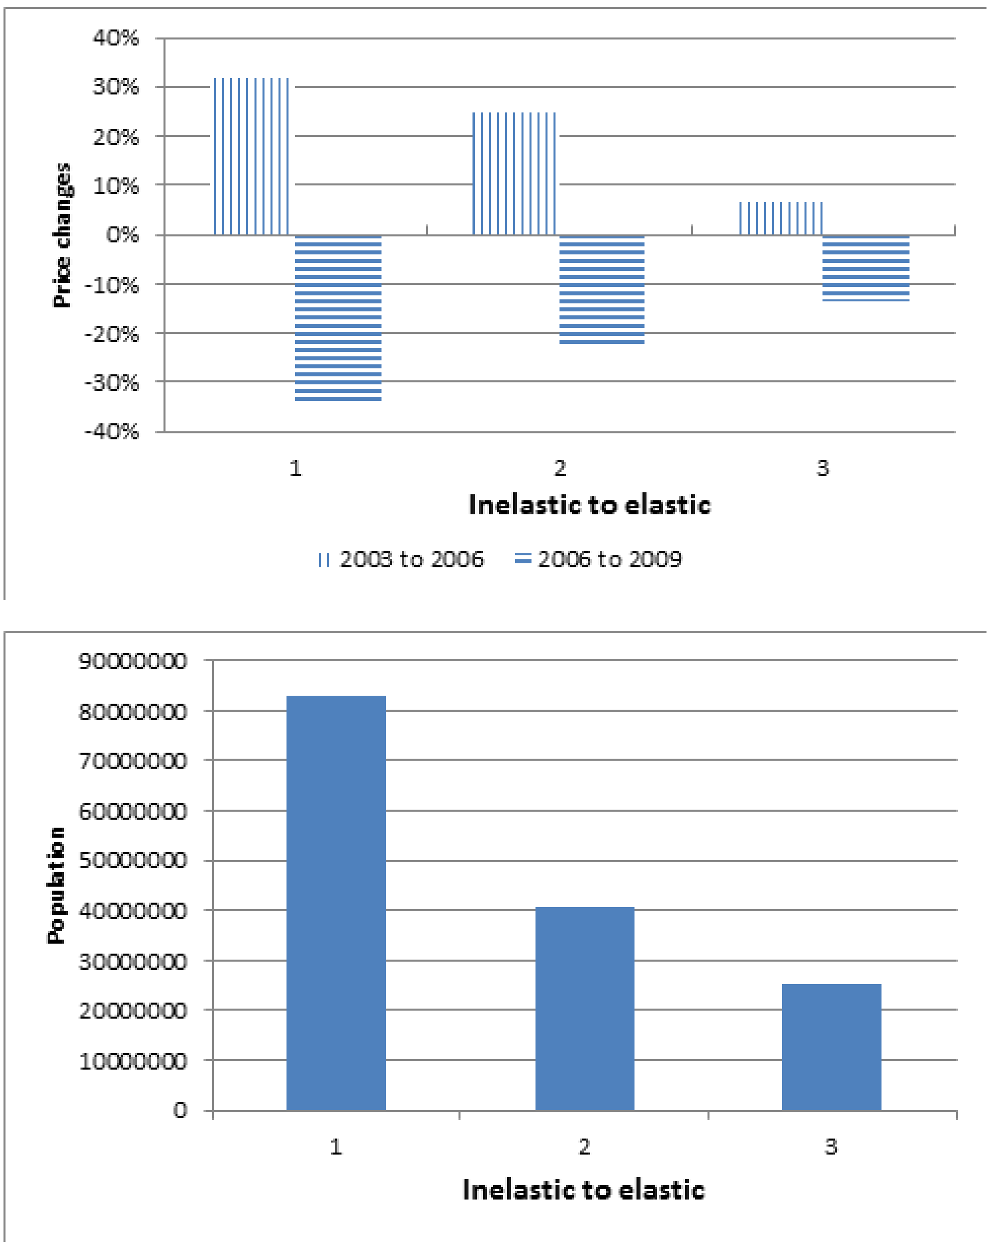

Figure 3.3: Housing Cycle across Ten Elasticity Groups with an Equal Number of Counties

This figure is constructed from sorting the counties in the U.S. into ten groups based on Saiz's (2010) housing supply elasticity measure, with each group holding an equal number of counties. The top panel depicts the average housing price expansion during the boom years of 2003 to 2006 and the average housing price contraction during the bust years of 2006 to 2009 in each of the groups. The bottom panel depicts the population in each group in the 2000 U.S. census.

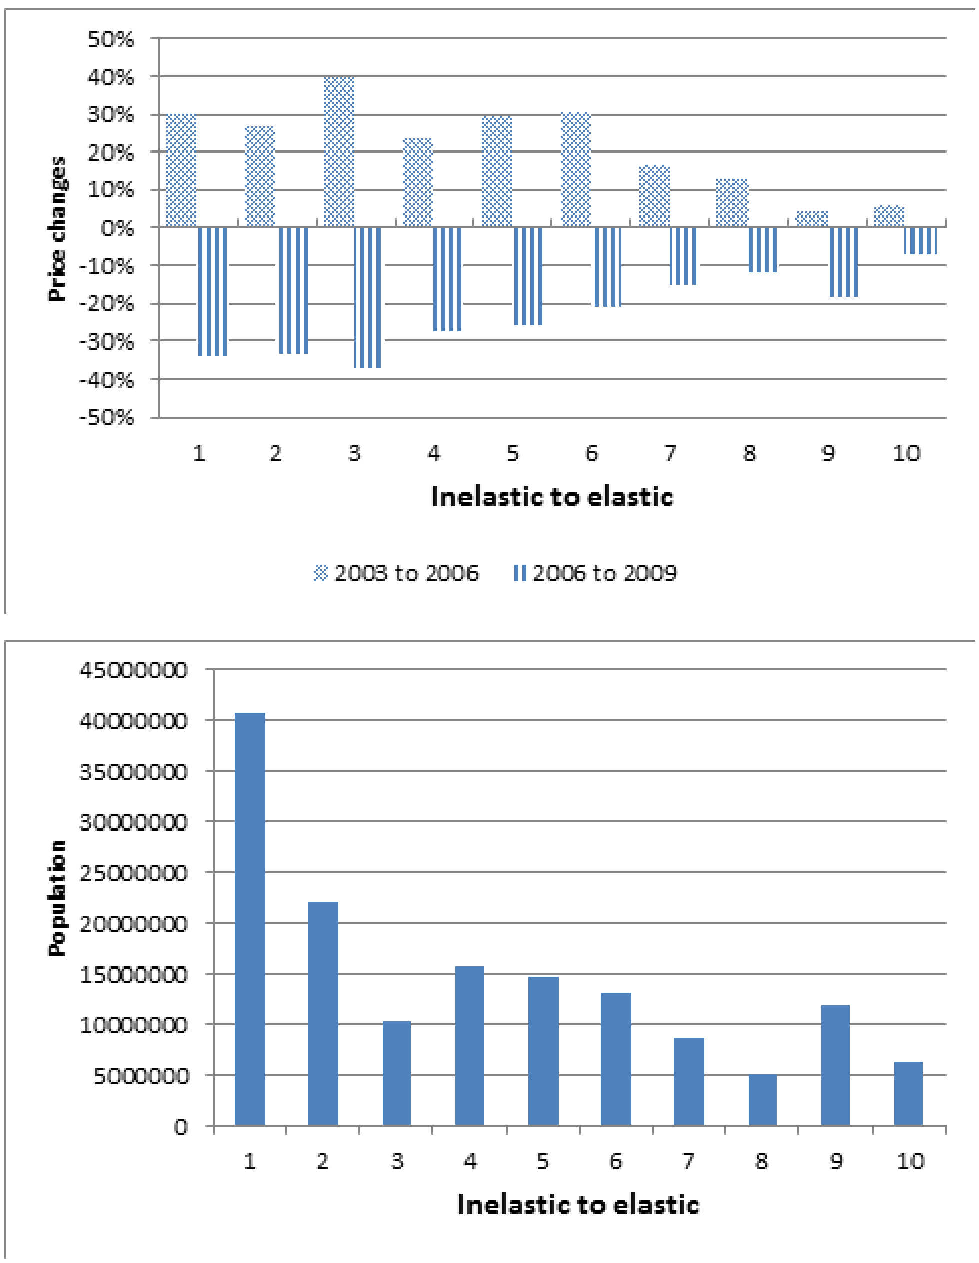

Figure 3.4: Housing Cycle across Three Elasticity Groups with an Equal Population

This figure is constructed from sorting the counties in the U.S. into three groups based on Saiz's (2010) housing supply elasticity measure, with each group holding an equal population. The top panel depicts the average housing price expansion during the boom years of 2003 to 2006 and the average housing price contraction during the bust years of 2006 to 2009 in each of the groups. The bottom panel depicts the population in each group in the 2000 U.S. census.

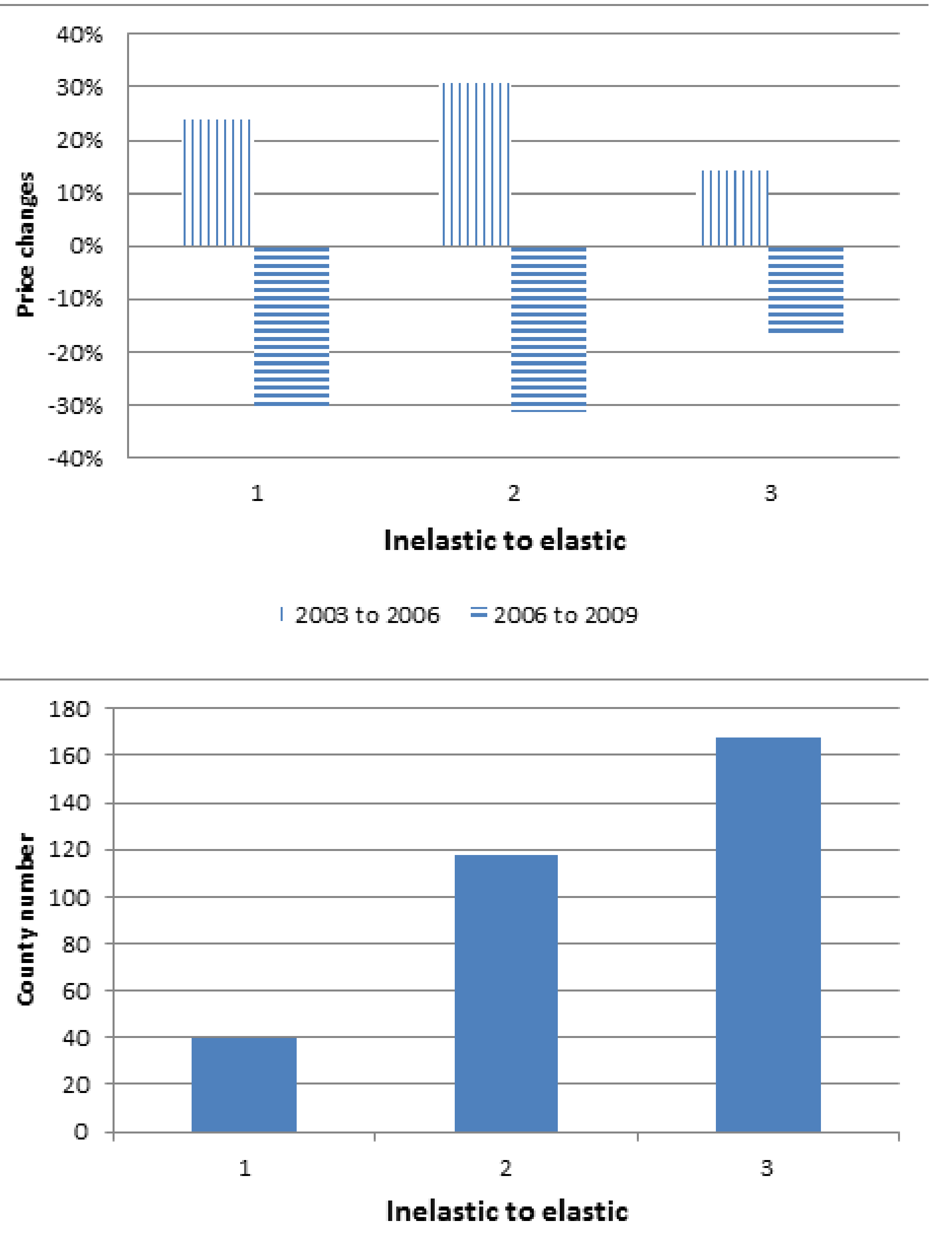

Figure 3.5: Housing Cycle across Ten Elasticity Groups with an Equal Population

This figure is constructed from sorting the counties in the U.S. into ten groups based on Saiz's (2010) housing supply elasticity measure, with each group holding an equal population. The top panel depicts the average housing price expansion during the boom years of 2003 to 2006 and the average housing price contraction during the bust years of 2006 to 2009 in each of the groups. The bottom panel depicts the population in each group in the 2000 U.S. census.

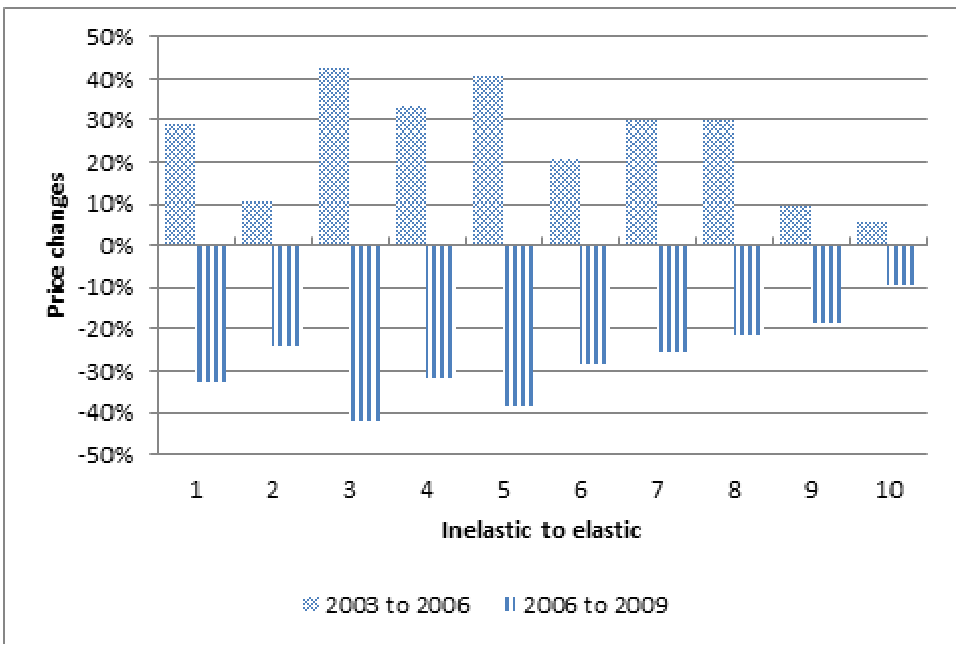

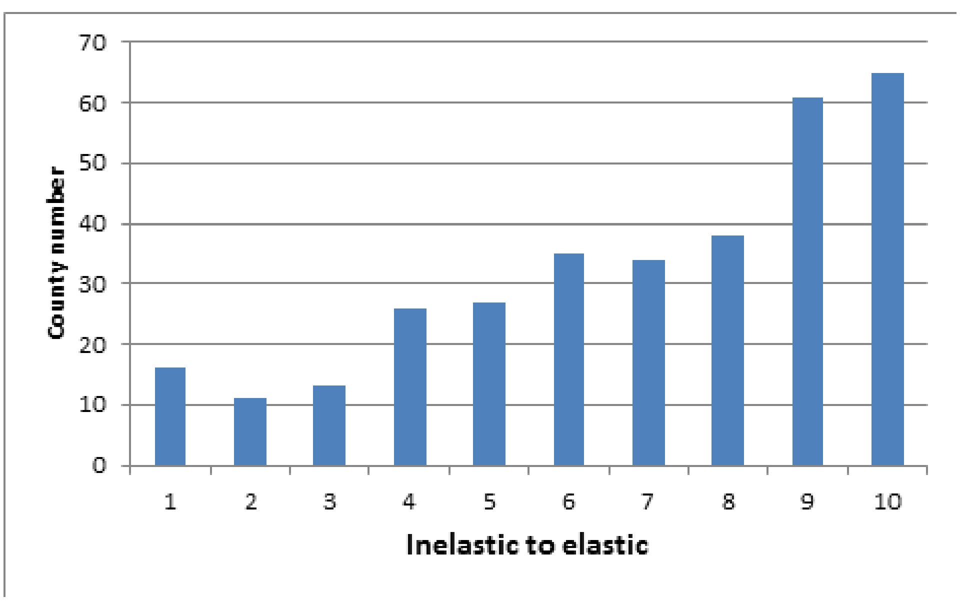

Figure 3.6: Housing Price Variance in the Baseline Model

This figure depicts the log-price variance in the baseline model against the supply elasticity, based on the following parameters: $\tau_\theta = 0.1, \tau_A = 1, \phi = 0.1, \eta_c = 0.5, \psi = 0.6, \eta_H = 0.9$. The solid line depicts the log-price variance in the presence of informational frictions, while the dashed line depicts the log-price variance in the perfect-information benchmark.

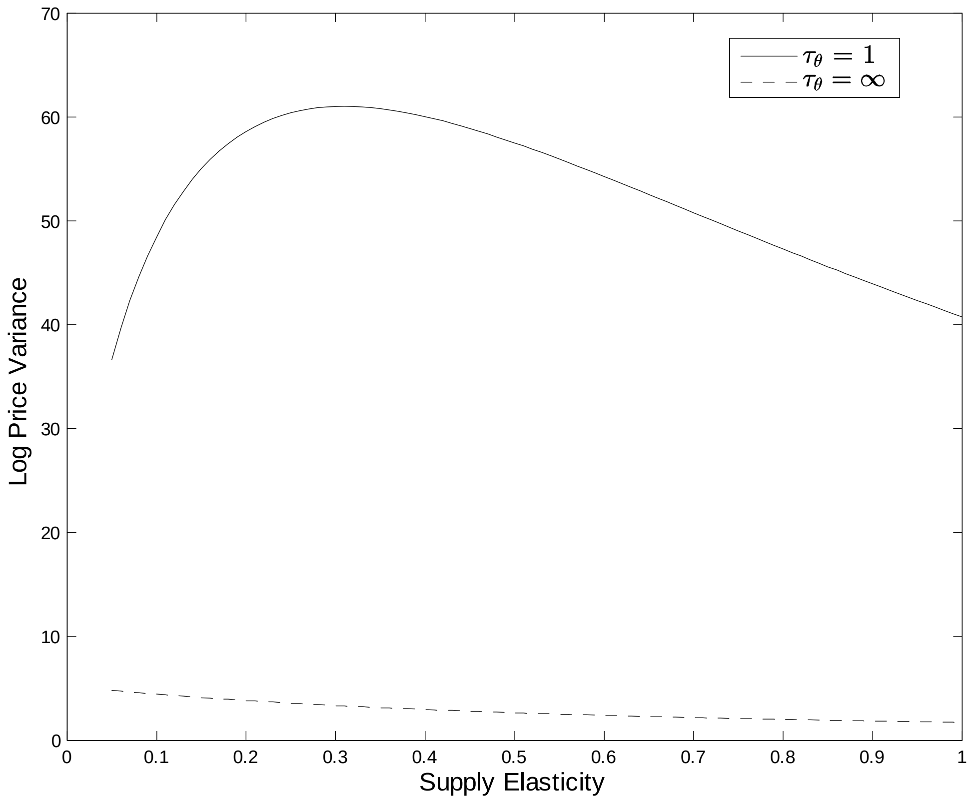

Figure 3.7: Variance of Primary and Secondary Housing Prices and Fraction of Secondary Homes in the Extended Model

This figure depicts the log-price variance of both primary and secondary homes in the extended model in the top two panels, the variance of the fraction of secondary home demand in the bottom left panel, and a scatter plot of the variance of the fraction and the variance of the secondary housing price in the bottom right panel based on the following parameters: $\tau_\theta = 0.1, \tau_A = 1, \phi = 0.1, \eta_c = 0.5, \psi = 0.6, \eta_H = 0.9, \alpha = 1, \tau_\varepsilon = 0.1$.

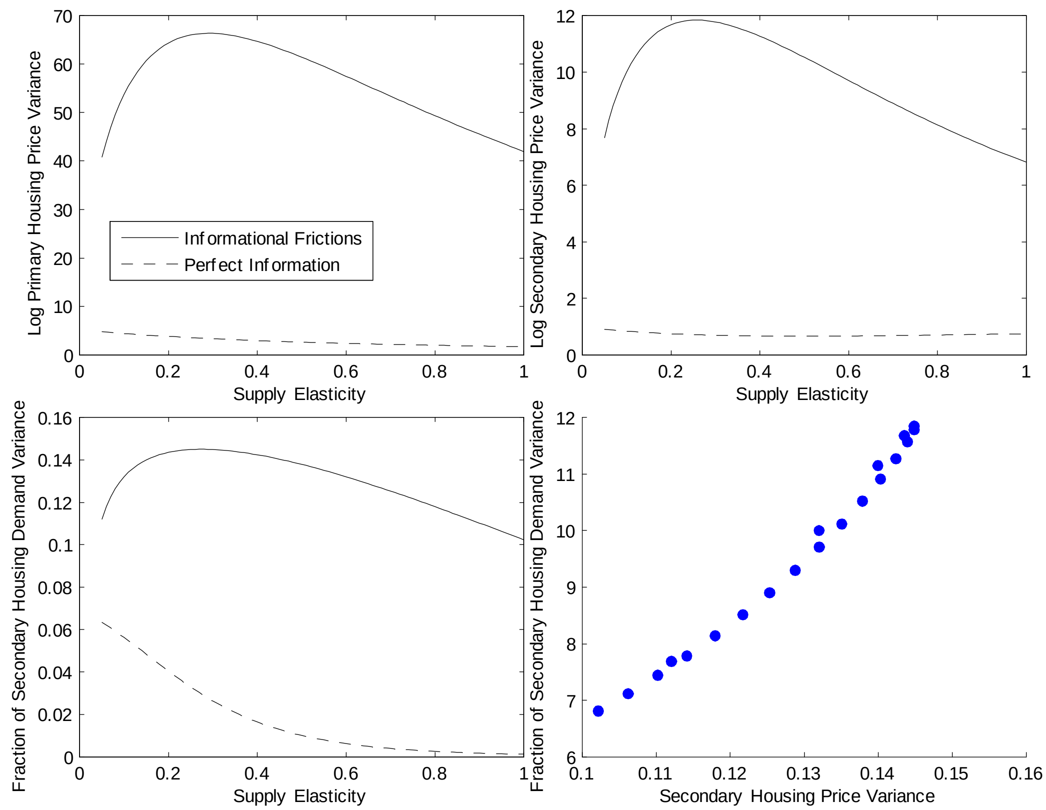

Figure 8: The Share of Non-Owner-Occupied Home Purchases

This figure plots the share of non-owner-occupied home purchases for the U.S. and three cities, New York, Las Vegas, and Charlotte.

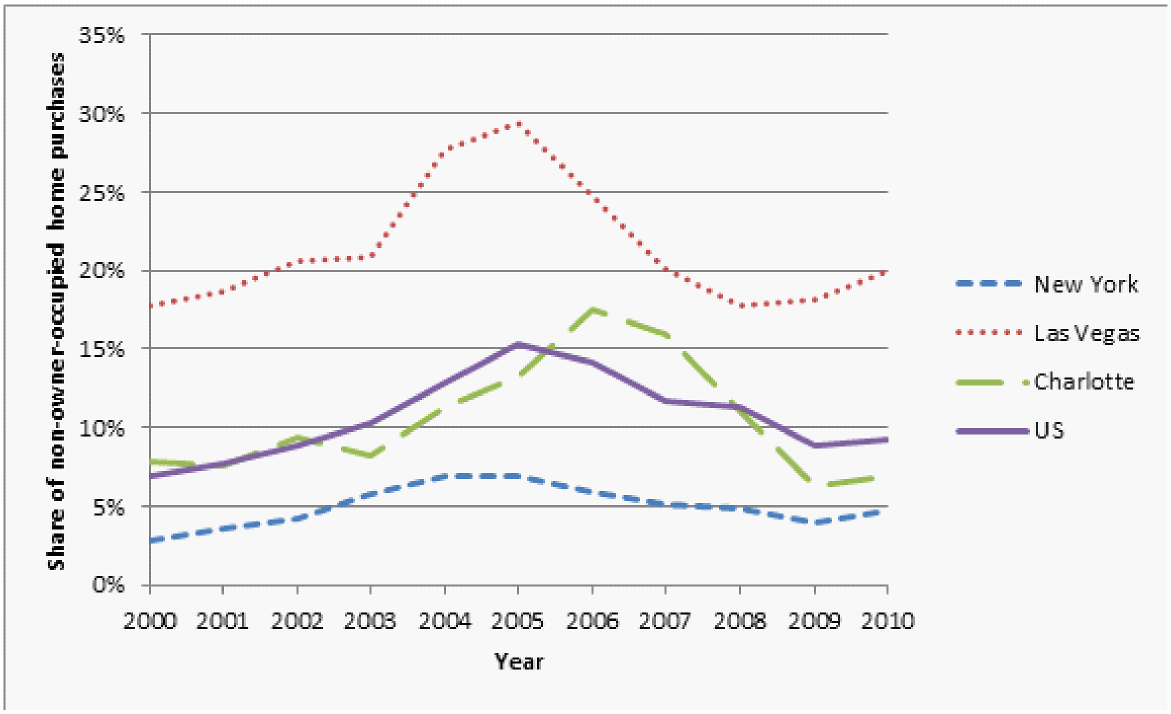

Figure 3.9: The Share of Non-Owner-Occupied Home Purchases in 2005 across Elasticity Groups with an Equal Number of Counties

We use Saiz's (2010) supply elasticity measure to sort the counties in our sample into three groups in the top panel and ten groups in the bottom panel, with each group holding the same number of counties. Each bar measures the average share of non-owner-occupied home purchases in 2005 in each group. The share of non-owner-occupied home purchases in each county is computed from the "Home Mortgage Disclosure Act" data set.

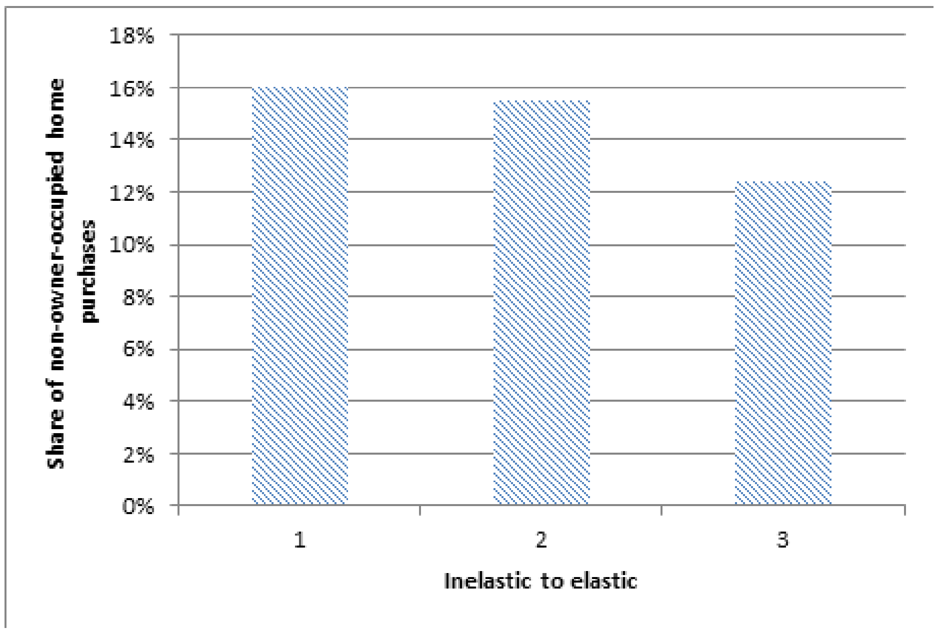

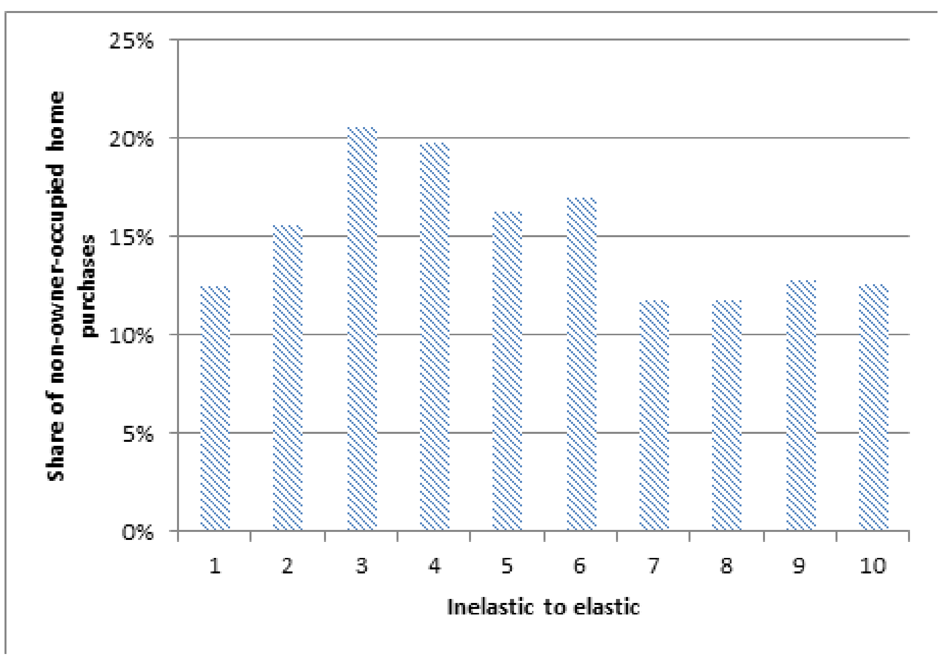

Figure 3.10: The Share of Non-Owner-Occupied Home Purchases in 2005 across Elasticity Groups with an Equal Population

We use Saiz's (2010) supply elasticity measure to sort the counties in our sample into three groups in the top panel and ten groups in the bottom panel, with each group holding the same population. Each bar measures the average share of non-owner-occupied home purchases in 2005 in each group. The share of non-owner-occupied home purchases in each county is computed from the "Home Mortgage Disclosure Act" data set.

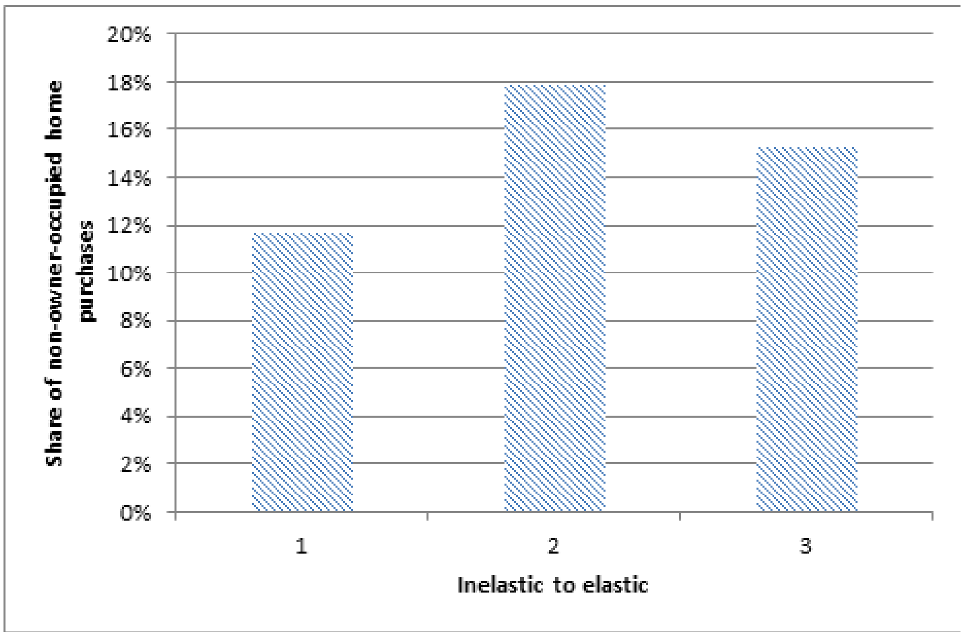

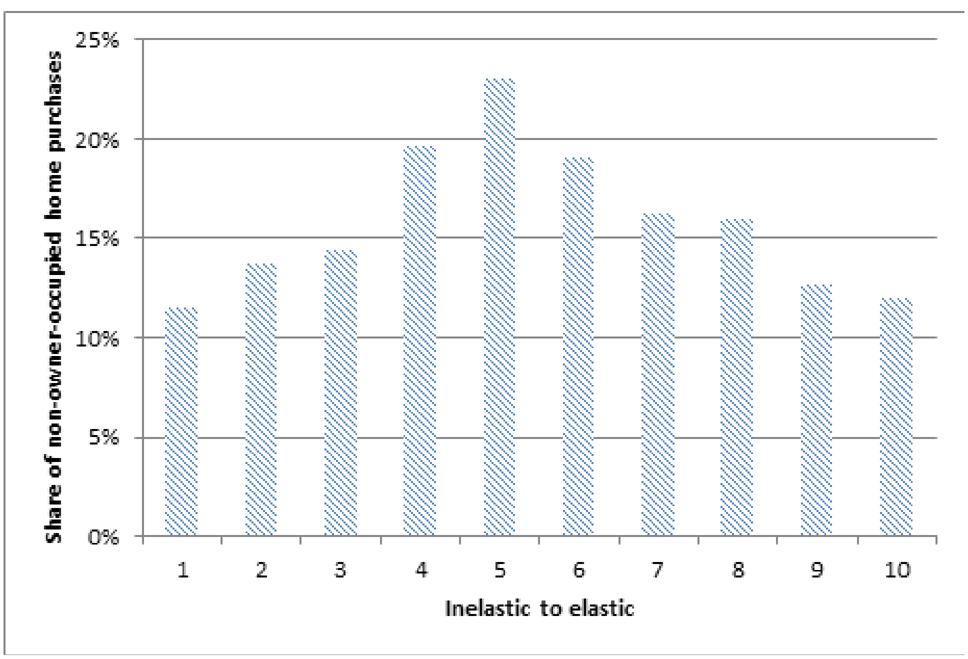

Figure 3.11: Change in the Fraction of Non-Owner-Occupied Home Purchases from 2001 to 2004 and the Recent Housing Cycle

The top panel plots the average housing price expansion during the boom years of 2003 to 2006 against the change of fraction of non-owner occupied home purchases from 2001-03 to 2004-06; the bottom panel plots the average housing price contraction during the bust years of 2006 to 2009 against the change of fraction of non-owner occupied home from 2001-03 to 2004-06.

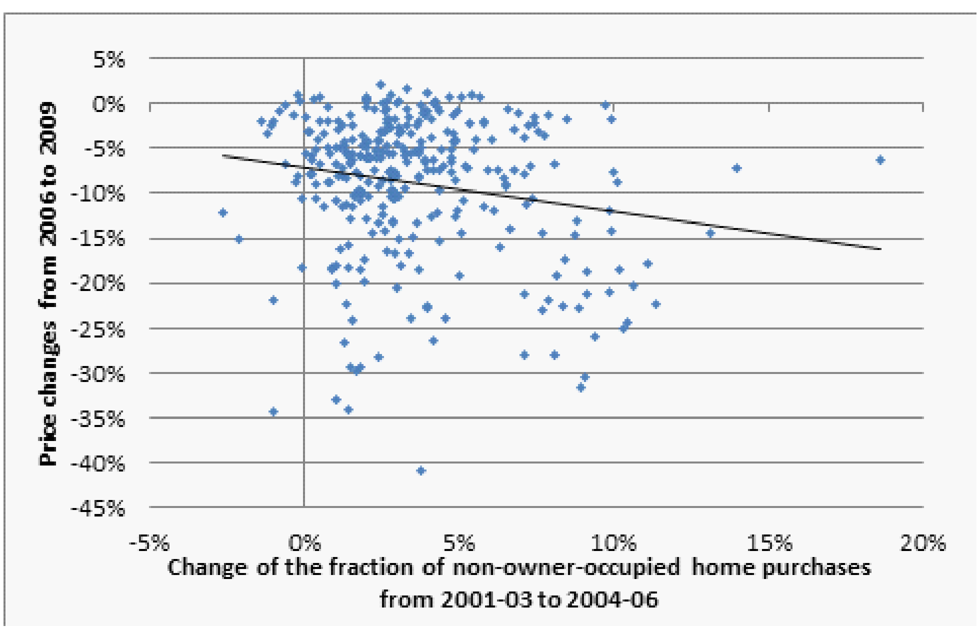

Table 1: Housing Boom and Bust during the Recent Cycle

This table presents coefficient estimates from regressing the change in real house price from 2003 to 2006 (housing boom period) and from 2006 to 2009 (housing bust period) on the dummies indicating whether a county is in the middle-elasticity group or the elastic group, with the inelastic group as the benchmark and a list of control variables. Robust standard errors are in parentheses. ***, **, * indicate coefficient estimates statistically distinct from 0 at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
	Annualized real house price change from 2003 to 2006		Annualized real house price change from 2006 to 2009	
Middle group dummy	0.0145*	0.0233***	-0.00559	-0.0387***
	(0.00871)	(0.00857)	(0.0119)	(0.0107)
Elastic group dummy	-0.0308***	-0.0173**	0.0591***	0.0125
	(0.00798)	(0.00834)	(0.0105)	(0.0102)
Fraction of subprime households in 2005		0.173***		-0.605***
		(0.0475)		(0.0636)
Annualized population change from 2003 to 2006		0.00325		
		(0.00696)		
Annualized per capita income change from 2003 to 2006		-0.0346		
		(0.0286)		
Annualized population change from 2006 to 2009				-0.00254
				(0.00954)
Annualized per capita income change from 2006 to 2009				0.0171
				(0.0431)
Constant	0.0677***	0.0297**	-0.118***	0.0150
	(0.00700)	(0.0119)	(0.00907)	(0.0153)
Observations	326	322	326	322
R-squared	0.146	0.209	0.160	0.476

Table 2: Change in the Fraction of Non-Owner-Occupied Home Purchases from 2001-03 to 2004-06

This table presents coefficient estimates from regressing the fraction of non-owner occupied home purchases in 2005 on the dummies indicating whether a county is in the middle-elasticity group or the elastic group, with the inelastic group as the benchmark and a list of control variables. Robust standard errors are in parentheses. ***, **, * indicate coefficient estimates statistically distinct from 0 at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)
	Change in the fraction of non-owner occupied home purchases from 2001-03 to 2004-06	
Middle group dummy	0.0249***	0.0242***
	(0.00424)	(0.00432)
Elastic group dummy	0.0145***	0.0124***
	(0.00348)	(0.00389)
Fraction of subprime households in 2005		-0.0318
		(0.0237)
Annualized population change from 2003 to 2006		0.00647
		(0.00444)
Annualized per capita income change from 2003 to 2006		-0.0385**
		(0.0173)
Constant	0.0190***	0.0257***
	(0.00283)	(0.00593)
Observations	323	319
R-squared	0.071	0.093

Table 3: Change in the Fraction of Non-Owner-Occupied Home Purchases from 2001-03 to 2004-06 and the Recent Housing Cycle

This table presents coefficient estimates from regressing the change in real house price from 2003 to 2006 (housing boom period) and from 2006 to 2009 (housing bust period) on the change of fraction of non-owner occupied home purchases from 2001-03 to 2004-06 and a list of control variables. Robust standard errors are in parentheses. ***, **, * indicate coefficient estimates statistically distinct from 0 at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
	Annualized real house price change from 2003 to 2006		Annualized real house price change from 2006 to 2009	
Change in the fraction of non-owner occupied home purchases from 2001-03 to 2004-06	0.636***	0.681***	-0.494***	-0.628***
	(0.115)	(0.120)	(0.165)	(0.138)
Annualized population change from 2003 to 2006		-0.00202		
		(0.00772)		
Annualized per capita income from 2003 to 2006		0.0136		
		(0.0299)		
Annualized population change from 2006 to 2009				0.00791
				(0.00991)
Annualized per capita income change from 2006 to 2009				-0.0593
				(0.0442)
Fraction of subprime households in 2005		0.241***		-0.679***
		(0.0429)		(0.0550)
Constant	0.0336***	-0.00717	-0.0704***	0.0432***
	(0.00475)	(0.00830)	(0.00693)	(0.00951)
Observations	323	319	323	319
R-squared	0.116	0.227	0.034	0.446

Bibliography

References

Aizenman, Joshua, Brian Pinto, and Vladyslav Sushko (2012), Financial Sector Ups and Downs and the Real Sector: Up by the Stairs and Down by the Parachute, mimeo, UCSC, The World Bank, and BIS.

Albagi, Elias, Christian Hellwig, and Aleh Tsyvinski (2012), A Theory of Asset Prices based on Heterogeneous Information, mimeo, USC Marshall, Toulouse School of Economics, and Yale University.

Albagi, Elias, Christian Hellwig, and Aleh Tsyvinski (2013), Dynamic Dispersed Information and the Credit Spread Puzzle, mimeo, USC Marshall, Toulouse School of Economics, and Yale University.

Albagi, Elias, Christian Hellwig, and Aleh Tsyvinski (2014), Risk-Taking, Rent-Seeking, and Investment when Financial Markets are Noisy, mimeo, USC Marshall, Toulouse School of Economics, and Yale University.

Allen, Frank, Stephen Morris, and Hyun Song Shin (2006), Beauty Contests and Iterated Expectations in Asset Markets, *The Review of Financial Studies* 19, 719-752.

Angeletos, Marios and Jennifer La'O (2009), Noisy business cycles, *NBER Macroeconomics Annual* 24, 319-378.

Angeletos, Marios and Jennifer La'O (2013), Sentiment, *Econometrica* 81, 739-780.

Angeletos, George-Marios, Guido Lorenzoni, and Alessandro Pavan (2010), Beauty contests and irrational exuberance: A neoclassical approach, mimeo, NBER.

Bacchetta, Philippe and Eric van Wincoop (2006), Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle, *American Economic Review* 96, 552-576.

Bacchetta, Philippe and Eric van Wincoop (2008), Higher Order Expectations in Asset Pricing, *Journal of Money, Banking, and Credit* 40, 837-866.

Baily, Martin Neil and Barry Bosworth (2013), The United States Economy: Why such a Weak Recovery?, Paper prepared for the Nomura Foundation's Macro Economy Research Conference, "Prospects for Growth in the World's Four Major Economies," September 11, 2013.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis (2013), Measuring Economic Policy Uncertainty, mimeo, Stanford University and University of Chicago Booth School of Business.

Bakke, Tor-Erik and Toni Whited (2010), Which Firms Follow the Market? An Analysis of Corporate Investment Decisions, *The Review of Financial Studies* 23, 1941-1980.

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny (1998), A model of investor sentiment, *Journal of Financial Economics* 49, 307-343.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer (2014), X-CAPM: An extrapolative capital asset pricing model, *Journal of Financial Economics,* forthcoming.

Barro, Robert J. (1990), The Stock Market and Investment, *Review of Financial Studies* 3, 115-131.

Barro, Robert J. and José F. Ursúa (2009), Stock Market Crashes and Depressions, mimeo, Harvard University.

Beaudry, Paul and Franck Portier (2006), Stock Prices, News, and Economic Fluctuations, *American Economic Review* 96, 1293-1307.

Beaudry, Paul and Franck Portier (2013), News Driven Business Cycles: Insights and Challenges, mimeo, Vancouver School of Economics and Toulouse School of Economics.

Bernanke, Ben (1983), Irreversibility, Uncertainty, and Cyclical Investment, *Quarterly Journal of Economics* 98, 85-106.

Bernanke, Ben (2014), The Federal Reserve: Looking Back, Looking Forward, AEA Meeting 2014.

Blanchard, Olivier J., Jean-Paul L'Huillier, and Guido Lorenzoni (2013), News, Noise, and Fluctuations: An Empirical Exploration, *American Economic Review* 103, 3045-3070.

Bloom, Nicholas (2009), The Impact of Uncertainty Shocks, *Econometrica* 77, 623-685.

Bond, Philip, Alex Edmans, and Itay Goldstein (2012), The real effects of financial markets, *Annual Review of Financial Economics* 4, 39-60.

Brown, David P. and Robert H. Jennings (1989), On Technical Analysis, *The Review of Financial Studies* 2, 527-551.

Bray, Margaret (1981), Futures trading, rational expectations, and the efficient markets hypothesis, *Econometrica* 49, 575-596.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo (2013), Understanding Booms and Busts in Housing Markets, mimeo, Duke University and Northwestern University.

Campbell, John Y. and Albert S. Kyle, (1993), Smart Money, Noise Trading and Stock Price Behavior, *Review of Economic Studies* 60, 1-34.

Campbell, John Y. and Luis M. Viceira, (2002), Strategic Asset Allocation, Cambridge University Press, Cambridge, UK.

Cao, Daniel (2011), Collateral Shortages, Asset Price and Investment Volatility with Heterogeneous Beliefs, mimeo, Georgetown University.

Case, Karl and Robert Shiller (1989), The Efficiency of the Market for Single Family Homes, *American Economic Review* 79, 125-137.

Case, Karl and Robert J. Shiller (2003), Is there a bubble in the housing market?, *Brookings Papers on Economic Activity* 2003(2): 299-362.

Chen, Qi, Itay Goldstein, and Wei Jiang (2007), Price Informativeness and Investment Sensitivity to Stock Price, *The Review of Financial Studies* 20, 619-650.

Cheng, Ing-haw and Wei Xiong (2014), The financialization of commodity markets, *Annual Review of Financial Economics,* forthcoming.

Cheng, Ing-haw, Andrei Kirilenko, and Wei Xiong (2012), Convective risk flows in commodity futures markets, mimeo, Princeton University.

Chinco, Alex and Christopher Mayer (2013), Distant speculators and asset bubbles in the housing market, mimeo, Columbia Business School.

Christiano, Lawrence J., and Terry J. Fizgerald (1998), The Business Cycle: It's Still a Puzzle, *Federal Reserve Bank of Chicago Economic Perspectives* 22, 56-83.

Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross (1985), A Theory of the Term Structure of Interest Rates, *Econometrica* 53, 385-408.

Daly, Mary C., Bart Hobijn, Asyengül Sahin, and Robert G. Valletta (2012), A Search and Matching Approach to Labor Markets: Did the Natural Rate of Unemployment Rise?, *Journal of Economic Perspectives* 26, 3-26.

Dang, Tri Vi, Gary Gorton, and Bengt Holmström (2013), Ignorance, Debt, and Financial Crises, mimeo, Columbia University, Yale University, and MIT.

Dangl, Thomas, and Michael Halling (2012), Predictive Regressions with Time-Varying Coefficients, *Journal of Financial Economics* 106, 157-181.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam (1998), Investor psychology and security market under- and over-reactions, *Journal of Finance* 53, 1839-1885.

David, Joel, Hugo A. Hopenhayn, and Venky Venkateswaran (2014), Information, Misallocation and Aggregate Productivity, mimeo, USC, UCLA, and NYU Stern.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger (2013), The Establishment-Level Behavior of Vacancies and Hiring, *Quarterly Journal of Economics,* forthcoming.

Detemple, Jerome and Shashidhar Murphy (1994), Intertemporal Asset Pricing and Heterogeneous Beliefs, *Journal of Economic Theory* 62, 294-320.

DeJong, David N. and Emilio Espino (2011), The Cyclical Behavior of Equity Turnover, *Quantitative Economics* 2, 99-133.

Diamond, Douglas W. and Robert E. Verrechia (1981), Information aggregation in a noisy rational expectations economy, *Journal of Financial Economics* 9, 221-235.

Durlauf, Steven (2004), Neighborhood Effects, *Handbook of Regional and Urban Economics* 4, 2173-2242.

Fajgelbaum, Pablo, Edouard Schaal, and Mathieu Taschereau-Dumouchel (2014), Uncertainty Traps, mimeo, UCLA, NYU, and Wharton.

Fama, Eugene F (1981), Stock Returns, Real Activity, Inflation, and Money, *American Economic Review* 71, 545–565.

Fama, Eugene F. (1990), Stock Returns, Expected Returns, and Real Activity, *Journal of Finance* 45, 1089-1108.

Fattouh, Bassam, Lutz Kilian, and Lavan Mahadeva (2012), The role of speculation in oil markets: What have we learned so far?, mimeo, University of Michigan.

Favara, Giovanni and Zheng Song (2014), House price dynamics with dispersed information, *Journal of Economic Theory* 149(1), 350-382.

FCIC (2011), The Financial Crisis Inquiry Report: Final Report of the National Commission of the Causes of the Financial and Economic Crisis of the United States.

Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, Keith Kuester, and Juan Rubio-Ramírez (2013), Fiscal Volatility and Economic Activity, mimeo, NBER.

Fujiwara, Ippei, Yasuo Hirose, and Mototsugu Shintani (2011), Can News Be a Major Source of Aggregate Fluctuations? A Bayesian DSGE Approach, *Journal of Money, Banking, and Credit* 43, 1-29.

Futia, Carl A. (1981), Rational Expectations in Stationary Linear Models, *Econometrica* 49, 171-192.

Foster, F. Douglas and S. Viswanathan (1996), Strategic Trading When Agents Forecast the Forecast of Others, *Journal of Finance* 51, 1437-1478.

Gao, Zhenyu (2013), Housing Boom and Bust with Elastic Supplies, mimeo, Princeton University.

Garbade, Kenneth and William L. Silber (1983), Price movements and price discovery in futures and cash markets, *Review of Economics and Statistics* 65, 289-297.

Garmaise, Mark and Tobias Moskowitz (2004), Confronting information asymmetries: Evidence from real estate markets, *Review of Financial Studies* 17, 405-437.

Gertler, Mark, and Cara S. Lown (1999), The Information in the High-Yield Bond Spread for the Business Cycle: Evidence and Some Implications, *Oxford Review of Economic Policy* 15, 132–150.

Gilchrist, Simon, V. Yankov, and Egon Zakrajsek (2009), Credit Market Shocks and Economic Fluctuations: Evidence from Corporate Bond and Stock Markets, *Journal of Monetary Economics* 56, 471-493.

Gilchrist, Simon and Egon Zakrajsek (2012), Credit Spreads and Business Cycle Fluctuations, *American Economic Review* 102, 1692-1720.

Glaeser, Edward (2013), A Nation of Gamblers: Real Estate Speculation and American History, *American Economic Review Papers and Proceedings* 103(3), 1-42.

Glaeser, Edward, Joseph Gyourko, and Albert Saiz (2008), Housing Supply and Housing Bubbles, *Journal of Urban Economics* 64, 198-217.

Glaeser, Edward I., Bruce Sacerdote, and José Scheinkman (2003), The Social Multiplier, *Journal of the European Economics Association* 1, 345-353.

Goldstein, Itay, Emre Ozdenoren, and Kathy Yuan (2011), Learning and complementarities in speculative attacks, *Review of Economic Studies* 78, 263-292.

Goldstein, Itay, Emre Ozdenoren and Kathy Yuan (2013), Trading frenzies and their impact on real investment, *Journal of Financial Economics* 109, 566-582.

Gorton, Gary and Guillermo Ordoñez (2012), Collateral Crises, mimeo Yale University.

Greenwald, Bruce and Joseph E. Stiglitz (1986), Externalities in Economies with Imperfect Information and Incomplete Markets, *Quarterly Journal of Economics* 101, 229-264.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell (1997), Long-Run Implications of Investment-Specific Technological Change, *American Economic Review* 87, 342-362.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell (2000), The Role of Investment-Specific Technological Change in the Business Cycle, *European Economic Review* 44, 91-115.

Greenwood, Jeremy and Boyan Jovanovic (1990), Financial Development, Growth, and the Distribution of Income, *Journal of Political Economy* 98, 1076-1107.

Grossman, Sanford and Joseph Stiglitz (1980), On the impossibility of informationally efficient markets, *American Economic Review* 70, 393-408.

Gyourko, Joseph (2009), Housing supply, *Annual Review of Economics* 1, 295-318.

Gunn, Christopher M. and Alok Johri (2013), An Expectations-Driven Interpretation of the "Great Recession", *Journal of Monetary Economics* 60, 391-407.

Hamilton, James (1983), Oil and the macroeconomy since World War II, *Journal of Political Economy* 91, 228-248.

Hamilton, James (2009), Causes and consequences of the oil shock of 2007-2008, *Brookings Papers on Economic Activity,* 215-259.

Hamilton, James and Cynthia Wu (2012), Effects of index-fund investing on commodity futures prices, mimeo, University of California, San Diego.

Hansen, Lars Peter and Thomas J. Sargent (1991), Two Difficulties in Interpreting Vector Autoregressions, in *Rational Expectations Econometrics* edited by Lars Peter Hansen and Thomas J. Sargent, Boulder CO: Westview Press, 77-119.

Hassan, Tarek and Thomas M. Mertens (2014), Information Aggregation in a DSGE Model, *NBER Macroeconomics Annual 2014*, forthcoming.

Hassan, Tarek and Thomas M. Mertens (2014), The Social Cost of Near-Rational Investment, mimeo, University of Chicago Booth School of Business and NYU Stern School of Business.

Haughwout, Andrew, Richard Peach, John Sporn, and Joseph Tracy (2012), The supply side of the housing boom and bust of the 2000s, in: *Housing and the Financial Crisis*, pages 69-104, National Bureau of Economic Research.

He, Hua and Jiang Wang (1995), Differential Information and Dynamic Behavior of Stock Trading Volume, *The Review of Financial Studies* 8, 919-972.

He, Zhiguo and Arvind Krishnamurthy (2012), Intermediary Asset Pricing, *American Economic Review* 103, 732-770.

Hellwig, Martin (1980), On the aggregation of information in competitive markets, *Journal of Economic Theory* 22, 477-498.

Hellwig, Christian, Sebastian Kohls, and Laura Veldkamp (2012), Information Choice Technologies, *American Economic Review Papers & Proceedings* 102(3), 35-40.

Henderson, Brian, Neil Pearson, and Li Wang (2012), New evidence on the financialization of commodity markets, mimeo, University of Illinois at Urbana-Champaign.

Henkel, Sam James., J. Spencer Martin, and Federico Nardari (2011), Time-Varying Short-Horizon Predictability, *Journal of Financial Economics* 99, 560–580

Himmelberg, Charles, Christopher Mayer, and Todd Sinai (2005), Assessing high house prices: Bubbles, fundamentals, and misperceptions, *Journal of Economic Perspectives* 19 (4), 67-92.

Hong, Harrison, and Jeremy Stein (1999), A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets, *Journal of Finance*, 54(6), 2143-2184.

Hu, Conghui and Wei Xiong (2013), The informational role of commodity futures prices, mimeo, Princeton University.

Ioannides, Yannis and Jeffrey E. Zabel (2003), Neighbourhood Effects and Housing Demand, *Journal of Applied Econometrics* 18, 563-584.

Judd, Kenneth (1985), The Law of Large Numbers with a Continuum of IID Random Variables, *Journal of Economic Theory* 35, 19-25.

Juvenal, Luciana and Ivan Petrella (2012), Speculation in oil market, mimeo, Federal Reserve Bank of Saint Loius.

Kilian, Lutz (2009), Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market, *American Economic Review* 99, 1053-1069.

Kilian, Lutz and Daniel Murphy (2012), The role of inventories and speculative trading in the global market for crude oil, mimeo, University of Michigan.

Kitagawa, G. (1996), Monte Carlo Filter and Smoother for Non-Gaussian Nonlinear State Space Models, *Journal of Computational and Graphical Statistics* 5, 1–25.

Knittel, Christopher and Robert Pindyck (2013), The simple economics of commodity price speculation, mimeo, MIT.

Kobayashi, K., and K. Nutahara (2007), Collateralized Capital and News-Driven Cycles, *Economics Bulletin* 5, 1-9.

Kobayashi, K., T. Nakajima, and M. Inaba (2012), Collateral Constraint And News-Driven Cycles, *Macroeconomy Dynamics* 16, 752-776.

Kogan, Leonid and Dimitris Papanikolaou (2013), Firm Characteristics and Stock Returns: The Role of Investment-Specific Shocks, *The Review of Financial Studies,* forthcoming.

Kurlat, Pablo (2013), Lemon Markets and the Transmission of Aggregate Shocks, *American Economic Review* 103, 1463-1489.

Kurlat, Pablo and Johannes Stroebel (2014), Testing for information asymmetries in real estate markets, mimeo, Stanford University and NYU.

Leduc, Sylvain, and Zheng Liu (2013), Uncertainty and the Slow Labor Market Recovery, *FRBSF Economic Letter* 2013-21.

Leeper, Eric M., Todd B. Walker, and Shu-Chun Susan Yang (2013), Fiscal Foresight and Information Flows, *Econometrica* 81, 1115-1145.

Li, Dan and Geng Li (2014), Are Household Investors Noise Traders? Evidence from Belief Dispersion and Stock Trading Volume, mimeo, Federal Reserve Board.

Lo, Andrew and Jiang Wang (2009), Stock Market Trading Volume, in *Handbook of Financial Econometrics* edited by Yacine Ait-Sahalia and Lars Peter Hansen, New York: North-Holland: Elsevier.

Lombardi, Marco J. and Ine Van Robays (2011), Do financial investors destabilize the oil price?, mimeo, ECB.

Lorenzoni, Guido (2009), A theory of demand shocks, *American Economic Review* 99, 2050-84.

Lucas, Robert. E, Jr. (1987), Models of Business Cycles, Basil Blackwell, New York.

Luo, Yuanzhi (2005), Do Insiders Learn from Outsiders? Evidence from Mergers and Acquisitions, *Journal of Finance* 60, 1951–1982.

Masters, Michael (2008), Testimony before the Committee on Homeland Security and Governmental Affairs, U.S. Senate, May 20.

Mian, Atif and Amir Sufi (2009), The consequences of mortgage credit expansion: Evidence from the U.S. mortgage default crisis, *Quarterly Journal of Economics* 124, 1449-1496.

Mian, Atif, and Amir Sufi (2012), What Explains High Unemployment? The Aggregate Demand Channel, mimeo, NBER.

Moreira, Alan and Alexi Savov (2013), The Macroeconomics of Shadow Banking, mimeo, Yale University and NYU Stern School of Business.

Morris, Stephen and Hyun Song Shin (2002), The social value of public information, *American Economic Review* 92, 1521-1534.

Moyen, Nathalie, and Stefan Platikanov, Corporate Investments and Learning, *Review of Finance* 17, 1437-1488.

Muir, Tyler (2014), Financial Crises and Risk Premia, mimeo, Yale School of Management.

Nakamura, Emi, Dmitriy Sergeyev, and Jon Steinsson (2012), Growth-Rate and Uncertainty Shocks in Consumption: Cross-Country Evidence, mimeo NBER.

Nathanson, Charles and Eric Zwick (2014), Arrested development: Theory and evidence of supply-side speculation in the housing market, mimeo, University of Chicago Booth School of Business and Kellogg School of Management.

Ng, Serena and Johnathan Wright (2013), Facts and Challenges from the Great Recession for Forecasting and Macroeconomic Modeling, *Journal of Economic Literature* 51, 1120-1154.

Nimark, Kristoffer (2012), Speculative Dynamics in the Term Structure of Interest Rates, mimeo, CREI.

Obstfeld, Maurice and Ken Rogoff (1996), *Foundations of International Macroeconomics*, MIT Press.

Ordoñez, Guillermo (2012), The Asymmetric Effects of Financial Frictions, mimeo, UPENN.

Ozdenoren, Emre and Kathy Yuan (2008), Feedback effects and asset prices, *Journal of Finance* 63, 1939-1975.

Piazzesi, Monika and Martin Schneider (2009), Momentum traders in the housing market: survey evidence and a search model, *American Economic Review Papers and Proceedings* 99(2), 406-411.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel (2007), Housing, consumption, and asset pricing, *Journal of Financial Economics*, 83, 531-569.

Pinkowitz, Lee, René M. Stulz, and Rohan Williamson (2013), Is there a U.S High Cash Holdings Puzzle after the Financial Crisis?, mimeo, Georgetown University and Ohio State University.

Philippon, T (2008), The Evolution of the US Financial Industry from 1860 to 2007: Theory and Evidence, mimeo, NYU Stern School of Business.

Poterba, James (1991), House price dynamics: the role of tax policy and demography?, *Brooking Papers on Economic Activity* 2, 143-203.

Pozsar, Zoltan, Tobias Adrian, Adam Ashcraft, and Hayley Boesky (2012), Shadow Banking, Federal Reserve Staff Report 458.

Reifschneider, Dave, William L. Wascher, David Wilcox (2013), Aggregate Supply in the United States: Recent Developments and Implications for the Conduct of Monetary Policy, mimeo Federal Reserve Board.

Reinhart, Carmen M. and Kenneth S. Rogoff (2009a), The Aftermath of Financial Crises, *American Economic Review* 99, 466-472.

Reinhart, Carmen M. and Kenneth S. Rogoff (2009b), This Time is Different: Eight Centuries of Financial Folly, Princeton and Oxford: Princeton University Press

Reinhart, Carmen M. and Kenneth S. Rogoff (2011), From Financial Crash to Debt Crisis, *American Economic Review* 101, 1676-1706.

Roll, Richard (1984), Orange juice and weather, *American Economic Review* 74, 861-880.

Saiz, Albert (2010), The Geographic Determinants of Housing Supply, *Quarterly Journal of Economics* 125(3), 1253-1296.

Sarolli, Gian Domenico (2014), Cleaning the Gears: Counter-cyclical Asset Trading with Financial Transactions Taxes, forthcoming *The Quarterly Review of Economics and Finance.*

Schwert, G. William (1990), Stock Returns and Real Economic Activity: A Century of Evidence, *Journal of Finance* 45, 1237-1257.

Singleton, Kenneth (2012), Investor flows and the 2008 boom/bust in oil prices, Working paper, Stanford University.

Sockin, Michael and Wei Xiong (2014), Information Frictions and Commodity Markets, forthcoming *Journal of Finance.*

Stoll, Hans and Robert Whaley (2010), Commodity index investing and commodity futures prices, *Journal of Applied Finance* 20, 7-46.

Straub, Ludwig and Robert Ulbricht (2013), Credit Crunches, Information Failures, and the Persistence of Pessimism, mimeo, Northwestern University.

Stock, James and Mark Watson (2003), Forecasting Output and Inflation: The Role of Asset Prices, *Journal of Economic Literature* 41, 788-829.

Stock, James and Mark Watson (2012), Disentangling the Channels of the 2007-2009 Recession, *Brookings Papers on Economic Activity* Spring 2012, 81-156.

Subrahmanyam, A and Sheridan Titman (2001), Feedback from stock prices to cash flows, *Journal of Finance* 56, 2389-2413.

Subrahmanyam, Avanidhar and Sheridan Titman (2013), Financial Market Shocks and the Macroeconomy, *The Review of Financial Studies* 26, 2687-2717.

Sun, Yeneng and Yongchao Zhang (2009), Individual Risk and Lebesgue Extension Without Aggregate Uncertainty, *Journal of Economic Theory* 144, 432-443.

Tang, Ke and Wei Xiong (2012), Index investment and financialization of commodities, *Financial Analysts Journal* 68 (6), 54-74.

Tinn, Katrin (2010), Technology Adoption and Exit in Imperfectly Informed Equity Markets, *American Economic Review* 10, 925-957.

Townsend, Robert (1983), Forecasting the Forecasts of Others, *Journal of Political Economy* 91, 546-588.

U.S. Senate Permanent Subcommittee on Investigations (2009), Excessive speculation in the wheat market, Committee on Homeland Security and Governmental Affairs, June 24.

Van Nieuwerburgh, Stijn and Laura Veldkamp (2006), Learning Asymmetries in Real Business Cycles, *Journal of Monetary Economics* 53, 753-772.

Van Nieuwerburgh, Stijn and Pierre-Olivier Weill (2010), Why Has House Price Dispersion Gone up?, *Review of Economic Studies* 77, 1567-1606.

Veldkamp, Laura (2005), Slow Boom, Sudden Crash, *Journal of Economic Theory* 124, 230–257.

Veldkamp, Laura and Justin Wolfers (2007), Aggregate Shocks or Aggregate Information? Costly Information and Business Cycle Comovement, *Journal of Monetary Economics* 54, 37-55.

Wang, Jiang (1993), A Model of Intertemporal Asset Prices Under Asymmetric Information, *Review of Economic Studies* 60, 249-282.

Wang, Jiang (1994), A Model of Competitive Stock Trading Volume, *Journal of Political Economy* 102, 127-168.

Uhlig, Harald (1996), A Law of Large Numbers for Large Economies, *Economic Theory* 8, 41-50.

Woodford, Michael (2003), Imperfect Common Knowledge and the Effects of Monetary Policy, in P.Aghion, R.Frydman, J.Stiglitz and M.Woodford, eds., Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps. Princeton, NJ: Princeton University Press, 25– 58.

Xiong, Wei and Hongjun Yan (2010), Heterogeneous Expectations and Bond Markets, *The Review of Financial Studies* 23, 1433-1466.