

Equilibrium Analysis in Behavioral One-Sector Growth Models*

Daron Acemoglu[†] and Martin Kaae Jensen[‡]

December 12, 2020

Abstract

Rich behavioral biases, mistakes and limits on rational decision-making are often thought to make equilibrium analysis much more intractable. We establish that this is not the case in the context of one-sector growth models such as Ramsey-Cass-Koopmans or Aiyagari models. We break down the response of the economy to a change in the environment or policy into two parts: the direct response at the given (pre-tax) prices, and the equilibrium response which plays out as prices change. Our main result demonstrates that under weak regularity conditions, regardless of the details of behavioral preferences, mistakes and constraints on decision-making, the long-run equilibrium will involve a greater capital-labor ratio if and only if the direct response (from the corresponding consumption-saving model) involves an increase in aggregate savings. One implication of this result is that, from a qualitative point of view, behavioral biases matter for long-run equilibrium if and only if they change the direction of the direct response. We show how to apply this result with the popular quasi-hyperbolic discounting preferences, self-control and temptation utilities and systematic misperceptions, clarifying the conditions under which usual comparative statics hold and those under which they are reversed.

Keywords: behavioral economics, comparative statics, general equilibrium, neoclassical growth.

JEL Classification: D90, D50, O41.

*We thank Xavier Gabaix for very useful discussion and comments. Thanks also to Drew Fudenberg, Marcus Hagedorn, David Laibson, Paul Milgrom and Kevin Reffett, as well as participants at the TUS-IV-2017 conference in Paris, and seminar participants at Lund University and the University of Oslo for helpful comments and suggestions. Acemoglu gratefully acknowledges financial support from the Army Research Office.

[†]Department of Economics, Massachusetts Institute of Technology and the Canadian Institute for Advanced Research.

[‡]School of Economics, University of Nottingham and School of Economics, University of Surrey.

1 Introduction

Most standard macro and growth models rely on very restrictive behavioral assumptions about households — infinitely lived, often representative, agents who are capable of solving complex maximization problems without any behavioral biases or limitations, and of implementing the optimal decisions without any inconsistencies or mistakes. It is an uncomfortable stage of introductory graduate courses when these assumptions are introduced and students rightfully ask whether everything depends on them. A natural conjecture is that these assumptions do matter and any degree of behavioral richness would render any general conclusions impossible. Not only do general equilibrium effects become notoriously complicated and the set of indirect effects correspondingly rich; we would also expect the specific departure from full rationality — *e.g.*, systematic mistakes, ambiguous beliefs, overoptimism or dynamic inconsistency — to have a first-order impact on the direction in which the economy responds to changes in policy or technology.

In this paper, we study one-sector growth models and establish that while it is true that at the individual level outcomes depend critically on the exact behavioral specification, robust predictions of long-run responses to changes in environment (policy, preferences or technology) can nonetheless be obtained in the presence of general behavioral preferences.¹ Specifically, we identify conditions that are sufficient — and when the steady-state equilibrium is unique or when changes are small, also necessary — for changes in environment to lead to comparative statics in line with the predictions of the baseline neoclassical growth models. These conditions depend only on the direction of the *direct response* to a change in environment, defined as the (partial equilibrium) impact on aggregate savings, computed from the consumption-saving problem of households, holding the pre-tax prices fixed at their initial steady-state values. Put simply, if the direct response to a change in environment is an increase in aggregate savings, then no matter how complex the general equilibrium interactions that will play out dynamically (as prices change), the long-run impact on the capital stock and output per capita will be positive. Conversely, if the direct response is a decrease in aggregate savings, then the long-run impact on the capital stock and output per capita will be negative.

Before we elaborate on this result further and provide an intuition, let us explain it in the context of a specific policy change — a reduction in the capital income tax rate. In baseline “neoclassical” settings, including the Ramsey-Cass-Koopmans model or the Aiyagari model,

¹To prove that these conclusions and intuitions hold under a broad class of specifications of mistakes and behavioral assumptions, we develop them in a general framework that nests a rich set of behavioral preferences, including those based on quasi-hyperbolic discounting, non-separable preferences, ambiguity, self-control, sparseness, systematic mistakes, under or overoptimism, and imperfect optimization.

this direct response is simply the “partial equilibrium” change in aggregate savings, holding prices at their initial steady-state values. This response is always positive under standard assumptions, ensuring that lower capital income taxes lead to higher capital-output ratio and output per capita in the long run. Taking this as a benchmark, our results can then be read as saying that *any* set of rich and more realistic behavioral preferences that do not reverse the direction of the direct response will leave the *qualitative* comparative statics of the steady-state equilibrium unchanged — the capital-labor ratio and output per capita will increase in response to lower capital income taxes.² These results apply with minimal assumptions, which in particular implies that subsets of individuals can have different types of behavioral preferences and can make various types of systematic mistakes in their expectations or optimization.


Conversely, our results also delineate robust conditions for behavioral preferences and systematic mistakes to *reverse* the direction of long-run comparative statics: When the direct response to a change in environment goes in the opposite direction of the direct response in benchmark neoclassical models, long-run (general equilibrium) comparative statics will go in the opposite direction of the conventional comparative statics — no matter how the various general equilibrium interactions play out. So if lower capital taxes reduce aggregate savings upon impact, they will lead to lower capital stock and output per capita in the long run.

Figure 1 presents these results diagrammatically. All four panels of the figure depict the key object in our analysis, “*the market correspondence*”, which summarizes the aggregate saving responses at different levels of the capital-labor ratio (see Section 2.5). Our main theorem amounts to saying that, for long-run comparative statics, it is sufficient to look at how the market correspondence shifts at the capital-labor ratio of the initial steady-state equilibrium. Panel 1 illustrates this point. Even though the market correspondence that applies for a new environment is not everywhere above the initial market correspondence, it is strictly above it at the original capital-labor ratio, and this is sufficient for us to establish that the change in environment will lead to a higher capital-labor ratio.


Panel 2 provides a complementary case. While in Panel 1 general equilibrium interactions reinforced the direct response, in this case they dampen it. In general, it is very difficult to determine, without explicit computations, whether Panel 1 or Panel 2 will apply — because general equilibrium interactions are difficult to characterize. Crucially, however, the direction of long-run comparative statics can be determined without this knowledge.

Panel 3 shows the converse case. Now the direct response is a reduction in aggregate savings. As a result, the figure shows that the long-run and output per capita will decline. Hence,


²Naturally, different distribution of preferences and mistakes across households will have *quantitative* implications. These are of course important for many applications, even though they are not our focus in the current paper.




Panel 1: Equilibrium adjustment reinforces the direct response



Panel 2: Equilibrium adjustment partially reverses the direct response



Panel 3: Reversal of the direct response implies reversal of the long-run outcome



Panel 4: Equilibrium adjustment reverses direct response (note the “downwards jumps”)

Figure 1: Panel 1 shows an instance in which general equilibrium effects amplify the direct response, while in Panel 2 they dampen it. In Panel 3 the direct response is a decline in aggregate savings, so the long-run impact incorporating general equilibrium effects is also negative. The scenario in Panel 4, where the direct response is positive and the long-run impact is negative, is impossible in the one-sector behavioral growth model because individual savings functions cannot “jump down” (equivalently, consumption functions cannot “jump up”). To overturn the (long-run) comparative statics in Panels 1-2, the direct response must be negative as in Panel 3.

if we think of Panel 1 as corresponding to the benchmark neoclassical growth model, Panel 3 would capture the case where behavioral preferences reverse the direction of the direct response, and as a result, lead to the complete opposite of the neoclassical long-run comparative statics. Finally, Panel 4 depicts the case ruled out by our theorems: appropriate, and quite weak, upper hemi-continuity assumptions preclude the possibility of downward jumps in the market correspondence.

These points are further clarified in Section 4, focusing on three types of economies featuring popular behavioral departures from neoclassical models: those with quasi-hyperbolic preferences as in Phelps and Pollak (1968), Laibson (1997) and Harris and Laibson (2001); those with self-control and temptation problems as in Gul and Pesendorfer (2004); and those with systematic misperceptions. We show that using our approach is straightforward in all three of these cases (as well as in several others presented in Appendix C). In fact, our analysis yields

general equilibrium comparative statics in this class of models that, to the best of our knowledge, do not exist in the literature. For example, with quasi-hyperbolic preferences, we show that when the extent of present-bias is limited and the intertemporal elasticity of substitution is not too high, models with quasi-hyperbolic preferences yield qualitatively similar comparative statics to those of neoclassical growth models (in other words, they can be represented by Panels 1 and 2 of Figure 1). Yet, with sufficiently high present-bias and intertemporal elasticity of substitution, the comparative statics are reversed (as in Panel 3 of Figure 1). In this case, for instance, lower capital income taxes lead to a lower capital stock and reduced output per capita in the long run. Once again, our approach enables the derivation of these general results without any explicit characterization of general equilibrium interactions — they follow directly from the inspection of the direct response to a change in environment.

To see that these findings are not driven by some strong implicit assumptions and to build an intuition for them, it is instructive to revisit the classic Aiyagari model with fully-rational (“neoclassical”) heterogeneous agents. In such an economy, the equilibrium adjustment following the direct response involves random/stochastic changes in individual asset holdings (the distribution of income), as well as prices and the aggregate capital stock as the economy settles into a new steady-state equilibrium. Even with fully-rational agents, this adjustment is complex: because of income effects, some households may change their savings in the opposite direction of the aggregate change as their income and the prices they face evolve. With behavioral preferences or biases, it is potentially even more so since we have to take into account not just the conventional income effects and price changes, but also any systematic mistakes in optimization or expectations, more complex intertemporal trade-offs and issues related to dynamic inconsistency. Nonetheless, our main theorems show that, even in such settings and exactly at the same level of generality as in the baseline Aiyagari economy, we can establish qualitative long-run comparative statics.³ However, as we emphasize in Section 4.4, though fairly general results about aggregate changes can be derived, there is a type of “indeterminacy” at the individual level — nothing much can be said about how individuals will behave and which individuals will go in the opposite direction of the aggregate economy. This observation also clarifies that our results are not a consequence of some (implicit) monotonicity assumption that ensures all households move in the same direction. On the contrary, our results are about aggregate out-

³As we explain later, a key requirement for our analysis is “time-stationarity” meaning that, facing the same problem starting at different points in time, individuals will make the same choices. As we point out in footnote 12, this does not preclude dynamic inconsistency, but rules out situation in which there are “temporary misperceptions”, which would entail beliefs changing systematically with calendar time. In the presence of such temporary misperceptions, our results would still apply but the relevant direct response may no longer correspond to the observed impact effect of a shock; in particular, in this case, the relevant direct response would need to be evaluated at the “long-run beliefs” of households (after temporary misperceptions have disappeared).

comes, without any knowledge or implication about individual adjustments.

We can now present the intuition for these results at two complementary levels. The first is economic in nature and it is related to an idea that already appears in Becker (1962) that “aggregation” disciplines economic behavior. Though we cannot say anything about individual behavior, we can determine the behavior of market-level variables (that is, aggregates such as the capital stock and income per capita). This is because even if many households respond in the opposite direction of the direct response, in equilibrium enough households have to move in the same direction as the direct response. The second intuition for our result is more mathematical. To develop this intuition, suppose that the steady-state equilibrium is unique, and focus on a policy change that increases aggregate savings at the initial capital-labor ratio. Then the only way the new steady-state equilibrium could have lower capital stock is when the equilibrium response goes in the opposite direction and more than offsets the initial increase in aggregate savings. This in turn can only be true if a higher capital stock induces lower savings. But even if this were the case, the equilibrium response could not possibly overturn the direct response. This is because the economic force leading to lower savings would not be present if the new steady-state equilibrium ended up with a lower capital stock, and thus the indirect equilibrium response would in this case reinforce rather than overturn the initial (positive) direct response. When there are multiple steady-state equilibria, this reasoning would not apply to all of them, but we develop a similar argument for extremal (greatest and least) steady-state equilibria, and under multiplicity, it is these equilibria to which our conclusions apply.

Our paper is related to several literatures. The first, already mentioned, is Becker (1962)’s seminal paper which argues that market demand curves will be downward sloping even if households are not rational because their budget constraints will put pressure for even random behavior to lead to lower demand for goods that have become more expensive. Machina (1982) makes a related type of observation about the independence axiom in expected utility theory. Though related to and inspired by these contributions, our main result is very different. While Becker’s argument is about whether an increase in price will lead to a (partial equilibrium) change in aggregate behavior consistent with “rational behavior”, our focus is about taking the initial change in behavior, whether or not it is rational, as given and then establishing that, under general conditions on the objectives and behavioral biases and constraints of households, the (general) equilibrium responses will not reverse this direct response.

The second literature we build on is robust comparative statics (*e.g.*, Topkis (1978), Vives (1990), Milgrom and Shannon (1994), Milgrom and Roberts (1994), Milgrom (1994), Quah (2007)). Not only do we share these papers’ focus on obtaining robust qualitative compara-

tive static results, but we also use similar tools, in particular a version of the “curve-shifting” arguments of Milgrom and Roberts (1994) (see also Acemoglu and Jensen (2015)) which allow us to derive robust results in non-monotone economies.⁴ Nevertheless, our main theorem is not an application of any result we are aware of; rather, it significantly extends and strengthens the approach used in the robust comparative statics literature (we provide a detailed technical discussion of the relationship with previous literature in Appendix B). Most significantly, in contrast to other approaches in the literature, our comparative static results only rely on “local information” — on behavior at a *specific* capital-labor ratio (or vector of prices) rather than the much stronger notions requiring that behavior increases or decreases savings for all prices.⁵ As a result, we are able to establish economically and mathematically stronger results: whenever the sum of the initial savings responses of all agents is positive at the initial capital-labor ratio, the full general equilibrium will involve an increase in the capital-labor ratio.

In this context, it is also useful to compare our results to those of our earlier paper, Acemoglu and Jensen (2015), where we analyzed a related setup, but with three crucial differences. First, and most importantly, there we focused on forward-looking rational households, thus eschewing any analysis of behavioral biases and their impacts on equilibrium responses. Second, and as a result of the first difference, we did not have to deal with the more general problem considered here, which requires a different mathematical approach. Third, we imposed considerably stronger assumptions to ensure that the direct response of *all* households went in the same direction at all prices, which we do not do in the current paper.

Finally, our paper is related to several recent works that incorporate rich behavioral biases and constraints into macro models. These include, among many others, Laibson (1997), Harris and Laibson (2001), Krusell and Smith (2003), Krusell, Kuruscu and Smith (2010), and Cao and Werning (2018) who study the dynamic and equilibrium implications of hyperbolic discounting (building on earlier work by Strotz (1956), and Phelps and Pollak (1968)). Particularly noteworthy in this context is Barro (1999) who shows that many of the implications of hyperbolic discounting embedded in a one-sector growth model are similar to those of standard preferences, but this is in the context of a model with a representative household and does not contain any comparative static results for this or other classes of behavioral preferences, which are our main contribution. Gul and Pesendorfer (2001, 2004) and Fudenberg and Levine (2006, 2012) de-

⁴See p.590 in Acemoglu and Jensen (2015) for additional discussion of such non-monotone equilibrium comparative statics results.

⁵See for example Lemma 1 (and Figures 1-3) in Milgrom and Roberts (1994) or Definition 5 in Acemoglu and Jensen (2015). Milgrom and Roberts (1994) also use local assumptions, but just to derive local comparative statics results (see Figure 7 and the surrounding discussion); this is different from our results, which are global despite being based on local assumptions.

velop alternative approaches to temptation and self-control and their implications for dynamic behavior. These latter two classes of models are discussed in detail in Section 4, where we show how our results can be applied to obtain new comparative statics in these settings.

Koopmans (1960), Epstein and Hynes (1983), Kreps and Porteus (1978), Lucas and Stokey (1984) and Epstein and Zin (1989, 1991) develop richer models of dynamic behavior with non-time-separable preferences, and Becker and Boyd (1997) and Backus, Routledge and Zin (2004) develop certain macroeconomic implications of such preferences. Gilboa (1987), Schmeidler (1989) and Gilboa and Schmeidler (1995) develop models of decision-making with max-min features resulting from lack of unique priors, and Hansen and Sargent (2001, 2010) and Hansen, Sargent and Tallarini (1999) discuss related preferences in various macroeconomic applications. Recent important work by Gabaix (2014, 2017) considers the macroeconomic implications of bounded rationality resulting from the inability of individuals to deal with complex problems and their need to reduce it to a sparse optimization problem, while Sims (2003) and Woodford (2013) consider the consequences of other complexity constraints on optimization. Finally, there are several examples of models featuring (systematic) mistakes and near-rational behavior including Simon (1956), Luce (1959), McFadden (1974), McKelvey and Palfrey (1995), and Train (2009). In the context of expectation formation and their implications for macroeconomics classic references include Cagan (1956), Nerlove (1958) and more recently Fuster, Herbert and Laibson (2012) and Beshears et al (2013). None of these papers develop comparative statics for macroeconomic models that apply under general behavioral preferences.

The rest of the paper is organized as follows. Section 2 describes the model and introduces the “market correspondence” (which is key to our analysis). Section 3 contains the main results and applications. Section 4 shows how our results can be applied with quasi-hyperbolic preferences, in the presence of self-control and temptation utility and with systematic misperceptions, in each case demonstrating the possibility that standard comparative static results can be reversed (for somewhat new reasons). We also establish in this section a type of indeterminacy result — while there are often unambiguous aggregate comparative statics, not much can be said about individual behavior. Section 5 concludes, while the Appendix contains an abstract discussion of our comparative statics results and the proofs of the results presented in the text.

2 Behavioral One-Sector Growth Models

This section introduces the general model. Most importantly, this allows for a broad variety of behavioral specifications of consumption-savings behavior in an (otherwise) textbook one-sector framework.

2.1 Production and Markets

The production side is the same as the canonical neoclassical growth model (e.g., Acemoglu (2009)) augmented with general distortions.

Labor is in fixed supply and normalized to unity so we can use capital, capital-labor ratio and capital-per-worker interchangeably and denote it by k . Markets clear at all times, and production is described by a profit maximizing aggregate constant returns firm with a smooth (per capita) production technology $y = f(k)$ that satisfies $f(0) = 0$, $f' > 0$, and $f'' < 0$. We also impose that there exists $\bar{k} > 0$ such that $f(k) < k$ all $k \geq \bar{k}$, which ensures compactness. This condition is implied by the standard Inada conditions when these are imposed. The rate of depreciation is $\Delta \in [0, 1]$.

We allow for taxes and distortions $\omega(k)$ and $\tau(k)$ on labor and capital. Throughout, “market prices” refer to pre-tax factor prices, $\hat{w}(k_t) \equiv f(k_t) - f'(k_t)k_t$ and $\hat{R}(k_t) \equiv f'(k_t)$. Hence, the after-tax (and after-distortion) wage and rate of return facing the households are

$$w_t = w(k_t) \equiv (1 - \omega(k_t))(f(k_t) - f'(k_t)k_t) , \quad (1)$$

and

$$R_t = R(k_t) \equiv (1 - \tau(k_t))f'(k_t) - \Delta . \quad (2)$$

The simplest example of such distortions are proportional taxes on capital and labor income, $\tau(k_t) = \tau$ and $\omega(k_t) = \omega$. Other examples include distortions from contracting frictions or markups due to imperfect competition. When $\tau(k) = \omega(k) = 0$ for all k , we recover the benchmark case with no distortions.

We allow proceeds from these distortions to be partially rebated to households (which will be the case when they represent taxes and some of the tax revenues are redistributed the households or when they result from markups that generate profits). The total amount of resources that is *not* rebated to households — that is, either consumed by the government, invested in public goods or wasted, in all cases in a way that does not affect marginal utilities — is denoted by

$$G = G(k_t) . \quad (3)$$

If nothing is rebated, then

$$G(k_t) = \omega(k_t)(f(k_t) - f'(k_t)k_t) + \tau(k_t)f'(k_t) . \quad (4)$$

On the other hand, if the only source of distortions is taxes because the government rebates everything back to consumers (e.g., in the form of lump-sum transfers), then $G(k_t) = 0$. This is the situation we focus on in our main applications in Section 4.

2.2 Households and Capital Markets

There is a continuum of households $[0, 1]$ with a typical household denoted by $i \in [0, 1]$. As in Aiyagari (1994), households are subject to borrowing constraints and any randomness is such that there is no aggregate uncertainty. Capital k_t is therefore deterministic and factor prices are given by (1) and (2) at all times. This specification nests the simpler homogenous agents, complete market models corresponding to there being no (ex-post) variation across agents and the borrowing limit being the “natural borrowing limit” (see Aiyagari (1994), p.666 or Cao and Werning (2018), p.809).

At date t each household i is influenced by a household specific shock z_t^i , where $(z_t^i)_{t=0}^\infty$ is a Markov process with invariant distribution μ_{z^i} . For concreteness, we suppose that $z_t^i = (\epsilon_t^i, l_t^i, \sigma_t^i)$ where l_t^i is the household’s labor endowment, ϵ_t^i is a random utility parameter, and σ_t^i captures any misperceptions as explained below. It is convenient to set $e_\tau^i = (w_\tau, R_\tau, T_\tau^i, z_\tau^i)$, where T_t^i denotes the transfers/rebates that household i receives at time t .⁶

At each date t , household i ’s objective is to maximize utility conditional on its beliefs (or expectations) about the future variables $(e_\tau^i)_{\tau=t+1}^\infty$ as well as its anticipated future savings behavior. Let us denote the “true model” by θ^M . This includes a complete description of all of this section’s contents, including current and future taxes, the stochastic process governing $(z_t^i)_{i \in [0,1]}$, equilibrium conditions, and so forth.

Household $i \in [0, 1]$ forms beliefs at date t on the basis of the true model θ^M and its observations of economic variables summarized in e_t^i , and this belief formation processes is summarized by the mapping $P_t^i : (w_t, R_t, T_t^i, z_t^i, \theta^M) \mapsto P_t^i(\cdot; w_t, R_t, T_t^i, z_t^i, \theta^M)$. Here beliefs, represented by $P_t^i(\cdot; w_t, R_t, T_t^i, z_t^i, \theta^M)$, define a probability measure on future outcomes; that is, for any measurable set B , the household believes that $(e_\tau^i)_{\tau=t+1}^\infty$ lies in B with probability $P_t^i(B; e_t^i, \theta^M) \in [0, 1]$.⁷ When this will cause no confusion, we omit (e_t^i, θ^M) and write $P_t^i(\cdot)$ instead of $P_t^i(\cdot; e_t^i, \theta^M)$. Rational expectations, for example, is the special case where the marginal distribution of exogenous parameters coincides with objective probabilities (as implied by the Markov process $(z_\tau^i)_{\tau=t+1}^\infty$), and the household uses the true model θ^M to correctly predict future prices. Other belief formation processes may completely ignore the true model and may generate beliefs on the basis of other variables summarized in e_t^i .⁸

⁶To simplify the notation, we have assumed that there are no shocks to capital income and to taxes/distortions. These can be incorporated into our analysis and included in z_t^i straightforwardly.

⁷See the paragraph just before Assumption 1 below for details on the measurable space to which this statement refers.

⁸A simple and familiar example is the Aiyagari model (Aiyagari (1994)), where z^i are *i.i.d.* labor endowment shocks $z^i \sim \mu_{z^i}$, and agents have rational expectations so beliefs about future prices coincide with actual (equilibrium) prices and beliefs about the future realizations of the labor endowment shock coincides with the objective probability measure, μ_{z^i} . For this reason, as in models with rational expectations more generally, beliefs can be sup-

future savings will be incorrectly perceived as well, introducing a dynamic inconsistency (Section 4.3). Moreover, since utility is random due to the parameter e_τ^i , the current self is generally uncertain, and may be wrong, and even systematically wrong, about the behavior of future selves. For example, we consider in Section 5 a household who systematically overestimates its future propensity to save. Finally, while current consumption c_t^i is fully pinned down by current income and savings a_{t+1}^i , the sequence of future consumption $(c_{t+1}^i, c_{t+2}^i, \dots)$ depends on the outcomes $(e_\tau)_{\tau=t+1}^\infty$ and is therefore *random*. Since the probabilities assigned to future outcomes are determined by the household's beliefs $P_t^i(\cdot) = P_t^i(\cdot, e_t, \theta^M)$, and these beliefs map current asset holdings a_t^i and gross savings a' to future consumption via (6), the distribution of (or beliefs about) future consumption $c^{i,t+1}$ is:

$$P_t^{c,i}(A|a', P_t^i(\cdot, e_t^i, \theta^M)) = P_t^i(\{(e_\tau^i)_{\tau=t+1}^\infty : (6) \text{ maps } (e_\tau^i)_{\tau=t+1}^\infty \text{ into } A \text{ when } a_{t+1}^i = a'\}).$$

In particular, beliefs matter both directly through expected income and indirectly because they shape future savings.

We follow Epstein and Zin (1989) and define utility directly on $(c_t^i, c_{t+1}^i, c_{t+2}^i, \dots)$. Hence given (deterministic) current consumption, $c_t^i \geq 0$, and the (random) future consumption sequence $c^{i,t+1} = (c_{t+1}^i, c_{t+2}^i, \dots)$, the household obtains utility $U^{e_t^i, i}(c_t^i, c^{i,t+1})$.

A (gross) savings level is *optimal*, denoted by $a_{t+1}^{i,*}$, if it maximizes $U^{e_t^i, i}(c_t^i, c^{i,t+1})$ subject to the previously described constraints (5)-(6), current assets a_t^i , current observations e_t^i , anticipated future savings behavior, and the current beliefs $P_t^i(\cdot) = P_t^i(\cdot, e_t^i)$. This may be expressed formally as follows:⁹

$$\begin{aligned} a_{t+1}^{i,*} \in \arg \max_{a' \in [\underline{a}^i, \bar{a}^i]} & U^{e_t^i, i}(c_t^i, c^{i,t+1}) \\ \text{subject to} & c_t^i = (1 + R_t)a_t^i + w_t l_t^i + T_t^i - a', \\ & c^{i,t+1} \sim P_t^{c,i}(\cdot | a', P_t^i(\cdot, e_t^i, \theta^M)), \text{ and} \\ & \text{given } a_t^i \text{ and } e_t^i. \end{aligned} \quad (7)$$

Given these choices, the distribution of future consumption in terms of beliefs can be written as $P_t^{c,i}(\cdot) = P_t^{c,i}(\cdot | a', P_t^i(\cdot, e_t^i, \theta^M))$.

In case $U^{e_t^i, i}(c_t^i, c^{i,t+1})$ has an expected utility representation with deterministic utility component $u^{e_t^i, i}$, (7) reduces to the more familiar formulation;

$$\begin{aligned} a_{t+1}^{i,*} \in \arg \max_{a' \in [\underline{a}^i, \bar{a}^i]} & \mathbb{E}_{P_t^{c,i}}[u^{e_{t+1}^i, i}(c_t^i, c^{i,t+1})] \\ \text{subject to} & c_t^i = (1 + R_t)a_t^i + w_t l_t^i + T_t^i - a', \\ & \text{given } a_t^i \text{ and } e_t^i. \end{aligned} \quad (8)$$

⁹To simplify the notation, we are omitting the non-negativity constraints on consumption (which are nevertheless imposed throughout).

Note however that in this formulation we are following Epstein and Zin (1989) and allowing for preferences not to have an expected utility representation. Throughout, we focus on two main cases. The first is the recursive specification studied by Epstein and Zin (1989),

$$U^{\epsilon^i, i}(c_t^i, c^{i, t+1}) = W \left(c_t^i, g^{-1} \left(\mathbb{E}_{P_t^{\epsilon^i, i}} [g(U^{\epsilon_{t+1}^i, i}(c^{i, t+1}))] \right) \right). \quad (9)$$

Here W is the aggregator and g a strictly increasing function (defining the certainty equivalent). As is well known, the Epstein-Zin specification permits risk attitudes to be disentangled from the degree of intertemporal substitutability through, for example, Kreps-Porteus preferences. In (9) when beliefs are not additive measures, the expectations operator is the Choquet integral (see Section 5 for more details on models featuring ambiguity).

The second case we focus on is the weakly additive specification:

$$U^{\epsilon^i, i}(c_t^i, \tilde{c}^{i, t+1}) = H \left(u^{\epsilon^i}(c_t^i) + \beta h \left(g^{-1} \left(\mathbb{E}_{P_t^{\epsilon^i, i}} [g(V^{\epsilon_{t+1}^i, i}(c^{i, t+1}))] \right) \right) \right). \quad (10)$$

Here u^{ϵ^i} is a strictly increasing and (weakly) concave function, $\beta > 0$ is a constant (“patience”), and h and H are strictly increasing functions.¹⁰ If $U^{\epsilon^i, i} = V^{\epsilon^i, i}$ and beliefs are correct, then (10) is dynamically consistent (in particular, recursive). However, our main interest in (10) stems from behaviors that are dynamically inconsistent, such as models of hyperbolic and more general delay discounting (Section 4.1), or versions of “sparse” optimization problems in the spirit of Gabaix (see Appendix C). The dynamic inconsistency implications of incorrect beliefs also feature in Section 4.3 in the context of recursive utility.

All sets are equipped with the Borel σ -algebra and the topology on probability measures and on random sequences is the weak convergence topology (Epstein and Zin (1989), p.940). Throughout, we impose the following very weak continuity conditions on utility and beliefs, as well as compactness on the set of random utility parameters.

Assumption 1 $U^{i, \epsilon^i}(c_0, c_1, c_2, \dots)$ is continuous in $(\epsilon^i, c_0, c_1, \dots)$, $P^i(B; w, R, z^i, \theta^M)$ is continuous in (w, R, z) for any measurable set B , and $\epsilon^i \in E^i$ where $E^i \subseteq \mathbb{R}$ is compact.

It is worth reiterating that we are greatly simplifying the description of the environment here by imposing *time-invariance*. First, this involves imposing time-invariant utility so that utility at time t , $U_t^{\epsilon^i}(c_t^i, c^{i, t+1}) = U^{\epsilon^i}(c_t^i, c^{i, t+1})$ for all t . Economically, this implies that households obtain the same (continuation) utility from the same consumption sequence starting from different points in time. Second, it involves assuming that the belief formation process is time-invariant, so that $P_t^i : (c_t^i, \theta^M) \mapsto P_t^i$, is time-invariant and can be written as

$$P_t^i = P^i \text{ for all } t = 0, 1, 2, \dots \text{ and all } i \in [0, 1]. \quad (11)$$

¹⁰Time-separable expected utility maximization is a special case of both formulations.

One justification is that the belief formation process may have already converged to a time-invariant limit starting from some initial condition.

These restrictions enable us to focus on the comparative statics of steady states — otherwise, saving decisions would not have time-invariant limits precluding the existence of non-stochastic steady states.¹¹ In particular, in the next subsection we establish that, under these assumptions, saving functions will be time-stationary, even though they will allow for dynamic inconsistency and discontinuities. Time-stationarity is crucial for our analysis, since without it there will typically be no steady states.

The time-invariance restriction is not without cost, and our results have to be applied with care in settings that are not time-invariant. For example, a policy change may create an initial period of belief confusion or mistaken perception, which becomes dissipated over time, inducing a specific type of time-dependence (Gabaix (2017)). If this is reversed in the course of the next $T < \infty$ periods, our analysis applies in principle but with some important caveats.¹²

2.3 Time-Stationary Savings Correspondences

Under time-invariance, there is no loss of generality in focusing on a household at date $t = 0$. To simplify notation, this is done in the following key definition where we also denote the observations at date 0 without subscript to further simplify (so $e_0^i = (w_0, R_0, T_0^i, z_0^i)$ is simply written as $e = (w, R, T^i, z^i)$). As a shorthand, we define an “environment”, denoted by $\theta = (\theta^M, (P^i)_{i \in [0,1]})$, to summarize the true model θ^M and beliefs $(P^i)_{i \in [0,1]}$.

Definition 1 (Time-Stationary Savings Functions and Savings Correspondences) $s_{w,R,T^i,z^i}^{\theta,i} : \mathbb{R} \rightarrow \mathbb{R}$ is a time-stationary savings function (TSSF) if for all initial levels of assets $a^i \in [\underline{a}^i, \bar{a}^i]$, for all (w, R, T^i) , almost all $z^i = (l^i, \epsilon^i, \sigma^i)$, and conditioning on these, the beliefs $P^i(\cdot; w, R, T^i, z^i, \theta^M)$ about the future random sequence $(w_t, R_t, T_t^i, z_t^i)_{t=1}^\infty$, $s_{w,R,T^i,z^i}^{\theta,i}(a^i)$ is an optimal savings level when all “future selves” adopt the same savings function:

$$s_{w,R,T^i,z^i}^{\theta,i}(a^i) \in \arg \max_{a' \in [\underline{a}^i, \bar{a}^i]: a' \leq (1+R)a^i + w l^i + T^i} U^{i,\epsilon^i}((1+R)a^i + w l^i + T^i - a', (1+R_1)a' + w_1 l_1^i + T_1^i - s_{w_1,R_1,T_1^i,z_1^i}^{\theta,i}(a'), (1+R_2)s_{w_1,R_1,T_1^i,z_1^i}^{\theta,i}(a') + w_2 l_2^i + T_2^i - s_{w_2,R_2,T_2^i,z_2^i}^{\theta,i}(s_{w_1,R_1,T_1^i,z_1^i}^{\theta,i}(a'), \dots)). \quad (12)$$

¹¹Equation (11) implies that, conditional on current observations (or initial conditions) a_t^i, w_t, R_t, T_t^i , and z_t^i , a household will face the same decision problem at date t as its future self will face at date $\tau > t$ if that future self faces the exact same initial conditions, $(a_\tau^i, w_\tau, R_\tau, T_\tau^i, z_\tau^i) = (a_t^i, w_t, R_t, T_t^i, z_t^i)$, and anticipates the same future savings behavior. In particular, the future self will then form the same beliefs about the future.

¹²This is because the relevant concept is no longer the “direct response” that takes place with the temporary beliefs, but the “hypothetical direct response” that would have obtained with the time-stationary beliefs (that apply after T periods) at the initial capital-labor ratio k^* .

The union of all time-stationary savings functions is called the time-stationary savings correspondence, $S_{e^i}^{\theta,i}(a_i) = \{s_{e^i}^{\theta,i}(a_i) : s_{e^i}^{\theta,i} \text{ is a TSSF}\}$.

This definition is quite broad and nests both dynamically consistent and dynamically inconsistent preferences, and in the latter case allows for both time-consistent and naive behavior (in particular, beliefs about future selves need not be correct). See Section 4 for explicit examples.

A correspondence is *measurable* if the inverse image of any open set is Borel-measurable (Aubin and Frankowska (1990), p.307). The savings correspondence $S_{w,R,z^i}^{\theta,i}(a^i)$ has a *compact range* when for fixed w and R , $S_{w,R,z^i}^{\theta,i}(a^i) \subseteq \bar{A}^i$, all z^i and a^i for some compact subset $\bar{A}^i \subset \mathbb{R}$ (note that \bar{A}^i may depend on w and R so it is possible for households' savings to go to infinity as prices go to 0 or infinity). In the Epstein-Zin formulation (9), supermodularity of the aggregator W has the usual meaning (e.g., Topkis (1978)) and is equivalent to assuming that consumption at different dates are Edgeworth-Pareto complements (Chipman (1977)).

The proof of the next lemma, as all other proofs in the paper, is presented in Appendix A.

Lemma 1 (Basic Properties of Savings Correspondences) *Let Assumption 1 hold and assume that utility is either given by the dynamically consistent Epstein-Zin formulation (9) where u^ϵ is increasing and concave for all ϵ , the aggregator $W(u, U)$ is concave in u , and increasing and supermodular in (u, U) , or by the weakly additive specification in (10) where H and h are strictly increasing functions and u_0^ϵ is concave for all ϵ . Then for each $i \in [0, 1]$, the (time-stationary) savings correspondence $S_{w,R,z^i}^{\theta,i}(a^i)$ exists, has a compact range, is upper hemi-continuous in w , R , T^i , and a^i , measurable in z^i , and its least and greatest selections are non-decreasing functions of assets a^i .*

The key observation is that under the general conditions of the one-sector behavioral growth model laid out above, savings correspondences must be “non-decreasing” in the standard sense of robust comparative statics (e.g., Topkis (1978), Vives (1990), Milgrom and Roberts (1994)), i.e., the least and greatest selections must be non-decreasing in assets. This is what rules out downward jumps in Figure 1 in the Introduction. A non-decreasing savings correspondence implies that the associated least and greatest consumption functions increase less than one-for-one with assets. As a result, any discontinuities must take the form of downward jumps — otherwise, there will be more than a one-for-one increase in consumption. Allowing for discontinuities is important since these are often common in the presence of dynamic inconsistencies (see e.g. Harris and Laibson (2001), p.937), and our Lemma 1 shows that, under standard assumptions, these will take the form of downward jumps.

2.4 Steady-State Equilibrium

In what follows, instead of the (time-stationary) saving correspondence $S_{w,R,T^i,z^i}^{\theta,i}(a^i)$, it is more convenient to work with the probability distribution of gross savings of a household induces by a stochastic/random savings decision \hat{a}^i . We refer to this as the *induced savings distribution* and denote it by $S_{w,R}^{\theta,i}(\hat{a}^i)$ (see the footnote for more details).¹³ Note that with this notation the transfers $(T_t^i)_{t=0}^\infty$ are subsumed into the environment θ .

The induced saving distribution captures individual randomness in a succinct way; in particular, an asset distribution \hat{a}^i is stationary (or invariant) for household i given the stationary market prices w and R and the environment θ , if and only if $\hat{a}^i \in S_{w,R}^{\theta,i}(\hat{a}^i)$.

Recall from (1)-(2) that w and R are the after-tax/distortions wage and rate of return, respectively. Hence they generally depend on the environment θ . Whenever this may cause confusion, we emphasize it by writing the market prices with the environment as a superscript. As in textbook treatments, we define steady-state equilibria directly in terms of the corresponding capital-labor ratio:

Definition 2 (Equilibrium) *The capital-labor ratio $k^* \in \mathbb{R}_+$ represents a (steady-state) equilibrium given the environment θ , if equilibrium prices $w^* = w^\theta(k^*)$ and $R^* = R^\theta(k^*)$ are given by (1) and (2), the gross savings (assets) of household i is given by (the random variable) $\hat{a}^{*,i} \in S_{w^*,R^*}^{\theta,i}(\hat{a}^{*,i})$ for almost every $i \in [0, 1]$, and the capital market clears, that is, $k^* = \int \hat{a}^{*,i} di$.*

In this definition we are implicitly assuming that the integral $\int \hat{a}^{*,i} di$ is well-defined by some version of the law of large numbers.¹⁴

2.5 The Market Correspondence

We are now ready to formally define the key theoretical innovation of this paper, namely the *market correspondence*. We will see that steady states in our model correspond to intersections of

¹³Formally, let $Q^{\theta,i}(a^i, B) = \int_{Z^i} 1_{S_{w,R,z^i}^{\theta,i}(a^i)}(B) \mu_{z^i}(dz^i)$ be the transition correspondence of savings (where μ_{z^i} denote the marginal distribution of the invariant distribution of $z_t = (z_t^i)_{i \in [0,1]}$). For a random variable \hat{a}^i on A^i with distribution η_t^i , we can now define $S_{w,R}^{\theta,i}(\hat{a}^i)$ as the set of random variables on A^i with distributions, $\eta_{t+1}^{\theta,i}(B) = \int_{a^i \in B} Q^{\theta,i}(a^i, B) \eta_t^{\theta,i}(da^i)$. Thus $S_{w,R}^{\theta,i}(\hat{a}^i)$ is the adjoint Markov correspondence (or rather, the set of random variable with distributions given by the adjoint; see the Appendix in Acemoglu and Jensen (2015) for more details).

¹⁴There is a large literature on laws of large numbers and their application in the presence of continuum of random variables as in our economy (Al-Najjar (2004), Uhlig (1996), Sun (2006)). Here and everywhere else in this paper we remain agnostic about precisely which formulation of the law of large numbers has been applied in the background. This “agnostic” approach is also the one taken in Acemoglu and Jensen (2015) where $\int a^i(k) di$ is simply assumed to equal (or be one-to-one) with a real number. This approach has the advantage of not committing to a specific interpretation and therefore comes with maximum generality. On the downside, we must be careful to not push the generality of the setting too far: In the Aiyagari model, for example, any sensible application of a law of large numbers will require that the labor endowments’ conditional distributions are at least pairwise independent conditioned on k . For further details and references, see Acemoglu and Jensen (2010, 2015)).

the market correspondence with the 45° line (Lemma 2), increases in (average) savings translates into shifts in the market correspondence (Lemma 3), and so forth.

Let $C_{w,R}^{\theta,i}(\hat{a}^i)$ denote the distribution of consumption implied by $S_{w,R}^{\theta,i}(\hat{a}^i)$ for the given environment θ .

Definition 3 (Equilibrium Asset Distributions and the Market Correspondence) *Let $C^{\theta,i}$ denote the consumption correspondence and $S^{\theta,i}$ denote the savings correspondence of household $i \in [0, 1]$. Also, let $G(k)$ denote government consumption and distortionary waste given the capital-labor ratio k .*

- A measurable mapping $\lambda : (i, k) \mapsto \lambda^i(k)$ where $\lambda^i(k)$ is a random variable on $A^i \subseteq \mathbb{R}$ is an equilibrium asset distribution, if

$$\lambda^i(k) = \frac{\hat{a}^i(k)}{\int \hat{a}^i(k) di} k, \text{ for all } (i, k) \quad (13)$$

where $(\hat{a}^i(k))_{i \in [0,1]}$ solve the fixed point problem,

$$\hat{a}^i(k) \in S_{w(k),R(k)}^{\theta,i} \left(\frac{\hat{a}^i(k)}{\int \hat{a}^i(k) di} k \right), i \in [0, 1]. \quad (14)$$

- The market correspondence $\mathcal{M}^\theta : \mathbb{R} \rightarrow 2^{\mathbb{R}}$ is

$$\mathcal{M}^\theta(k) = \{f(k) + (1 - \delta)k - G(k)\} - \left\{c \in \int C_{w(k),R(k)}^{\theta,i}(\lambda^i(k)) di : \lambda^i(k) \text{ is an equilibrium asset distribution} \right\}. \quad (15)$$

Note that the equilibrium asset holdings $\lambda^i(k)$, $i \in [0, 1]$, may be correlated across households (this will happen, for example, if households are subject to correlated shocks). But conditional on k , the definition of \mathcal{M} requires that the integral $\int C_{w(k),R(k)}^{\theta,i}(\lambda^i) di$ has a degenerate distribution ($\int C_{w(k),R(k)}^{\theta,i}(\lambda^i) di$ denotes both this distribution and its point of unit mass). With a representative household with consumption correspondence C , (13) reduces to $\lambda^i = k$ for all i , (14) becomes redundant, and (15) becomes $\mathcal{M}^\theta(k) = f(k) + (1 - \Delta)k - \int C_{w(k),R(k)}^\theta(k) di$. If, in addition, C^θ is single-valued and we suppress θ , we obtain the standard expression for aggregate savings in representative agent models:

$$\mathcal{M}(k) = f(k) + (1 - \Delta)k - G(k) - c_{w(k),R(k)}(k). \quad (16)$$

The next lemma establishes that we can work directly with the market correspondence without specifying the equilibrium asset distribution. It also confirms that fixed points of the market

correspondence will be steady-state equilibria. The proof of this lemma uses the fixed point comparative statics theorem of Acemoglu and Jensen (2015) (Theorem 4, p.601), which itself builds on Smithson’s generalized fixed point theorem as well as Richter’s theorem (Aumann (1965)). However, the most critical component of the proof is the observation that for a given k , $\mathcal{M}^\theta(k)$ equals the set of fixed points of a convex valued correspondence whose least and greatest selections are decreasing, and therefore it is itself convex-valued.

Lemma 2 (*Properties of the Market Correspondence*) *Suppose that all households satisfy the conditions of Lemma 1. Then the market correspondence \mathcal{M}^θ is a compact- and convex-valued upper hemi-continuous correspondence that begins above and ends below the 45° line. Furthermore, $k \in \mathcal{M}^\theta(k)$ if and only if k is a steady-state equilibrium.*

The market correspondence being convex-valued is an important and non-trivial property. This property does *not* simply follow from a convexification argument as in Aumann (1965), but depends critically on the fact that savings correspondences are increasing in the sense of Lemma 1 and so, in particular, on the fact that they have no jumps down. If, in fact, S^i were to have jumps down for a subset of agents of positive measure, then the correspondence $\int \mathcal{A}_k^{\theta,i}(\cdot) di$ in the proof would have jumps down as well. In that case, the market correspondence would not necessarily be convex-valued and this paper’s main result that the average direct response determines the long-run outcome would become invalid.

3 Main Results

This section contains our main results. Generalizations are provided in Appendix B and these results are applied in the context of specific behavioral models in Section 4.

Recall that θ^M denotes the “true model”, $(P^i)_{i \in [0,1]}$ denotes the households’ beliefs, and that the environment $\theta = (\theta^M, (P^i)_{i \in [0,1]})$ therefore contains all of the exogenous variables, parameters and policy variables of the model as well as specifications of how beliefs about exogenous or endogenous objects are formed. This section studies changes in the environment and the *set of possible environments* Θ is taken to be an ordered set to facilitate this perspective.

For a given environment, $\theta^* \in \Theta$, say, we know from Lemma 2 that steady-state equilibria (Definition 2) correspond to points where the market correspondence intersects with the 45°-line, *i.e.*, k^* is a steady state if and only if $k^* \in \mathcal{M}^{\theta^*}(k^*)$. This was illustrated in Figure 1 in the Introduction in the case where the market correspondence is single-valued (or we consider an appropriate selection from it).

We are now ready to define the direct responses discussed in detail in the Introduction. Before presenting our main definition (which concerns the aggregate response), it is helpful to state the more familiar individual household level definition.

Definition 4 (Individual Direct Responses) Let k^* be an equilibrium given the environment $\theta^* \in \Theta$ and denote by \hat{a}^i household i 's associated steady-state assets. Let $\theta^{**} \in \Theta$ be a different environment. Then we say that household i 's direct response is positive if its asset holdings increase at k^* when the environment changes from θ^* to θ^{**} , i.e., if

$$S_{w^{\theta^{**}}, R^{\theta^{**}}(k^*), z^i}^{\theta^{**}, i}(a^i) \geq S_{w^{\theta^*}, R^{\theta^*}(k^*), z^i}^{\theta^*, i}(a^i), \text{ a.e. } (z^i, a^i) \in Z^i \times \text{Support}(\hat{a}^i). \quad (17)$$

If the inequality is reversed, then household i 's direct response is instead negative.

Note that if savings correspondences are not single-valued, then the inequality in (17) refers to the strong set order, that is, the least and greatest optimal savings levels must increase. This convention is adopted throughout the rest of the paper.

The next definition specifies the meaning of direct responses in the aggregate.

Definition 5 (Direct Responses) Let k^* be an equilibrium given the environment $\theta^* \in \Theta$ and consider a different environment $\theta^{**} \in \Theta$. We say that the direct response is positive if the mean asset holdings of households increase at k^* when the environment changes from θ^* to θ^{**} , i.e., if $\int \hat{a}^{\theta^{**}, i}(k^*) di \geq \int \hat{a}^{\theta^*, i}(k^*) di$. If the inequality is reversed so that the mean asset holdings decrease at k^* when the environment changes from θ^* to θ^{**} , the direct response is negative.

The definition is intuitive: We average over the asset holdings (or gross savings) of households in the old and new environments holding the capital-labor ratio k^* (hence prices) fixed, and trace the direction of change. As we illustrate in Section 4, the definition makes direct reference to the associated consumption-savings model. In particular, for given k^* , the relevant asset holdings can be computed without any knowledge of (general) equilibrium changes in prices or quantities that follow from the change in environment. Clearly, if individual direct responses in Definition 4 are uniformly positive, the (aggregate) direct response in Definition 5 is positive.

Note that in both Definitions 4 and 5, (pre-tax) market prices are fixed at their initial steady-state values. For example, if the only change in environment is a change in the capital tax rate ($\theta = \tau$), then we have $w^{\theta^{**}}(k^*) = w^{\theta^*}(k^*) = f(k^*) - f'(k^*)k^*$, and $R^{\theta^*}(k^*) = (1 - \tau^*)f'(k^*) - \Delta$ and $R^{\theta^{**}}(k^*) = (1 - \tau^{**})f'(k^*) - \Delta$. So when investigating whether a change in environment leads to a positive or negative direct response, it is sufficient to consider the consumption-savings problem in steady state, with given prices. By comparison, the standard approach in

the robust comparative statics literature — including in our own work, Acemoglu and Jensen (2015) — is to impose positive direct responses in the sense of Definition 4 uniformly across all households and for all market prices (all capital-labor ratios).¹⁵ In Section 4, we illustrate how the direction of the direct response can be determined in growth models with quasi-hyperbolic preferences, self-control and temptation utilities and systematic misperceptions, and in all of these cases such results are made possible by the fact that we only need to determine the direction of the direct response, without taking into account any general equilibrium changes in prices.

We can now state the simplest version of our main result, which establishes that the long-run equilibrium outcome is pinned down by the direct response.

Theorem 1 (Main Theorem, Unique Steady State) *Assume that households satisfy the conditions of Lemma 1. For environments $\theta^*, \theta^{**} \in \Theta$ let k^* and k^{**} denote associated non-trivial steady-state equilibria and assume that these are unique. Then $k^{**} \geq k^*$ if and only if the direct response is positive when the environment changes from θ^* to θ^{**} . Similarly, $k^{**} \leq k^*$ if and only if the direct response is negative when the environment changes from θ^* to θ^{**} .*

Although uniqueness is a special case, the theorem captures this paper’s main message: In one-sector growth models, long-run outcomes are entirely pinned down by the average of the *direct responses*. Misperceptions, biases and other departures from standard, fully-rational and time-separable preferences thus impact long-run outcomes in so far as they influence household decisions at given prices. This result also implies that such departures can easily lead to “paradoxical” comparative statics (which reverse those of the standard neoclassical growth model) provided that they change the sign of the direct response. Conversely, when they do not do so, despite the very rich and potentially complex general equilibrium interactions that these behavioral preferences may spawn, they will not affect the qualitative properties of the long-run equilibrium. In the next section, we use this theorem in economies with quasi-hyperbolic preferences, self-control and temptation utilities and systematic misperceptions to investigate the direction of comparative statics with respect to changes in taxes (how our results can be applied with other classes of behavioral preferences and biases is discussed in Appendix C).

The remainder of this subsection generalizes Theorem 1 to situations with multiple equilibria and extends the discussion of the intuition and the mathematical arguments from the

¹⁵For example, in Aiyagari’s model, one can use the results in Light (2018) who shows that households will increase their savings if preferences are CRRA, the coefficient of relative risk aversion is less than one, and the rate of return increases (see his Theorem 1). In contrast, we will not impose such uniform positive or negative direct responses. Rather, our approach relies on the considerably weaker condition that at the (initial) steady-state capital-labor ratio k^* , the direct response is positive (or negative).

introductory section.

We next show that both necessity and sufficiency in our main result remain valid when there are multiple equilibria provided that we focus on the least or the greatest steady state and the exogenous changes we are considering are “small” (meaning that we can choose them to be small enough in the usual implicit function theorem sense).

Theorem 2 (*Greatest and Least Steady States under Multiplicity I*) Assume that households satisfy the conditions of Lemma 1 and let $k_-^* = \inf\{k : k \in \mathcal{M}^{\theta^*}(k)\}$ denote the least steady state and $k_+^* = \sup\{k : k \in \mathcal{M}^{\theta^*}(k)\}$ the greatest steady state when the environment is $\theta^* \in \Theta$, and analogously k_-^{**} and k_+^{**} when the environment is $\theta^{**} \in \Theta$. Assume in addition that \mathcal{M}^θ is upper hemi-continuous in $\theta \in \Theta$ (where now Θ is a topological space). Consider an infinitesimal change in the environment to θ^{**} . Then, $k_-^{**} \geq k_-^*$ if and only if the direct response is positive at k_-^* when the environment changes from θ^* to θ^{**} , and $k_+^{**} \geq k_+^*$ if and only if the direct response is positive at k_+^* when the environment changes from θ^* to θ^{**} .

If there are multiple equilibria and the change in environment is not “small” (or we are unwilling or unable to place a topology on the set of possible environments Θ), the sufficiency part of our main result will still hold for the greatest equilibrium when the direct response is positive (and for the least equilibrium when the direct response is negative):

Theorem 3 (*Greatest and Least Steady State under Multiplicity II*) Assume that households satisfy the conditions of Lemma 1 and consider $k^* = \sup\{k : k \in \mathcal{M}^{\theta^*}(k)\}$ (the greatest steady state) of the environment $\theta^* \in \Theta$. Then if the direct response is positive at k^* when the environment changes from θ^* to a new environment $\theta^{**} \in \Theta$, the economy’s greatest steady state increases, i.e., $\sup\{k : k \in \mathcal{M}^{\theta^{**}}(k)\} \geq k^*$. Analogously, consider $k^* = \inf\{k : k \in \mathcal{M}^{\theta^*}(k)\}$ (the least steady state) of the environment $\theta^* \in \Theta$. Then if the direct response is negative at k^* when the environment changes from θ^* to the new environment $\theta^{**} \in \Theta$, the economy’s least steady state decreases, i.e., $\inf\{k : k \in \mathcal{M}^{\theta^{**}}(k)\} \leq k^*$.

Appendix B contains additional results along the lines of the previous two theorems. Although important for theoretical applications, the details are less central to our substantive results, hence their relegation to the Appendix. In addition, we also provide there a detailed comparison with the related equilibrium comparative statics results in Milgrom and Roberts (1994) and Acemoglu and Jensen (2013).

The intuition for the results presented in this section was already discussed in the Introduction. Here we had elaborate their mathematical and conceptual underpinnings. Most impor-




Figure 2: A positive direct response shifts the market correspondence up at k_+^* (shown by the move from the red to the green dot) and leads to a higher steady-state capital-labor ratio (shown by the blue dot). The figure depicts a case in which there are multiple steady states both before and after the change in environment.

tantly, our approach enables us to represent any model that falls within the general one-sector behavioral growth model with a *market correspondence* \mathcal{M}^θ . From Lemmas 1 and 2, savings correspondences have no jumps down which, which implies that the market correspondence will be compact- and convex-valued, upper hemi-continuous and begin above and end below the 45° line. Crucially, a positive direct response will raise (or “shift up”) the market correspondence at the initial capital-labor ratio k^* as illustrated in Figure 2 (this is proved in the key Lemma 3 in the Appendix). As this figure also illustrates, the new equilibrium k^{**} must then be above k^* , regardless of whether the market correspondence shifts up or down at other capital-labor ratios $k \neq k^*$. This result implies that the direct response of aggregate savings at k^* pins down the direction of change for the steady-state.¹⁶

4 Applications

In this section, we illustrate how the general framework we have developed so far can be applied, focusing on three classes of models: those with quasi-hyperbolic preferences as in Phelps and Pollak (1968), Laibson (1997) and Harris and Laibson (2001); those with self-control and temptation utility as in Gul and Pesendorfer (2004); and those featuring systematic misperceptions. We emphasize, in particular, how the direct response can be characterized and how, from

¹⁶Note also that the figure illustrates the “general” case, in the sense that there are multiple steady states both before and after the change in the environment, and we focus on the largest ones corresponding to k_+^* and k_+^{**} in Theorems 2 and 3.

this analysis, we obtain simple conditions for the equilibrium comparative statics of benchmark neoclassical models to generalize or to be reversed under these behavioral considerations.

4.1 Quasi-Hyperbolic Preferences

We start with the quasi-hyperbolic preferences studied in Phelps and Pollak (1968), Laibson (1997) and Harris and Laibson (2001). Despite the popularity and the broad range of applications of these preferences, we are unaware of any analysis of general equilibrium comparative statics in this context.¹⁷ We will see that using the approach developed in this paper, the general characterization of long-run equilibrium and comparative statics is straightforward.

We embed these preferences in an Aiyagari economy: Suppose that the economy is inhabited by a continuum of households, all of which have quasi-hyperbolic CRRA preferences with discount rate δ , present-bias $\beta \leq 1$, and rate of risk aversion $\gamma > 0$. At every date, each household has labor endowment $\hat{l} + l_t$ where $\hat{l} \in \mathbb{R}_{++}$ and $l_t \sim \mu(\cdot)$ where $\mu(\cdot)$ is *i.i.d.* with mean zero and bounded support in $(-\hat{l}, +\infty)$. Borrowing is limited only by the present expected value of future labor income (see Aiyagari (1994), p.666).¹⁸

Suppose that there is no labor taxation, while capital is taxed at the rate τ as in (2). Whether tax receipts are fully rebated back to households is not central, but for specificity, we suppose that they are fully rebated in a lump-sum fashion (so that $G(k_t) = 0$ for all t in (3)). We assume full depreciation for notational simplicity and denote the pre-tax rental rate of capital by \hat{R}_t , so that the after-tax rate of return is $R_t = (1 - \tau)\hat{R}_t - 1$, and the lump-sum transfer is $T_t = \tau\hat{R}_t k_t$. Let us define the “environment” as $\theta = 1 - \tau \in \Theta = [0, 1]$, highlighting that our main comparative statics will be with respect to the capital income tax rate τ .

We now investigate whether and when an increase in θ (a lower tax on capital) increases mean savings at given prices. This involves studying the corresponding consumption-savings model where we fix w and \hat{R} at their initial steady-state values. The generalized Euler equation of the consumption-savings model is

$$u'(C_\theta(z_t)) = \theta\hat{R}\delta \int (1 + (\beta - 1)D_Y C_\theta(z_{t+1})) u'(C_\theta(z_{t+1})) \mu(dl_{t+1}), \quad (18)$$

where $z_t = \theta\hat{R}a_t + w[\hat{l} + l_t] + T_t$, $z_{t+1} = \theta\hat{R}[z_t - C_\theta(z_t)] + w[\hat{l} + l_{t+1}] + T_{t+1}$, a_t is assets, and C_θ the consumption function. Note that in terms of the paper’s general notation, $s_{w,R,\hat{l}+l_t}^{i,\theta}(a_t^i) = \theta\hat{R}a_t^i + w[\hat{l} + l_t] + T_t - C_\theta(\theta\hat{R}a_t^i + w[\hat{l} + l_t] + T_t)$, and because households are ex-ante identical, we omit the i

¹⁷The exception is for the deterministic logarithmic utility case, which is observationally equivalent to neoclassical growth models with standard preferences, as noted in Barro (1999) and Krusell, Kuruscu and Smith (2002).

¹⁸This ensures that the borrowing constraints will not bind in equilibrium, simplifying the analysis and enabling us to directly apply the implicit function theorem.

superscripts. As in Harris and Laibson (2001) or Aiyagari (1994), a partial equilibrium stationary distribution is any joint distribution μ^* over assets and endowments that is an invariant measure for the stochastic difference equation:

$$a_{t+1} = s_{w,R,\hat{l}+l_t}^\theta(a_t), l_t \sim \mu. \quad (19)$$

To apply our main results all we need is to determine whether aggregate savings are increasing or decreasing in θ at given (pre-tax) factor prices. We can further simplify the analysis of this question and clarify the economics by focusing on the case where the random component of labor income originating from $\mu(\cdot)$ is “small”, which implies that precautionary savings are small.¹⁹ In this case, saving functions become almost linear, and the generalized Euler equation simplifies to

$$((1 - D_Y C_\theta(z))\hat{R})^\gamma = \theta \hat{R} \delta (1 + (\beta - 1) D_Y C_\theta(z)) \Omega, \quad (20)$$

where z is a level of wealth (inclusive of discounted future labor income) in the steady-state support set from above, and $\Omega \approx 1$.²⁰ Applying the implicit function theorem,²¹ we immediately obtain the key result that the marginal propensity to save $1 - D_Y C_\theta(z)$ of a household with total wealth z is increasing in θ *if and only if*

$$(1 - D_Y C_\theta(z))^{-1} > \frac{1 - \beta}{\beta} \frac{1 - \gamma}{\gamma}. \quad (21)$$

Note that one advantage of this special case with negligible precautionary savings is that when (21) holds, the consumption function of all agents will shift in the same direction as the change in environment θ (since $D_Y C_\theta(z)$ is approximately constant for z in the support of the steady-state asset distribution μ^*).

Notably, if this inequality is reversed, then the the marginal propensity to save is decreasing in θ (for all steady-state wealth levels) and thus households’ direct response to a tax reduction is *negative*. When the steady-state capital-labor ratio is unique, we can apply Theorem 1 and conclude (proof in the text):²²

¹⁹Formally, we require that the distance between $\mu(\cdot)$ and the degenerate measure with unit mass on 0 is sufficiently small in the Levy-Prokhorov metric.

²⁰Here $\Omega = \int ((\hat{R}\hat{y})/(\hat{R}\hat{y} + w l_{t+1} + \tau \hat{R} \cdot (\tilde{n}^* - \hat{y})))^\gamma \mu(dl_{t+1})$, where \tilde{n}^* is steady-state capital inclusive of future discounted labor income, $\hat{y} = \hat{R}[z - C_\theta(z)]$ and z is such that $\hat{y} = \tilde{n}^*$. Since l_{t+1} has mean zero and is close to 0, it is clear that $\Omega \approx 1$.

²¹In this application of the implicit function theorem, we keep z fixed and treat $\bar{s} = 1 - D_Y C_\theta(z)$ as the unknown variable. We thus get $\gamma \hat{R}^\gamma \bar{s}^{\gamma-1} d\bar{s} = \Omega \left([\hat{R}\delta(1 + (\beta - 1)[1 - \bar{s})] d\theta - [\theta \hat{R}\delta(\beta - 1)] d\bar{s} \right)$. Because (20) holds, we can divide through with this equation and rearrange to see that $d\bar{s}/d\theta > 0 \Leftrightarrow (21)$.

Note also that when there are nontrivial precautionary savings, we would need to apply an infinite-dimensional implicit function theorem as in Yano (1989). We adopted the assumption that shocks to labor endowments are small to attain this simplification.

²²Without uniqueness, we could instead use Theorem 2.

Proposition 1 *In the behavioral growth model with quasi-hyperbolic preferences, a reduction in the capital income tax rate increases the steady-state capital-labor ratio if (21) holds and reduces it if (21) holds with the inequality reversed.*

Several points are worth noting. First, the above argument (hence the proposition) heavily relies on our main theorems, which establish that even with rich general equilibrium interactions, the impact of changes in environment on long-run equilibrium will be in the same direction as the direct response at the steady-state prices. The equilibrium adjustment triggered by the changes in prices and the asset distribution following the direct response involves a complex interplay between income, wealth and substitution effects and quasi-hyperbolic preferences. Nonetheless, our main result ensures that these dynamic equilibrium adjustments cannot overturn the direction of the direct response.

Second, the proposition shows that with quasi-hyperbolic discounting (and limited precautionary motives), a simple condition, (21), is necessary and sufficient for a reduction in the tax rate on capital income to increase the long-run capital stock as in standard neoclassical growth models. This further implies that when this condition is not satisfied, in contrast to standard neoclassical models, a reduction in the tax rate on capital income reduces the long-run capital stock.

Third, the proposition and condition (21) also clarify when this paradoxical result happens. Specifically, because the left-hand side of (21) is bounded from above by $1 + \beta^{-1}((\theta\delta)^{-1} - 1)$, $\gamma < \theta(1-\beta)\delta$ is sufficient to ensure that (21) holds with the inequality reversed.²³ Thus, provided that the initial tax on capital income is not too high ($\theta^{-1} = 1/(1-\tau)$ not too high), households are sufficiency present-biased (β low enough), and the intertemporal elasticity of substitution is high relative to this (γ is low enough), (21) will be violated, the direct response to a reduction in the capital income tax rate will be towards lower mean savings, and consequently, the steady-state capital stock will decrease.

Conversely, when $\beta = 1$ (so that preferences are dynamically consistent) or when $\gamma \geq 1$ (including in the case where preferences are logarithmic as in Barro (1999) and Krusell, Kuruscu and Smith (2002)), (21) is always satisfied and thus we have the conventional results that lower capital taxes increase the long-run capital stock.

Fourth, condition (21) also gives us an intuition about why lower capital taxes can reduce long-run capital-labor ratio. The generalized Euler equation (18) shows that a higher marginal utility in the future will be associated with a higher marginal utility today — hence, more sav-

²³The upper bound follows because in steady state we must have $\hat{R}\bar{s}(\theta, R) < 1$ and $\theta\hat{R}\delta(1 + (\beta - 1)D_Y C_\theta(Y_{t+1})) \approx 1$

ings in the future will go together with more savings today. This linkage will be particularly strong when γ is low (high intertemporal elasticity) and $\theta\delta(1 - \beta)$ is high, thus being amplified by the reduction in the capital income tax rate (a higher θ). In particular, when $\gamma < \theta(1 - \beta)\delta$, current consumption becomes so sensitive to future decisions that greater future savings induced by lower taxes will lead to even more savings today. But this cannot be sustained in steady state, implying that steady-state savings will have to decrease.

Finally, we should remark that there are various ways in which Proposition 1 can be extended. The conclusion is *a fortiori* true when we allow for more systematic mistakes, such as incorrect expectations about what a change in policy means and what it implies for future transfers and rebates. Similar results also apply when we introduce different groups of consumers with different types of preferences (for example, a subset of the consumers having quasi-hyperbolic preferences, while the rest have standard preferences or utility functions from another class). Lastly, the analysis of the model in the presence of non-CRRA preferences and with substantial precautionary savings because of future uncertainty is similar as well, except that the direct response to changes in policy is no longer characterized by a single generalized Euler equation (independent of future income as in (20)). In this case, it is not typically possible to have an explicit necessary and sufficient condition as in (21),²⁴ but the general analysis remains similar and as in Proposition 1, a reduction in the capital income tax rate may increase or reduce the long-run capital stock.

4.2 Self-Control and Temptation

We now present a similar analysis for the dynamic self-control and temptation preferences introduced in Gul and Pesendorfer (2004). Household i has utility function u^i , temptation cost v^i , discount factor δ^i , and (stationary) labor endowment l^i , at every date (the labor supply is thus $L = \int l^i di$). A borrowing constraint prevents households from holding negative gross assets (this is useful for linking temptation to current wealth as opposed to the borrowed discounted value of the entire future labor income stream).

As in the previous section, capital is taxed at the rate τ , generating a tax revenue of $\tau \hat{R}k_t$ at date t which is rebated back to the households through lump-sum transfers, $T_t = \tau \hat{R}k_t$. The relevant environment is again represented by $\theta = 1 - \tau$.

As discussed in Gul and Pesendorfer (2004) (see in particular their Section 6), dynamic self-control and temptation preferences yield unique optimal payoffs and therefore well-defined recursive dynamic programs. The model satisfies the conditions of Lemma 1, provided that

²⁴An exception, which still allows an explicit necessary and sufficient condition, is when preferences are CRRA and all uncertainty comes from variation in future rates of returns.

overall utility, $u^i(c) + v^i(c)$, is concave, increasing and continuous. To simplify we are also going to assume that u^i is strictly concave and that u^i and v^i are twice continuously differentiable on \mathbb{R}_+ . In addition, we assume that v^i is either strictly convex or strictly concave with positive third derivative. As in Gul and Pesendorfer (2004) (p.119), the consumption-savings model can be summarized by the Bellman equation:

$$W^i(z_t^i) = \max_{0 \leq y_t^i \leq z_t^i} u^i(z_t^i - y_t^i) + v^i(z_t^i - y_t^i) + \delta^i W^i((1 + R)y_t^i + wl^i) - v^i(z_t^i). \quad (22)$$

Here $z_t^i = (1 + R)a_t^i + wl^i$ is the household's current wealth, y_t^i its current savings, w and R are steady-state market prices given the initial capital tax rate $\tau = 1 - \theta$, and W^i is the value function (which is unique and strictly concave).

Let $S_\theta^i(z_t^i)$ denote the (unique) maximizer on the right-hand side of (22) so that, in terms of our notation $s_{w,R,l^i}^{i,\theta}(a_t^i) = S_\theta^i((1 + R)a_t^i + wl^i)$. Denote household i 's steady-state wealth and consumption by $z^{i,*} = \theta \hat{R} a^{i,*} + wl^i$ and $c^{i,*} = (\theta \hat{R} - 1)a^{i,*} + wl^i$, respectively. Note that these are uniquely determined because consumption stationarity implies $\delta^i \theta \hat{R} > 1$ (Gul and Pesendorfer (2004), p.137).

We now show how we can use Definition 5 to obtain comparative statics in this general equilibrium model with self-control and temptation utility. First, we derive the Euler equation from (22). Keeping the initial prices and capital-labor ratio fixed, we obtain the partial equilibrium responses $\partial \hat{a}^{\theta^*,i} / \partial \theta$ to a reduction in capital taxes (an increase in $\theta = 1 - \tau$). The average of these is positive, *i.e.*, $\partial \int \hat{a}^{\theta^*,i}(k^*) di / \partial \theta > 0$, where k^* is the initial steady-state capital-labor ratio if and only if:²⁵

$$\int \frac{\frac{1}{\delta^i \theta^2 \hat{R} - \theta} + (a^{i,*} - k^*) \hat{R} \left[\frac{u''^{i,i}(c^{i,*}) + v''^{i,i}(c^{i,*})}{u'^{i,i}(c^{i,*}) + v'^{i,i}(c^{i,*})} - \frac{v''^{i,i}(z^{i,*})}{v'^{i,i}(z^{i,*})} \right]}{\frac{u''^{i,i}(c^{i,*}) + v''^{i,i}(c^{i,*})}{u'^{i,i}(c^{i,*}) + v'^{i,i}(c^{i,*})} - \theta \hat{R} \left[\frac{u''^{i,i}(c^{i,*}) + v''^{i,i}(c^{i,*})}{u'^{i,i}(c^{i,*}) + v'^{i,i}(c^{i,*})} - \frac{v''^{i,i}(z^{i,*})}{v'^{i,i}(z^{i,*})} \right]} di > 0. \quad (23)$$

We therefore conclude (proof in the text):

Proposition 2 *Consider the behavioral growth model with dynamic self-control preferences. Then a reduction in the capital income tax rate increases the steady-state capital-labor ratio if (23) holds and reduces it if (23) holds with the inequality reversed.*

As in the previous subsection, this result heavily exploits our main comparative static theorems. It shows how basic neoclassical comparative statics generalize to a model with dynamic

²⁵Omitting the household index, the Euler equation is $-u'(z_t - S_\theta(z_t)) - v'(z_t - S_\theta(z_t)) + \delta \theta \hat{R} \{u'(\theta \hat{R} S_\theta(z_t) + wl + T_t - S_\theta(\theta \hat{R} S_\theta(z_t) + wl + T_t)) + v'(\theta \hat{R} S_\theta(z_t) + wl + T_t - S_\theta(\theta \hat{R} S_\theta(z_t) + wl + T_t)) - v'(\theta \hat{R} S_\theta(z_t) + wl + T_t)\} \leq 0$, with equality whenever $S_\theta(z) > 0$. To obtain the household's response, evaluate in the initial steady state where $T_t = (1 - \theta) \hat{R} k^*$, $a^* = S_\theta(z^*)$ and the borrowing constraint does not bind. Then apply the implicit function theorem.

self-control preferences under a range of conditions, but it also highlights how these comparative static results can be reversed because of self-control considerations.

Though the condition in (23) is straightforward to compute and easy to understand, its intuition becomes clearer when we abstract from the redistributive consequences of the tax system. For this purpose, let us consider the special case of our model where households are homogeneous ($u^i = u$, $v^i = v$ and $\delta^i = \delta$) and always have the same labor endowment, $l^i = l$, so that $a^{i,*} = k^*$ in the initial steady state. In this case, the second term in the numerator of (23) is equal to zero, and (23) becomes equivalent to

$$\left(\theta\hat{R} - 1\right) \frac{u''(c^*) + v''(c^*)}{u'(c^*) + v'(c^*)} < \theta\hat{R} \frac{v''(z^*)}{v'(z^*)}, \quad (24)$$

Since $\theta\hat{R} > \delta\theta\hat{R} > 1$ and $u + v$ is strictly concave, it is clear that when households' temptation utilities are convex, (24) holds and a lower capital income tax rate always increases the capital-labor ratio in steady state. Conversely, suppose that v is strictly concave. Now, if $\frac{v''(z^*)}{v'(z^*)} / \frac{u''(c^*) + v''(c^*)}{u'(c^*) + v'(c^*)}$ is greater than $1 - \delta$, (24) holds with the inequality reversed, and the direct response to a reduction in the capital income tax rate will be negative, leading to a lower capital-labor ratio in the long run. This comparative static reversal does not require far-fetched conditions. For example, $u = v$ and these functions exhibiting either a constant or an increasing absolute rate of risk aversion ensure such a reversal.

We can also provide an economic intuition for this paradoxical result. With dynamic self-control and temptation preferences and v concave, in addition to the usual consumption smoothing motive, households have an incentive to smooth their wealth — a smoother wealth profile leads to lower temptation costs. If this wealth smoothing motive is sufficiently strong relative to the consumption smoothing motive, a reduction in the capital income tax rate makes it optimal to reduce savings today in order to achieve this smoother wealth profile. This explanation also highlights that, though Gul and Pesendorfer (2004) focus on the cases where v is convex, there is nothing in the economics that precludes concave v . But this type of concavity can create powerful behavioral effects that can reverse the direction of general equilibrium comparative statics. Once again, we are not aware of any results in the literature of this type, which are enabled by our general results.

4.3 Systematic Misperceptions

Our final application is to a heterogeneous household growth model with systematic misperceptions. The description of markets and taxes is as in the quasi-hyperbolic model. To simplify,

we assume here that households are homogenous with labor endowment l , which they supply inelastically, so that $L = l$. As in baseline neoclassical models, there is geometric discounting, no non-separabilities and no self-control problems. The only behavioral element is that households can have misperceptions, and we simplify the discussion here by assuming that these only concern future transfers. Specifically, we assume that instead of the actual future transfers T_{t+1}, T_{t+2}, \dots , households believe that they will receive transfers $\alpha T_{t+1}, \alpha T_{t+2}, \dots$ where $\alpha > 0$. This misperception is systematic in the sense that households never correct their mistaken beliefs. We impose no binding borrowing constraints on households, so when $\alpha = 1$, we are back to the standard neoclassical growth model. On the other hand, when $\alpha < 1$, households expect to receive a lower transfer than what is actually the case, and they will underestimate the impact of a change in capital taxes on their future incomes.

Clearly, this modified one-sector growth model satisfies the conditions of Lemma 1, so, once more, dynamic general equilibrium effects cannot overturn direct responses, and we can again focus on the (partial equilibrium) consumption-savings version of the model.

Fixing an initial level of the tax $\tau \in (0, 1)$, a household's Euler equation is:

$$\frac{u'(C_\theta(z_t))}{u'(C_\theta(z_{t+1}^E))} = (1 - \tau)\hat{R}\delta. \quad (25)$$

Here $z_t = (1 + R)a_t + wl + T_t$ denotes *actual* income at date t and $z_{t+1}^E = (1 + R)[z_t - C_\theta(z_t)] + wl + T_{t+1}^E$ is *expected* income at the following date given the *expected* transfers $T_{t+1}^E = \alpha T_{t+1}$. Note that the key feature of this model is that whenever $\alpha \neq 1$, z_{t+1}^E will be different from actual income at date $t + 1$, z_{t+1} . As a consequence, planned consumption $C_\theta(z_{t+1}^E)$ will differ from actual consumption $C_\theta(z_{t+1})$ whenever beliefs are incorrect.²⁶ Then, using (25), we have that $\theta\hat{R}\delta \neq 1$ in steady state whenever $\alpha \neq 1$, ensuring that, as in the previous subsection, each households' savings is uniquely determined by equilibrium prices and its current assets.

Applying the implicit function theorem, we have:

$$\frac{dS_\tau(z^*)}{d\tau} = \frac{(1 - \tau)\hat{R}\delta u''(c^{E,*}) \left((\alpha - 1)\hat{R}a^* - S'_\tau(z^{E,*})(\alpha - 1)\hat{R}a^* \right) - \hat{R}\delta u'(c^{E,*})}{-u''(c^*) - (1 - \tau)\hat{R}\delta u''(c^{E,*}) \left[(1 + (\alpha - 1)\tau)\hat{R}(1 - S'_\tau(z^{E,*})) - 1 - o(|z^{E,*} - z^*|) \right]},$$

where $z^{E,*}$ and $c^{E,*}$ denote expected income and consumption at the next date in steady state, a^* is households' steady-state assets and $\lim_{x \rightarrow 0} o(x) = 0$. Now supposing that τ is not too high and α is not too far from 1 (so that misperceptions are not too large), we can conclude that household i 's direct response to an increase in the capital income tax rate is negative if and only

²⁶This introduces a dynamic inconsistency into the model. The current description parallels "naive behavior" in the sense of Strotz (1956), except that here the utility objective itself does not cause the dynamic inconsistency.

if:

$$(1 - \tau) \frac{u''(c^{E,*})}{u'(c^{E,*})} (\alpha - 1) \hat{R}a^* (1 - S'_\tau(z^{E,*})) < 1. \quad (26)$$

Using our main results we thus have (proof in the text):

Proposition 3 *Consider the behavioral growth model with misperceptions and assume that the initial tax is not too high and misperceptions are not too strong. Then a reduction in the capital income tax rate increases the steady-state capital-labor ratio if (26) holds and reduces it if (26) holds with the inequality reversed.*

Just as in the previous two models, we again see how basic neoclassical comparative statics generalize to alternative behavioral specifications, and how they can be reversed. Specifically, note that (26) necessarily holds if $\alpha \geq 1$ (since the left-hand side is then negative) — in particular, we unambiguously obtain the usual comparative statics if beliefs are either correct ($\alpha = 1$) or the households overestimate their future transfer payments ($\alpha > 1$). But if $\alpha < 1$ and the intertemporal elasticity of substitution is low enough (or equivalently, the absolute rate of risk aversion is high enough), this inequality is reversed and a lower tax on capital income leads to a lower capital-labor ratio in the long run.

Intuitively, when $\alpha > 1$, households overestimate the decline in future income caused by the tax reduction — hence the tax cut makes them “feel relatively poorer” in the future, motivating them to shift income towards future dates. Since the usual substitution effect pushes in the same direction, the conventional comparative statics conclusion remains intact independently of how strong (or weak) the consumption smoothing motive is. On the other hand, if $\alpha < 1$, households underestimate the declines in their future incomes resulting from a lower capital income tax. As a result, following such a tax reduction, they feel relatively richer in future dates and reduce their savings and increase their consumption today in response, especially if the consumption smoothing motive is sufficiently strong (or the intertemporal elasticity of substitution is sufficiently low).

4.4 Indeterminacy

In this subsection, we show that, although our framework leads to sharp results on how aggregate variables change in response to certain changes in the environment (for example, lower taxes on capital income leading to greater capital-labor ratios), nothing much can be said about individual behavior. Specifically, we will establish a type of indeterminacy result, proving that it is in general impossible to know which households will increase their savings and which ones

will reduce them. Put differently, fixing a set of households with given preferences, we can always find the production function and preferences for the rest of the households such that this initial set will reduce their savings while aggregate savings increase.

For transparency, we will state this “indeterminacy” result for the Aiyagari model with intertemporally separable neoclassical preferences, which guarantees that a reduction in the capital income tax rate increases the long-run capital-labor ratio. Clearly, indeterminacy in this model implies, *a fortiori*, that there will also be indeterminacy with more complex behavioral preferences.

Let $\eta(\cdot)$ denote the Lebesgue measure on the set of households $[0, 1]$ so that $\eta(J)$ is the mass of a (measurable) subset of households $J \subseteq [0, 1]$. Let us also refer to the ergodic average of the savings of a household in a steady state simply as its *ergodic savings*. Then, we have:

Proposition 4 *Consider the Aiyagari model and suppose that all households have CRRA preferences with some coefficient of risk aversion less than one (potentially different across households). Consider a reduction in the capital income tax rate. Fix a set of households $J \subseteq [0, 1]$ with $\eta(J) \leq B$ for some $B > 0$ sufficiently small. Then there exist a production function f and utility functions for the remaining households such that a lower capital income tax will lead to lower ergodic savings for all households in J , while the steady-state capital-labor ratio and aggregate savings increase.*

In sum, though aggregate savings will necessarily increase, we can say nothing about how the savings of a fixed set of households will change between the current and the new steady states. They may increase or decrease depending on the fine details of the utility functions of the rest of the households. In fact, Proposition 4 shows that the savings of any (albeit small) subset of households may *uniformly* decrease, while aggregate savings increase. This indeterminacy finding reiterates that our main results are not driven by some hidden monotonicity assumptions — they are truly a consequence of the discipline that this class of models, despite rich behavioral preferences, imposes on aggregate variables, while placing little or no restrictions on individual behavior.

We also note that the same indeterminacy result holds with the other behavioral models discussed in this paper — in particular with quasi-hyperbolic and self-control and temptation preferences. An increase in the capital-labor ratio of the economy does not pin down which households will increase their savings, and any given household can increase or reduce its savings depending on the production function and the preferences of other households in the economy.

5 Concluding Remarks and Future Directions

A common conjecture is that equilibrium analysis becomes excessively challenging in the presence of behavioral preferences and biases, thus implicitly justifying a focus on models with time-additive, dynamically consistent preferences and rational expectations. In this paper, we demonstrated that, in the context of one-sector behavioral growth models, this conjecture is not necessarily correct. Results concerning the direction of change in the long run (or “robust comparative statics” for the steady-state equilibrium) can be obtained for a wide range of behavioral preferences and rich heterogeneity. Put simply, our main results state the following: for any change in policy or underlying production or preference parameters of the model, we first determine whether, at the initial capital-labor ratio (or at the initial pre-tax/distortion vector of prices), aggregate savings increases or decreases; this step involves no equilibrium analysis, but only the determination of what the average of individual optimization decisions given prices is. Critically, this needs to be done only at a single vector of prices (or at a single capital-labor ratio), because our condition is completely “local”. Then under fairly mild regularity conditions (satisfied for all behavioral preferences we have discussed in this paper), no matter how complex the equilibrium responses are, they will not overturn the direction of the initial change and thus the steady-state equilibrium will involve a greater capital-labor ratio (and the changes in prices that this brings). Conversely, if the initial change is a decline in aggregate savings at the initial capital-labor ratio, the long-run capital-labor ratio will decline.

At the root of this result is a simple and intuitive observation: in the one-sector model, the only way the direction of the impact of the initial impetus can be reversed is by having the equilibrium response to this initial shock to go strongly in the opposite direction. For example, savings could decline strongly in response to a higher capital-labor ratio. But either such an equilibrium response would still not overturn the initial increase in aggregate savings, in which case the conclusion about the steady-state equilibrium applies. Or it would overturn it and reduce the long-run capital-labor ratio, but in this case the perverse effect would go in the direction of strengthening, not reversing, the initial increase in savings.

We illustrated these comparative statics by working through one-sector growth models embedding three popular behavioral effects: quasi-hyperbolic preferences, self-control and temptation utilities and systematic mistakes. In all three cases, we showed that our approach can be applied relatively straightforwardly and leads to results that are, to the best of our knowledge, new in the literature. Moreover, we found that, for a broad range of parameters, comparative statics in these behavioral models are similar to those in standard neoclassical growth models. However, importantly, we also identified economically intuitive conditions under which these

comparative statics are reversed. In each case, this reversal takes place along the lines of our main result: behavioral preferences change the direction of the direct response, and, despite potentially complex general equilibrium interactions, this initial impetus then leads to a change in the same direction in the long-run equilibrium.

This intuition also clarifies the limitations of our results. A similar logic would not apply if the economy had multiple state variables rather than the single state variable as in our (one-sector) behavioral growth model. In such richer circumstances, similar results would necessitate at least some supermodularity conditions for the set of state variables or a result that in the relevant problems the vector of state variables could be reduced to be functions of a single overall state variable. One example in which this latter approach can be used straightforwardly is an extension of our setup to a multi-sector neoclassical growth model. For brevity, we did not develop the details of this model, but the main idea is simple. Suppose that we have a n -sector growth model with no irreversibilities, neoclassical production functions in each sector and competitive capital markets (though distortions that differ across sectors can be introduced for additional generality). Then the marginal return to capital has to be equalized across different sectors, which determines an allocation of the overall capital stock across sectors and enables us to have a reduced-form problem just as a function of the overall capital stock. Then similar comparative static results can be developed for this overall capital stock in this type of multi-sector environment. Beyond this case, extending our results to other settings with multiple state variables is far from trivial, and would typically necessitate strong supermodularity/monotonicity conditions (in contrast, our current results require no such monotonicity assumptions).

Another evident limitation of our approach bears repeating at this point: our focus has been on comparative statics, and thus on qualitative rather than quantitative results. Many questions in modern macroeconomics necessitate quantitative analysis, and the quantitative impact of a policy change may critically depend on behavioral biases and the exact structure of preferences even if the direction of long-run change does not. An obvious but challenging area for future research is to investigate when certain quantitative conclusions may not depend on appropriately introduced behavioral biases or heterogeneity (for example, in the sense that as behavioral assumptions are modified, quantitative change in some key variables remains near changes implied by a benchmark model).

Appendix A: Proofs

Proof of Lemma 1. We begin with the weakly additive specification (10). Throughout, the superscript θ as well as the household index i are omitted to simplify notation. The monotonic transformation H may be ignored and (12) written as

$$s_{w,R,T,z}(a) \in \arg \max_{a': a' \leq (1+R+\sigma)a + wl + T} u_0^\epsilon((1+R+\sigma)a + wl + T - a') + M(a').$$

Here M is a function that generally depends on w, R, T, z , and a' but it does not depend on current assets a . Under Assumption 1, M is continuous function (in particular, beliefs vary continuously with w, R, T , and z). Since $u_0^\epsilon((1+R+\sigma)a + wl - a') + M(a')$ is supermodular in (a, a') if and only if u_0^ϵ is concave, it follows from Topkis' theorem (Topkis (1978)) that the least and greatest optimal savings functions are non-decreasing in current assets a .

If $S_{w,R,T,z}(a)$ is upper hemi-continuous in z , it is measurable in z (Aubin and Frankowska (1990), Proposition 8.2.1). To establish both the upper hemi-continuity and measurability requirements of Lemma 1, it therefore suffices to show that under Assumption 1, $S_{w,R,T,z}(a)$ is upper hemi-continuous in w, R, T, a , and z . The proof is the same in each case and in fact, the statement is true if we consider (w, R, T, a, z) jointly. Nonetheless, to simplify notation we establish the claim only for a . We begin with the case where u_0 is strictly concave. Let $a_n \rightarrow a$, and $b_n = s_{w,R,T,z}(a_n) \in S_{w,R,T,z}(a_n)$ for all n where $s_{w,R,T,z}$ is a time-stationary savings function (TSSF). Without loss of generality, index again by n a subsequence with $b_n \rightarrow b$. We first show that $b \in S_{w,R,T,z}(a)$. By definition,

$$b_n \in \arg \max_{a': (1+R+\sigma)a_n + wl + T \geq a'} u_0^\epsilon((1+R+\sigma)a_n + wl + T - a') + M(a')$$

Under Assumption 1, M is continuous from below (respectively, above) if and only if $s_{w,R,T,z}(\cdot)$ is continuous from below (respectively, above). In particular, M is continuous if and only if $s_{w,R,T,z}(\cdot)$ is continuous. For any a , we can pass to yet another subsequence (again indexed by n) such that the convergence $a_n \rightarrow a$ is monotone. Since u_0 is strictly concave, $s_{w,R,T,z}$ is increasing, and it follows then that $b_n \rightarrow b$ monotonically. In case (a_n) (and therefore (b_n)) is an increasing sequence, the conclusion that

$$b \in \arg \max_{a': (1+R+\sigma)a + wl \geq a'} u_0^\epsilon((1+R+\sigma)a + wl - a') + M(a')$$

follows by a standard continuity argument provided that $s_{w,R,T,z}$, and therefore M is continuous from below. In the second case of decreasing (a_n) and (b_n) the conclusion follows if $s_{w,R,T,z}$ is continuous from above. Crucially, it may be shown that if $s_{w,R,T,z}$ is a TSSF, then

so is both its lower continuous and its upper continuous closures.²⁷ Further, since an increasing function is continuous except for at an at most countable number of points, $s_{w,R,T,z}(a)$ coincides with its lower and upper continuous closures nearly everywhere (as a minimum, at all points of continuity). Because of this we may from the beginning of the argument above replace $s_{w,R,T,z}$ with, as appropriate, the lower or upper continuous closure without having to change the sequences (a_n) and (b_n) . But then the previous argument may be applied to conclude that $b \in S_{w,R,T,z}(a)$. If there are only a finite number of TSSFs, this argument implies upper hemi-continuity of $S_{w,R,T,z}(a)$ (since then for any sequence (a_n) and any sequence (b_n) with $b_n \in S_{w,R,T,z}(a_n)$ all n , there exist convergent subsequences with $b_n = s_{w,R,T,z}(a_n)$ all n for some fixed TSSF). If $S_{w,R,T,z}$ is the union of an infinite family of TSSFs, we instead use that if s^n is a sequence of TSSFs, then its pointwise limit is also a TSSF. Finally, to extend the proof from the case where u_0 is strictly concave to the case where u_0 is merely assumed to be concave, one uses a standard approximation argument: consider a sequence of strictly concave functions u_0^n that converge pointwise to u_0 ; repeat the above argument for all n ; use that continuity of the maximum operator implies that the limit is optimal for u_0 . That $S_{w,R,T,z}$ has a compact range follows immediately from upper hemi-continuity and boundedness of the set of feasible savings levels.

Next, consider the recursive Epstein-Zin specification (9). It is convenient to adopt Epstein and Zin's notation and denote the certainty equivalent by μ_P where P is the associated measure, and also, it is convenient to suppress ϵ and the income transfer T . Let $V(a, e) = \max_{a'} W(u((1+R)a+wl-a'), \mu_P[V(a', e')])$ denote the value function where $P(\cdot) = P(\cdot, e, \theta^M)$ are beliefs formed conditional on e and θ^M . It is clear that the value function is weakly increasing in a under the conditions of the Lemma and, furthermore, V is a continuous function under Assumption 1. From continuity of the value function follows that the savings correspondence is upper hemi-continuous (and hence have a compact range and be measurable in z as in the weakly additive specification considered previously). It remains therefore only to be shown that the savings correspondence is increasing in a . This follows by the same argument as in the previous weakly additive case if we can show that $W(u((1+R)a+wl-a'), \mu_P[V(a', e')])$ is supermodular in a and a' . Since the objective is concave in a , it is differentiable almost everywhere in a and when the derivative exists it equals: $(1+R)W'_1(u((1+R)a+wl-a'), \mu_P[V(a', e')]) \cdot u'((1+R)a+wl-a')$.

²⁷Note that since we consider a continuum of consumers, isolated points on the graph of an individual's savings correspondence are irrelevant and this paper's construction of savings correspondences adds these onto the graph (but note that every proof in this paper would remain valid had we instead considered increasing selections from the savings correspondences). Let $s_{w,R,T,z}$ be an increasing function. Then the lower continuous closure is defined by $\underline{s}_{w,R,T,z}(a) = \lim_{a^n \uparrow a} s_{w,R,T,z}(a^n)$. The upper continuous closure is defined similarly, replacing $a^n \uparrow a$ with $a^n \downarrow a$. These are always well-defined for increasing functions since an increasing function is continuous except at an at most countable number of points (and the points of discontinuity are of the jump type).

By Theorem 4 in Jensen (2007), it is sufficient for increasing differences/supermodularity in a and a' that this term is increasing in a' between any two points where it is well-defined. Since u is concave, $u'((1+R)a + wl - a')$ is increasing in a' . Since $W_1'(u, U)$ is decreasing in u , and increasing in U , and $u((1+R)a + wl - a')$ is decreasing in a' and $\mu_P[V(a', e')]$ is increasing in a' , $W_1'(u((1+R)a + wl - a'), \mu_P[V(a', e')])$ is increasing in a' . The conclusion now follows because the product of two increasing functions is an increasing function. ■

Proof of Lemma 2. First use that $f(k) + (1 - \delta)k - G(k)$ equals aggregate income after taxes and net of any waste to write the market correspondence in terms of the savings correspondences:

$$\begin{aligned} \mathcal{M}^\theta(k) &= f(k) + (1 - \delta)k - G(k) - \int ((1 + R(k))a^i + l^i w(k) - S_{w(k), R(k)}^{\theta, i}(\frac{\hat{a}^i(k)}{\int \hat{a}^{\theta, i}(k) di} k) di \\ &= \int S_{w(k), R(k)}^{\theta, i}(\frac{\hat{a}^i(k)}{\int \hat{a}^i(k) di} k) di . \end{aligned}$$

By definition of the equilibrium asset distributions, these are solutions to the fixed point problem in (14), and therefore the market correspondence can also be written as:

$$\mathcal{M}^\theta(k) = \left\{ \int \hat{a}^i(k) di : \hat{a}^i(k) \in S_{w(k), R(k)}^{\theta, i}(\frac{\hat{a}^i(k)}{\int \hat{a}^i(k) di} k), \text{ a.e. } i \in [0, 1] \right\} . \quad (27)$$

From here it is seen that $k \in \mathcal{M}^\theta(k)$ if and only if there exists $(\hat{a}^i(k))$ which satisfies (14) and such that $k = \int \hat{a}^i(k) di$. Substitute this into (14) to see that $\hat{a}^i(k) \in S_{w(k), R(k)}^{\theta, i}(\hat{a}^i(k))$ which means that $\hat{a}^i(k)$ is an invariant distribution for household i . Comparing with Definition 2, we conclude that $k \in \mathcal{M}^\theta(k)$ if and only if k is an equilibrium.

For $K > 0$, let $\mathcal{A}_k^i(K) \equiv \{\hat{a}^i \in \mathcal{P}(\bar{A}^i) : \hat{a}^i \in S_{w(k), R(k)}^{\theta, i}(\hat{a}^i \frac{k}{K})\}$ where $\mathcal{P}(\bar{A}^i)$ is the set of probability measures on the compact range $\bar{A}^i \subseteq \mathbb{R}$ of the savings correspondence equipped with the weak *-topology. By Lemma 1, $S_{w, R, z^i}^{\theta, i}(a^i)$ is increasing and upper hemi-continuous in a^i , hence the induced savings distributions (the adjoint Markov correspondence) $S_{w, R}^{\theta, i}(\hat{a}^i)$ is type I and type II monotone and upper hemi-continuous in \hat{a}^i (see the Appendix in Acemoglu and Jensen (2015)). It follows from the fixed point comparative statics Theorem 3 in Acemoglu and Jensen (2015) that $\mathcal{A}_k^i(K)$ is type I and type II monotone in K^{-1} . By Theorem 4 in that same paper, $\int \mathcal{A}_k^i(\cdot) di$ has decreasing least and greatest selections. Since $\int \mathcal{A}_k^i(\cdot) di$ is convex valued by Richter's theorem (see Aumann (1965)), and a convex and real-valued correspondence whose least and greatest selections are decreasing must have a convex set of fixed points, $\mathcal{M}^\theta(k) = \{K : K \in \int \mathcal{A}_k^{\theta, i}(K) di\}$ is therefore convex. Note that the fact that the least and greatest savings functions have no jumps down is the critical property for the previous argument (that the savings correspondence is increasing matters only because it ensures this).

To see that the market correspondence $\mathcal{M}^\theta(k) = \{K : K \in \int \mathcal{A}_k^{\theta,i}(K) di\}$ is upper hemi-continuous, note that its graph is $\{(k, K) : (K, k, K) \in \text{Graph}[\int \mathcal{A}_k^{\theta,i}(K) di]\}$ where $\text{Graph}[\int \mathcal{A}_k^{\theta,i}(K) di] = \{(K, k, Z) : Z \in \int \mathcal{A}_k^{\theta,i}(K) di\}$ is a closed set since $\int \mathcal{A}_k^{\theta,i}(K) di$ is upper hemi-continuous in k and K . That $\mathcal{M}^\theta(k)$ is compact follows now from boundedness (savings correspondences have compact ranges). Finally, $\mathcal{M}^\theta(k)$ begins above the 45° line and ends below it. The former is obvious since $f(0) = 0$ and therefore $\mathcal{M}^\theta(0) = \{0\}$. The latter is true since consumption is non-negative, hence $\mathcal{M}^\theta(k) \leq f(k)$, and for sufficiently large k , $f(k) \leq k$ because the production technology is effectively compact. ■

Proof of Theorem 1. The proof relies on the following lemma.

Lemma 3 (Mean Asset Holdings and Shifts in the Market Correspondence) *Assume that households satisfy the conditions of Lemma 1, and let $k^* \in \mathcal{M}^{\theta^*}(k^*)$ be either the least steady state $\inf\{k : k \in \mathcal{M}^{\theta^*}(k)\}$ or the greatest steady state $\sup\{k : k \in \mathcal{M}^{\theta^*}(k)\}$ given an environment $\theta^* \in \Theta$. Consider a different environment $\theta^{**} \in \Theta$. Then the population's mean asset holdings increase at k^* when the environment changes from θ^* to θ^{**} if and only if the market correspondence “shifts up” at k^* (i.e., provided there exists $\tilde{k} \in \mathcal{M}^{\theta^{**}}(k^*)$ with $\tilde{k} \geq k^*$). Similarly, the population's mean asset holdings increase at k^* when the environment changes from θ^* to θ^{**} if and only if the market correspondence “shifts down” at k^* (there exists $\tilde{k} \in \mathcal{M}^{\theta^{**}}(k^*)$ with $\tilde{k} \leq k^*$).*

Proof. From the proof of Lemma 2 we know that $\mathcal{M}^\theta(k) = \{K : K \in F_k(K, \theta)\}$ where

$$F_k(K, \theta) = \left\{ \int a^i di : a^i \in S_{w(k), R(k)}^{\theta,i} \left(\frac{a^i}{K} k \right), a.e. i \right\}.$$

$F_{k^*}(\cdot, \theta^*)$ and $F_{k^*}(\cdot, \theta^{**})$ are upper hemi-continuous, convex valued, and necessarily begin above and end below the diagonal since they are decreasing correspondences (see the proof of Lemma 2). For clarity, we first consider the case where households' assets distributions are uniquely determined in steady state (in principle, different assets distributions might support the same steady state).

Let k^* be the greatest equilibrium. If the population's mean asset holdings increase at k^* , $F_{k^*}(k^*, \theta^{**}) \geq k^*$, then there exists $K \in F_{k^*}(K, \theta^{**})$ with $K \geq k^*$. Since $K \in \mathcal{M}^{\theta^{**}}(k^*) \Leftrightarrow K \in F_{k^*}(K, \theta^{**})$, it follows that the market correspondence shifts up at k^* . This argument also applies if k^* is the least equilibrium since F is decreasing in K .

To see that an increase in mean asset holdings is also necessary for the market correspondence to shift up, use that if there does not exist $\hat{k} \in F_{k^*}(k^*, \theta^{**})$ with $\hat{k} \geq k^*$, then because $F_k(K, \theta^{**})$ is convex valued with least and greatest selections that are decreasing in K , there

is not a $K \in F_{k^*}(K, \theta^{**})$ with $K \geq k^*$, and so the population's mean asset holdings does not increase as θ^* changes to θ^{**} .

If households' steady-state asset distributions are not uniquely determined from k , one considers instead the greatest mean asset holdings: $A_+^{\theta,i}(k) = \sup\{\mathbb{E}[\hat{a}^i] : \hat{a}^i \in S_{w(k),R(k)}^{\theta,i}(\hat{a}^i)\}$, and define the greatest average asset holdings across the agents (given θ and the steady state k): $A_+^\theta(k) = \int A_+^{\theta,i}(k) di$. Then the change in environment from θ^* to θ^{**} shifts the market correspondence up at k^* if and only if $k^* \leq A_+^{\theta^{**}}(k^*)$ (see the proof of Proposition 3). Note that trivially the left-hand side of this inequality, k^* , is the average asset holding across the households at the steady state k^* . So, the necessary and sufficient condition is that the greatest average asset holding after the change in environment is above the average asset holdings before the change. ■

We are now ready to prove Theorem 1. Only the case where the population's mean asset holdings increase is considered (the case where the population's mean asset holding decrease is proved by an analogous argument).

Sufficiency: Since the change to θ^{**} shifts the market correspondence up at k^* , there exists $\tilde{k} \in \mathcal{M}^{\theta^{**}}(k^*)$ with $\tilde{k} \geq k^*$. Since $\mathcal{M}^{\theta^{**}}$ ends below the diagonal, it must begin above and end below the 45° line on the interval $[k^*, +\infty)$. $\mathcal{M}^{\theta^{**}}$ is also upper hemi-continuous and convex valued (Lemma 2), hence it intersects the 45 degree line at some point k^{**} on $[k^*, +\infty)$. This yields a steady-state equilibrium $k^{**} \geq k^*$ given environment θ^{**} , and by assumption, this is the unique steady-state equilibrium. Parenthetically, note that the same conclusion follows by instead considering a single-valued market correspondence that is continuous but for jumps up (see Appendix B).

Necessity: Assume that $k^{**} \geq k^*$ and that the change from θ^* to θ^{**} does not increase the households' mean asset holdings. By Lemma 3, the market correspondence then does not shift up at k^* . So $\sup \mathcal{M}^{\theta^{**}}(k^*) < k^*$ since the market correspondence is closed. But then since the market correspondence ends below the 45° line and is upper hemi-continuous and convex valued, $\mathcal{M}^{\theta^{**}}$ must intersect with the 45° at least twice on the interval $[k^*, +\infty)$. This contradicts that the economy has a unique interior steady state given θ^{**} . ■

Proof of Theorem 2. Since the market correspondence is compact-valued, a sufficiently small change in the environment can lead to existing equilibria disappearing but not to the creation of new equilibria. In particular, no new equilibrium can be created below the least equilibrium which must therefore increase by the argument used to prove Theorem 1. This argument obviously also applies to the greatest equilibrium; and in both cases necessity follows by the

argument from Theorem 1 as well. ■

Proof of Theorem 3. Let k^* denote the greatest steady state. Repeating the argument used to prove the “sufficiency” part of Theorem 1, $\mathcal{M}^{\theta^{**}}$ must have a fixed point on $[k^*, +\infty)$. The result for the least steady-state is proved analogously. ■

Proof of Proposition 4. Denote that pre-tax rental rate of capital by \hat{R} , the rate of capital tax by $\tau \in [0, 1)$ and set $\theta = 1 - \tau$. Since the coefficients of relative risk aversion (RRAs) are below 1, the individual savings functions are increasing in the after-tax rate of return $R = \theta\hat{R} - \Delta$ and decreasing in the (after-tax) wage w , and the steady-state equilibrium is unique (Light (2018)). From Lemma 1, the savings functions are also non-decreasing in assets. Therefore, ergodic savings are increasing in R and decreasing in w . Furthermore, all individual direct responses to a lower capital income tax are positive, and thus Theorem 1 implies that the long-run capital-labor ratio increases following such a tax reduction. Consequently, the steady-state market interest rate (*i.e.*, \hat{R}) declines, and the wage increases.




Figure 3: Blue curves show the demand for capital from the production side, $k = (f')^{-1}((R + \Delta)/\theta)$, before (solid) and after (dashed) the reduction in the capital income tax rate from $\tau^* = 1 - \theta^*$ to $\tau^{**} = 1 - \theta^{**}$, $\tau^* > \tau^{**}$. The red curves show the supply of savings from the household side, $k = s(R, w((R + \Delta)/\theta))$, where $w((R + \Delta)/\theta) = f((f')^{-1}((R + \Delta)/\theta)) - f'((f')^{-1}((R + \Delta)/\theta))(f')^{-1}((R + \Delta)/\theta)$. The figure depicts an economy where a large fraction of households have coefficients of relative risk aversion close to 0, so the supply of savings is very elastic.

Now fix the preferences of the households in J . We first note that, regardless of their exact RRAs, the ergodic savings of each household in this set declines if the reduction in the market interest rate is small enough for a given increase in the wage rate w (because, in this case, the income effect dominates the substitution effect). We next prove that we can choose the prefer-

ences of the households in $[0, 1] \setminus J$ and the production technology f so that the reduction in the market interest rate is arbitrarily small and the increase in the wage rate is large. Since $\eta(J)$ is small, we can ensure that the equilibrium of this economy is arbitrarily close to the equilibrium of a hypothetical economy without the household in the set J . Now choose the RRAs of all households in the set $[0, 1] \setminus J$ to be arbitrarily close to 0. This ensures that the savings curve in the Aiyagari diagram depicted in Figure 3 is arbitrarily close to a vertical line. Now consider an increase in θ (a reduction in the capital income tax rate). This leads to an arbitrarily small shift in the savings curve (shown in red), while for any production function where $f''(\cdot) < 0$, the demand for capital from the production side shifts up. This ensures that the wage rate, w , increases, while the change in the after-tax rate of return R is arbitrarily small. This establishes the desired result. ■

Appendix B: Changes in the Environment: A Topological Approach, Discussion of Related Literature

Since this section's observations may be of independent interest and apply not only to market correspondences, we are going to view the market correspondence $\mathcal{M} : K \times \Theta \rightarrow 2^{\mathbb{R}}$, $K \subseteq \mathbb{R}$, more abstractly and impose any necessary assumptions directly. Denote by $m_S^\theta(k) = \inf \mathcal{M}^\theta(k)$ and $m_L^\theta(k) = \sup \mathcal{M}^\theta(k)$ the least and greatest selections, and by $k_S^\theta = \inf\{k \in K : k \in \mathcal{M}^\theta(k)\}$ and $k_L^\theta = \sup\{k \in K : k \in \mathcal{M}^\theta(k)\}$ the least and greatest fixed points (when they exist, which of course they do if \mathcal{M} is a market correspondence). Now equip Θ with an order as well as a topology (in the simplest situation where we consider a change in just a single parameter, Θ may be taken to be a subset of \mathbb{R} , and these would therefore be the usual/Euclidean order and topology, respectively). We also introduce some additional terminology: A function $m : \Theta \rightarrow \mathbb{R}$ is (i) *increasing* if $\theta \leq \hat{\theta} \Rightarrow m(\theta) \leq m(\hat{\theta})$ for all $\theta, \hat{\theta} \in \Theta$, and (ii) *locally increasing at $\theta^* \in \Theta$* if $\theta \leq \hat{\theta} \Rightarrow m(\theta) \leq m(\hat{\theta})$ for all $\theta, \hat{\theta}$ in an open neighborhood of θ^* . Finally, \mathcal{M} *begins above and ends below the 45° line* if $m_*(\inf K, \theta) \geq \inf K$ and $m^*(\sup K, \theta) \leq \sup K$.

Theorem 4 (Abstract Shifts in Fixed Point Correspondences) *Consider an upper hemi-continuous and convex valued correspondence $\mathcal{M} : K \times \Theta \rightarrow 2^{\mathbb{R}}$ where K is a compact subset of \mathbb{R} and Θ is a compact subset of an ordered topological space. Suppose that the graph begins above and ends below the 45° line for all $\theta \in \Theta$. Then the least and greatest fixed points k_S^θ and k_L^θ are increasing in θ if for all $\theta^* \in \Theta$, $m_L^\theta(k_L^{\theta^*})$ and $m_S^\theta(k_S^{\theta^*})$ are locally increasing in θ at θ^* .*

Proof. Consider the greatest fixed point $k_L^{\theta^*}$ given some $\theta^* \in \Theta$. To simplify notation, we take $\Theta \subseteq \mathbb{R}$ (but the argument is true in general). Since $m_L^\theta(k_L^{\theta^*}) \geq m_L^{\theta^*}(k_L^{\theta^*}) = k_L^{\theta^*}$ for $\theta^* + \epsilon >$

$\theta > \theta^*$, $m_L^\theta(\cdot)$ begins above the 45° line and ends below it on the interval $[k_L^{\theta^*}, \sup K]$. Since \mathcal{M} has convex values, $\mathcal{M}^\theta(\cdot)$ therefore has a fixed point on this interval, and so $k_L^\theta \geq k_L^{\theta^*}$. This argument clearly extends to any $\theta > \theta^*$ (not necessarily in a neighborhood) since we may reach any such θ in a finite number of steps (Θ is compact so any open cover contains a finite subcover). The more difficult case is when $\theta^* - \epsilon < \theta < \theta^*$. Assume for a contradiction that $k_L^\theta > k_L^{\theta^*}$. Consider θ_n , where $\theta < \theta_n < \theta^*$. Since $\theta_n > \theta$, it follows from the first part of the proof that $k_L^{\theta_n} \geq k_L^\theta > k_L^{\theta^*}$. Note that these inequalities hold for any $\theta_n \in (\theta, \theta^*)$. Since K is compact, we may consider a sequence $n = 0, 1, 2, \dots$ with $\theta_n \uparrow \theta^*$ and such that $\lim_{n \rightarrow \infty} k^*(\theta_n)$ exists. $k_L^{\theta_n} \in \mathcal{M}^{\theta_n}(k_L^{\theta_n})$ for all n and \mathcal{M} has a closed graph, hence $\lim_{n \rightarrow \infty} k_L^{\theta_n} \in \mathcal{M}^{\theta^*}(\lim_{n \rightarrow \infty} k_L^{\theta_n})$. But since $\lim_{n \rightarrow \infty} \lim_{n \rightarrow \infty} k_L^{\theta_n} \geq k_L^\theta > k_L^{\theta^*}$, this contradicts that $k_L^{\theta^*}$ is the greatest fixed point. The parallel statement for the least fixed point $k_S^{\theta^*}$ is shown by a dual argument (in this case the situation where $\theta^* - \epsilon < \theta < \theta^*$ is simple while the limit sequence argument must be used for the case where $\theta^* + \epsilon > \theta > \theta^*$). ■

The following corollary is immediate by combining Theorem 4 with Lemma 3, as we did in the proof of Theorem 1:

Corollary 1 (Main Comparative Statics Result, Topological Case) *Let the assumptions of Theorem 1 hold and assume in addition that Θ is a compact subset of an ordered topological space and that the market correspondence $\mathcal{M}^\theta(k)$ is upper hemi-continuous in (θ, k) . Then the greatest and least steady states k_S^θ and k_L^θ are increasing in θ if for all $\theta^* \in \Theta$ and all $\theta_a < \theta_b$ in a neighborhood of θ^* , the change in the environment from θ_a to θ_b raises mean savings at $k_L^{\theta^*}$ as well as at $k_S^{\theta^*}$.*

Note that in all cases, “curve shifting theorems” such as Theorem 4 can be used in our setting because (i) Lemma 2 has established the requisite properties of the market correspondence; and (ii) Lemma 3 allows us to relate increases in mean savings/assets with “shifts up” in the market correspondence. In that sense, curve shifting arguments are the last (and simplest) step in the proofs of our results; in particular, (i) and (ii) are clearly the more substantial technical contributions of this paper.²⁸ Nonetheless, it is useful to briefly contrast our use of curve shifting arguments to similar results in the literature because there are both significant economic and mathematical differences.

Most of the results in the literature are similar to Corollary 2 in Milgrom and Roberts (1994) which shows that when the equivalent of our market correspondence \mathcal{M} is “continuous but for

²⁸As mentioned several times in the main text, the key technical contribution is the definition of the market correspondence in terms of an auxiliary fixed point problem which allows us to conduct the analysis without taking asset distributions (explicitly) into account and ensures that the market correspondence is convex-valued.

jumps up” and its graph shifts up (meaning that $m_L^\theta(k)$ and $m_S^\theta(k)$ are increasing in θ for all k), then the least and the greatest fixed points increase.²⁹ Let us refer to this well-known result as the “for all k curve shifting theorem”. The key thing to note is that since the “curve” must shift up for all k (for all capital-labor ratios in our setting), it requires information not only about how savings change for the prices determined in the original steady state; it requires that we have such information for (all) capital-labor ratios/prices. Acemoglu and Jensen (2015), which relies on such a standard curve shifting theorem, defines “local positive shocks” as changes in parameters that increase savings for all capital-labor ratios. In conventional settings with rational expectations, such requirements can be imposed, even if they are quite demanding. When the economic problems involve rich and variegated behavioral preferences and biases, they become essentially untenable. It is against this background that Theorem 4 should be evaluated. It shows that if \mathcal{M} is upper hemi-continuous in (k, θ) (rather than just in k , cfr. footnote 29), the same conclusion requires only that the correspondence shifts up at the least and the greatest fixed points, k_S^θ and k_L^θ . The results presented in Section 3 similarly require only local shifts in steady states. That we only need to verify that \mathcal{M} shifts up *locally*, in particular, *at* the steady states, enables us to separate direct responses (or the “all-else-equal” behavior) from equilibrium responses.

To explain a little further, let us consider a particularly simple case where a dynamic economy can be reduced to a fundamental equation of the form

$$G(k_t, k_{t-1}, \theta) = 0, \quad (28)$$

where $\theta \in \mathbb{R}$ is an exogenous parameter, $k_t \in \mathbb{R}$ is capital, or the capital-labor ratio, at date t and $G : \mathbb{R}^3 \rightarrow \mathbb{R}$ a smooth function. In this case, the market correspondence can be defined as

$$\mathcal{M}^\theta(k) = \{\hat{k} : G(\hat{k}, k, \theta) = 0\}. \quad (29)$$

In the Ramsey-Cass-Koopmans model, for example, $G(k_t, k_{t-1}, \theta) = 0 \Leftrightarrow k_t = g(k_{t-1}, \theta)$, and then $\mathcal{M}^\theta(k) = g(k, \theta)$. Clearly, k^* is a steady state given θ if and only if $k^* \in \mathcal{M}^\theta(k^*)$. Note, however, that (28) — even in the more general form $0 \in G(k_t, k_{t-1}, \theta)$ where G is a correspondence — is not general enough to nest our one-sector behavioral growth model (because we also need to condition on the distribution of assets). Nevertheless, (28) is useful to provide the technical intuition for our main results since both in the case of (29) and our Definition 3, the market correspondence is constructed by conditioning on the information that the capital-labor

²⁹ \mathcal{M} is *continuous but for jumps up* if it has convex values, $\limsup_{x^n \uparrow x} m^*(x^n, t) \leq m_L^\theta(k)$, and $\liminf_{x^n \downarrow x} m_*(x^n, t) \geq m_S^\theta(k)$.

Acemoglu and Jensen (2013) proves that if \mathcal{M} is upper hemi-continuous in k and has convex values, then it is continuous but for jumps up.

ratio in question, k , has to be consistent with a steady-state equilibrium. In particular, the fact that, with the conditioning on the steady state k^* , (29) a one-dimensional fixed point problem allows us to use “curve shifting” arguments without imposing any type of monotonicity on the dynamical system defined by (28) (see also Acemoglu and Jensen (2015) for a related discussion of non-monotone methods). Given $\mathcal{M}^\theta(k)$ and this construction, Theorem 4 and the results presented in Section 3 enable us to predict how the greatest and the least steady states vary with θ when $\mathcal{M}^\theta(k)$ shifts up locally starting at these steady states (and provided that \mathcal{M} satisfies the relevant theorem’s regularity conditions).

The added generality and flexibility is considerable significant. In many applications, including the problem of equilibrium analysis in the behavioral growth model we focus on in this paper, the conditions for the “for all k curve shifting theorem” will not hold even if (28) applies. This is for both substantive and technical reasons. Substantively, as already mentioned, in economies such as the one-sector behavioral growth model the possible heterogeneity in the responses of agents to changes in the environment would often preclude such uniform shifts. To see the technical problem, suppose that we were checking these conditions using the implicit function theorem. That would amount to verifying that $\frac{dk}{d\theta} > 0$ for all \tilde{k} while $G(k, \tilde{k}, \theta) = 0$ holds. But since the implicit function theorem requires as a minimum that $D_k G(k, \tilde{k}, \theta) \neq 0$, and “running through all \tilde{k} ’s” will almost invariably violate this condition for some \tilde{k} , this method will generally fail (order theoretic methods are of no help here either; and of course, it is *not* enough to show that $\frac{dk}{d\theta} > 0$ for almost every \tilde{k} because any point we fail to check may precisely be a point where the market correspondence “jumps”). When we only need to check local conditions, these difficulties are bypassed.

Appendix C: Additional Examples of Behavioral Growth Models

We now show that our results can also be applied with other popular models of behavioral preferences and biases than those considered in Section 4. In particular, we briefly outline how these models satisfy the conditions of Lemma 1, and therefore our main results, Theorems 1-3, can be readily applied when households have (a mixture of) these preferences.

Random Utility, Mistakes, Approximate Rational, and Satisficing Behavior

Consider an additive objective with delay discounting but assume that utility at each date is random: $U(c_0, c_1, c_2, \dots) = u^{\epsilon_0}(c_0) + f(1)u^{\epsilon_1}(c_1) + f(2)u^{\epsilon_2}(c_2) + \dots$. The random utility parameter ϵ_t reflects the household’s idiosyncratic tastes or biases (McFadden (1974), p.108). There are two (mathematically equivalent) interpretations. The first is that the household is uncertain about its

future preferences, and if the objective is dynamically inconsistent, it is consequently uncertain about the behavior of future selves.³⁰ In the second interpretation, u^{ϵ_t} is a self's "true utility" at date t if and only if $\epsilon_t = 0$, hence if $\epsilon_t \neq 0$ the household will make a mistake because it will maximize an objective that departs from its true objective. In either situation, the household's savings function will be a behavioral process in the sense of Train (2009), p.3. In the second interpretation, this behavioral process describes approximate rational behavior (Luce (1959)), which if the distribution of ϵ_t is uniform on $[-a, a]$, $a > 0$, can also be interpreted as satisficing/ ϵ -optimizing behavior in the sense of Simon (1956).³¹ Finally, if ϵ parametrizes selves' subjective beliefs, a time-stationary savings function will be the quantal response equilibrium (McKelvey and Palfrey (1995)) of the game the current self plays with future selves.

Such (generalized) random utility models fit straightforwardly into (10). If the only source of uncertainty is the random utility parameter ϵ , we have $z = \epsilon$ and conditioned on z_t, w_t, R_t , and T_t , the household's beliefs about future (after-tax) prices, transfers, and about $(\epsilon_\tau)_{\tau=t}^\infty$ are given by $P_t(\cdot) = P(\cdot; w_t, R_t, T_t, z_t, \theta^M)$ where θ^M is the true model. For the conditions of Lemma 1 to hold we need that $u^\epsilon(c)$ is concave and continuous in c , continuous in ϵ and we must impose the belief related continuity conditions of Assumption 1.³²

An interesting example, which was mentioned in Section 2.2, is when ϵ_t objectively (*i.e.*, as expressed through μ_z which is part of the true model θ^M) is *i.i.d.* with mean 0 but the household through $P_t(\cdot)$ (incorrectly) believes that ϵ_t first-order stochastically dominates this objective probability, and in particular has mean greater than 0. If savings increases in ϵ_t , the household (or agent) is then "over-optimistic" about its future frugality which causes it to save less today than it would if the beliefs were correct. At the following date, the household will of course be "disappointed" for not living up to its own expectations but with the time-stationary savings function of Definition 1, it goes on to assume that *next year* it will start saving more, and so on, year after year. This is an example of a systematic bias or misperception (which can in our model be driven both by time-consistent and naive behavior).

Sparse Maximization and Inattention

A household faced with an infinite (or even just a long) time horizon may, optimally or as a rule-of-thumb, opt to keep down mental costs involved in estimating, assessing and using

³⁰The game between temporal selves will in this situation be a Bayesian game.

³¹Whether a household maximizes a function that is ϵ away from the true objective or ϵ -maximizes the true objective amounts to the same as long as the decision function is continuous in ϵ .

³²If the objective is dynamically consistent, it would also fit into (9) and we would be able to consider non-additive utility. In this case, we must in addition impose complementarity conditions on the aggregator to ensure that goods are normal (Chipman (1977)). See the discussion immediately prior to Lemma 1.

objective probabilities and calculating the optimal decision (Sims (2003)). One way to capture this in dynamic consumption and saving problems is to take as objective $\sum_{t=1}^T \beta^t u(c_t)$ where T is finite; so that the household looks only T periods into the future at any point in time. This specification fits into (10) by taking H , h , and g equal to the identity function, $u_0^\epsilon(c) = u(c)$, and $\tilde{U}^\epsilon(c_1, c_2, \dots) = \sum_{t=1}^T \beta^{t-1} u(c_t)$. Note that because at any future date, the household also looks T periods into the future, such preferences are dynamically inconsistent; and unless the current self is naive in the sense of Strotz (1956), it will thus take as given the expected (inattentive) behavior of future selves. In the deterministic case, the conditions of Lemma 1 will hold if u is continuous and concave, and in models with uncertainty we must in addition impose continuity conditions on the belief formation as described in the previous model.

Finite time horizon objectives may be interpreted as a simple version of “sparse maximization” in the sense of Gabaix (2014, 2017). It can be combined with the random utility model above by taking the time-horizon of future selves as an idiosyncratic characteristic of the household, so that the maximization problem becomes $\sum_{t=1}^\epsilon \beta^t u(c_t)$, where $\epsilon \in \{1, 2, \dots, \hat{T}\}$ and the probability distribution over ϵ reflects the household’s (subjective) beliefs about future selves’ time-horizon (as reflected by the beliefs $P(\cdot)$). Here, the sparsity of the planning horizon at future dates is uncertain from the point of view of today, and the household is uncertain about how inattentive/sparse future selves will be. Other, richer types of sparsity constraints following Gabaix (2014, 2017) can also be incorporated, for example, by reducing the set of choice variables. A particularly fruitful approach is to replace the max operator in (12), with the “sparse max” operator of Gabaix. As explained in Gabaix (2017), the “sparse max” formulation is quite tractable and implies a “sparse” Bellman operator which is a monotone contraction (see Gabaix (2017), Lemma 3.6). Using this, general savings correspondences for sparse maximization are easily shown to satisfy the conclusions of Lemma 1.

Ambiguity

If a household has incomplete information about the objective probabilities governing the random disturbances in z , then even under rational expectations ($P(\cdot) = P(\cdot; w, R, T, z, \theta^M)$ probabilistically correct given the model θ^M and current observations) it may not have unique subjective beliefs unless it satisfies the axioms of Savage (1954) (note that this has nothing to do with whether the subjective beliefs are right or wrong as discussed in the previous example). Since we have allowed the beliefs $P(\cdot)$ to be non-additive, in which case the expectations in (9)-(10) refer to the Choquet integral, most models of ambiguity are readily covered by our specification. In particular, non-additive $P(\cdot)$ may reflect “multiple priors” since Choquet expected utility with

convex capacities equals the minimum expected utility over the probabilities in the capacity's core (Gilboa (1987), Schmeidler (1989), Gilboa and Schmeidler (1995)).

Since we are following Epstein and Zin, ambiguity does not lead to any new complications for Lemma 1 because the Choquet integral has the same continuity properties as the Lebesgue integral.

Rules-of-Thumb Behavior

Our results critically depend the fact that Lemma 1 implies that the savings correspondence has increasing extremum selections. One would typically verify the conditions of this lemma starting from behavioral microfoundation as we have done so far. But our results can also be applied directly by taking the savings function (or a savings correspondence) as the starting point, and then impose simple decision rules (rules of thumb) without an explicit micro-foundation. In this context, the conclusions of Lemma 1 become assumptions on the saving function (and these assumptions are equivalently weak). As an example, suppose that the household saves a fraction of current (or perceived current) income, with that fraction depending on some current variables such as a measure of the environment's variability. Just like the systematically wrong beliefs of Section 4.3, such rules-of-thumb may include "highly irrational" behaviors.

References

- Acemoglu, D. (2009), *Introduction to Modern Economic Growth*, Princeton University Press, 2009.
- Acemoglu, D. and M.K. Jensen (2010): "Robust Comparative Statics in Large Static Games", *IEEE Proceedings on Decision and Control* 49.
- Acemoglu, D. and M.K. Jensen (2013): "Aggregate Comparative Statics", *Games and Economic Behavior* 81, 27-49.
- Acemoglu, D. and M.K. Jensen (2015): "Robust Comparative Statics in Large Dynamic Economies", *Journal of Political Economy* 123, 587-640.
- Ahn, H., G. Kaplan, B. Moll, T. Winberry and C. Wolf (2018): "When Inequality Matters for Macro and Macro Matters for Inequality", NBER macroeconomics annual 32, 1-75.
- Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2018): "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach", *NBER Working Paper Series*.
- Aiyagari, S.R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving", *Quarterly Journal of Economics* 109, 659-684.
- Al-Najjar, N.I. (2004): "Aggregation and the Law of Large Numbers in Large Economies", *Games and Economic Behavior* 47, 1-35.

- Aubin, J.-P. and H. Frankowska, *Set-Valued Analysis*, Birkhauser, Boston, 1990.
- Aumann, R.J. (1965): "Integrals of Set-Valued Functions", *Journal of Mathematical Analysis and Applications* 12, 1-12.
- Backus, D.K., B.R. Routledge, and S.E. Zin (2004): "Exotic Preferences for Macroeconomists", *NBER Macroeconomics Annual* 19, 319-390.
- Balbus, L., K. Reffett, and L. Wozny (2015): "Time Consistent Markov Policies in Dynamic Economies with Quasi-hyperbolic Consumers", *International Journal of Game Theory* 44, 83-112.
- Barro, R.J. (1999): "Ramsey Meets Laibson in the Neoclassical Growth Model", *Quarterly Journal of Economics* 114, 1125-1152.
- Barro, R.J. and X.I. Sala-i-Martin, *Economic Growth*, MIT Press, 2004. 2
- Becker, G.S. (1962): "Irrational Behavior and Economic Theory", *Journal of Political Economy* 70, 1-13.
- Becker, R.A. and J.H. Boyd, *Capital Theory, Equilibrium Analysis, and Recursive Utility*, Blackwell, 1997.
- Bernheim, B.D. and A. Rangel (2004): "Addiction and Cue-Triggered Decision Processes", *American Economic Review* 94, 1558-1590.
- Beshears J., J.J. Choi, A. Fuster, D. Laibson and B.C. Madrian (2013): "What Goes Up Must Come Down? Experimental Evidence on Intuitive Forecasting", *American Economic Review* 103, 570-574.
- Cagan, P. (1956): "The Monetary Dynamics of Hyperinflation", In Friedman, M. (ed.). *Studies in the Quantity Theory of Money*. Chicago: University of Chicago Press.
- Cao, D. and I. Werning (2018): "Saving and Dissaving with Hyperbolic Discounting", *Econometrica* 86, 805-857.
- Chipman, J.S. (1977): "An Empirical Implication of Auspitz-Lieben-Edgeworth-Pareto Complementarity", *Journal of Economic Theory* 14, 228-231.
- Epstein, L.G. and J.A. Hynes (1983): "The Rate of Time Preference and Dynamic Economic Analysis", *Journal of Political Economy* 91, 611-635.
- Epstein, L.G. and S.E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework", *Econometrica* 57, 937-969.
- Epstein, L.G. and S.E. Zin (1991): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis", *Journal of Political Economy* 99, 263-286.
- Fudenberg, D. and D.K. Levine (2006): "A Duel Self Model of Impulse Control", *American Eco-*

- conomic Review* 96, 1449-1476.
- Fudenberg, D. and D.K. Levine (2012): "Timing and Self-Control", *Econometrica* 80, 1-42.
- Fuster A., B. Hebert and D. Laibson (2012): "Investment Dynamics with Natural Expectations", *International Journal of Central Banking* 8. 243-265.
- Gabaix, X. (2014): "A Sparsity Based Model of Bounded Rationality", *Quarterly Journal of Economics* 129, 1661-1710.
- Gabaix, X. (2017): "Behavioral Macroeconomics via Sparse Dynamic Programming", Working Paper, Harvard.
- Gilboa, I. (1987): "Expected Utility with Purely Subjective Non-Additive Probabilities", *Journal of Mathematical Economics* 16, 65-88.
- Gilboa, I. and D. Schmeidler (1995): "Case-Based Decision Theory", *Quarterly Journal of Economics* 110, 605-639.
- Gul, F. and W. Pesendorfer (2001): "Temptation and Self-Control", *Econometrica*, 69, 1403-143.
- Gul, F. and W. Pesendorfer (2004): "Self-Control and the Theory of Consumption", *Econometrica*, 72, 119-158.
- Halevy, Y. (2015): "Time consistency: Stationarity and Time Invariance", *Econometrica* 83, 335-352.
- Hansen, L.P. and T. Sargent (2001): "Robust Control and Model Uncertainty", *American Economic Review* 91, 60-66.
- Hansen, L.P. and T. Sargent (2010): "Fragile Beliefs and the Price of Uncertainty", *Quantitative Economics* 1, 129-162.
- Hansen, L.P., T. Sargent and T. Tallarini (1999): "Robust Permanent Income and Pricing", *Review of Economic Studies* 66, 873-907.
- Harris, C. and D. Laibson (2001): "Dynamic Choices of Hyperbolic Consumers", *Econometrica* 69, 935-957.
- Huggett, M. (1997): "The One-Sector Growth Model with Idiosyncratic Shocks: Steady States and Dynamics", *Journal of Monetary Economics* 39, 385-403.
- Jensen, M.K. (2007): "Monotone Comparative Statics in Ordered Vector Spaces", *The B.E. Journal of Theoretical Economics* 7 (Topics), Article 35.
- Jensen, M.K. (2018): "Distributional Comparative Statics", *Review of Economic Studies* 85, 581-610.
- Jensen, M.K. (2020): "The Ego Loss Approach to Dynamic Inconsistency", Working Paper, University of Nottingham.
- Koopmans, T. C. (1960): "Stationary ordinal utility and impatience", *Econometrica* 28, 287-309.

- Kreps, D.M. and L. Porteus (1978): "Temporal Resolution of Uncertainty and Dynamic Choice Theory", *Econometrica* 46, 185-200.
- Krusell, P., B. Kuruscu and A.A. Smith (2002): "Equilibrium Welfare and Government Policy with Quasi-geometric Discounting", *Journal of Economic Theory* 105, 42-72.
- Krusell, P. and A. A. Smith (2003), "Consumption-Savings Decisions with Quasi-Geometric Discounting", *Econometrica* 71, 365-375.
- Krusell, P., B. Kuruscu, and A. A. Smith (2010), "Temptation and Taxation", *Econometrica* 78, 2063-2084.
- Laibson, D. (1996): "Hyperbolic Discount Functions, Undersaving, and Savings Policy", *NBER Working Paper Series* 5635.
- Laibson, D. (1997): "Golden Eggs and Hyperbolic Discounting", *Quarterly Journal of Economics* 112, 443-477.
- Light, B. (2018): "Uniqueness of Equilibrium in a Bewley-Aiyagari model", *Economic Theory* 69, 435-450.
- Light, B. and G.Y. Weintraub (2018): "Mean Field Equilibrium: Uniqueness, Existence, and Comparative Statics", *Operations Research* (forthcoming).
- Loewenstein, G., T. O'Donoghue and S. Bhatia (2015): "Modeling the Interplay Between Affect and Deliberation", *Decision* 2, 55-81.
- Loewenstein, G. and D. Prelec (1992): "Anomalies in Intertemporal Choice: Evidence and an Interpretation", *Quarterly Journal of Economics* 107, 573-598.
- Lucas, R.E. and N.E. Stokey (1984): "Optimal Growth with Many Consumers", *Journal of Economic Theory* 32, 139-171.
- Luce, R., *Individual Choice Behavior*, Wiley, New York, 1959.
- Machina, M. (1982): "Expected Utility Analysis without the Independence Axiom", *Econometrica* 50, 277-323.
- McFadden, D. (1974): "Conditional logit analysis of qualitative choice behavior", in P. Zarembka, ed., *Frontiers in Econometrics*, Academic Press, New York, 105-142.
- McKelvey, R. And T. Palfrey (1995): "Quantal Response Equilibria for Normal Form Games", *Games and Economic Behavior* 10, 6-38.
- Milgrom, P. (1994): "Comparing Optima: Do Simplifying Assumptions Affect Conclusions?", *Journal of Political Economy* 102, 607-615.
- Milgrom, P. and J. Roberts (1994): "Comparing Equilibria", *American Economic Review* 84, 441-459.
- Milgrom, P. and C. Shannon (1994): "Monotone Comparative Statics", *Econometrica* 62, 157-180.

- M. Nerlove (1958): "Adaptive Expectations and Cobweb Phenomena", *Quarterly Journal of Economics* 72, 227-240.
- Phelps, E. S. and R. A. Pollak (1968): "On Second-Best National Saving and Game Equilibrium Growth", *Review of Economic Studies* 35, 185-199.
- Pollak, R.A. (1971): "Additive Utility Functions and Linear Engel Curves", *Review of Economic Studies* 38, 401-414.
- Quah, J. K.-H. (2007): "The Comparative Statics of Constrained Optimization Problems", *Econometrica* 75, 401-431.
- Savage, L.J., *Foundations of Statistics*, Wiley, New York, 1954.
- Schmeidler, D. (1989): "Subjective Probability and Expected Utility without Additivity", *Econometrica* 57, 571-587.
- Simon, H.A. (1956): "Rational Choice and the Structure of the Environment", *Psychological Review* 63, 129-138.
- Sims, C.A. (2003): "Implications of Rational Inattention", *Journal of Monetary Economics* 50, 665-690.
- Strotz, R.H. (1956): "Myopia and Inconsistency in Dynamic Utility Maximization", *Review of Economic Studies*, 165-180.
- Sun, Y. (2006): "The Exact Law of Large Numbers via Fubini Extension and Characterization of Insurable Risks", *Journal of Economic Theory* 126, 31-69.
- Tarski, A. (1955): "A Lattice Theoretic Fixed Point Theorem and its Applications", *Pacific Journal of Mathematics* 5, 285-309.
- Thaler, R. and H. Shefrin (1981): "An Economic Theory of Self-Control", *Journal of Political Economy* 89, 392-406.
- Topkis, D.M. (1978): "Minimizing a Submodular Function on a Lattice", *Operations Research*, 26, 305-321.
- Train, K., *Discrete Choice Methods with Simulation*, 2nd Edition, Cambridge University Press, 2009.
- Uhlig, H. (1996): "A Law of Large Numbers for Large Economies", *Economic Theory* 8, 41-50.
- Vives, X. (1990): "Nash Equilibrium with Strategic Complementarities", *Journal of Mathematical Economics* 19, 305-321.
- Woodford, M. (2013): "Macroeconomic Analysis Without the Rational Expectations Hypothesis", *Annual Review of Economics* 5, 303-346.
- Yano, M. (1989): "Comparative Statics in Stochastic Dynamic Models", *Journal of Mathematical Economics* 18, 169-185.