ESSAYS ON PRICE PROTECTION AND INVESTOR ATTENTION

GABRIEL BUCHBINDER

A DISSERTATION

PRESENTED TO THE FACULTY

OF PRINCETON UNIVERSITY

IN CANDIDACY FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE

BY THE DEPARTMENT OF ECONOMICS

ADVISOR: STEPHEN MORRIS

SEPTEMBER 2019

Abstract:

This collection of essays includes a theoretical evaluation of price protection and a study of investor attention using a new measure based on Google searches.

Chapter 1 introduces price protection to a dynamic pricing environment with homogeneous goods and different agent types. I find that prices regularly drop to allow the lowest valuation to make a purchase. I characterize how prices fluctuate during repeating cycles in equilibrium and the firm optimal length of such cycles for different parameters under no commitment. I find that price protection lowers social welfare by delaying price drops.

Chapter 2 introduces a measure of investor attention. The finance literature on investor attention has recently measured interest in stocks by using the search volume index (SVI) of companies' ticker symbols on Google. I use the major stylized fact of local bias (LB) in household stock-portfolio choices to test the precision of this measure. Constructing a database of Google searches on stock tickers at the metropolitan level, I show that stock-ticker-SVIs fail to exhibit the property of LB. I therefore propose an alternative measure of investor attention to stocks which seem to reflect more naturally the way retail investors search for information on Google and is consistent with LB behavior. This measure increases estimated LB effects by a factor of 10, suggesting that one standard deviation increase in distance lowers investor interest by around 15%.

In Chapter 3, which I co-author with Ioannis Branikas, we use this measure to test for the effect of product advertisements on investor attention. Using Super Bowl as an experiment and a new measure of local investment interest on stocks from Google searches, we study the effects of advertising expenses on investor attention. We find that the post-game Monday attention of investors in areas with high viewership ratings increases significantly for the stocks of companies that aired commercials, regardless of whether these are local or non-local. Distant firms with high advertis-

ing exposure in a region attract more interest than local firms with low exposure, suggesting that advertising has a stronger effect on investor attention than local bias.

Acknowledgments:

I am extremely grateful to my advisor, Stephen Morris, for his guidance, support, encouragement, and patience. I am also thankful for helpful comments and suggestions from the other members of my dissertation committee, Faruk Gul and Wei Xiong.

I also thank Alan Blinder, Roland Bénabou, Jakub Kastl, Wolfgang Pesendorfer and Can Urgun for multiple useful discussions on the topics of my dissertation and other projects. Thanks to Laura Hedden, for administrative support.

A special thanks to my friend and co-author Ioannis Branikas, who along with my good friends Jakob Schlockerman and Dmitry Mukhin helped me along this journey.

Thanks to my Master's advisor, Vinicius Carrasco, whose support was fundamental for me to enter this PhD program.

Most importantly, thanks to my parents Ricardo and Betina for support throughout, and to my brother, Felipe, who has always been there for me.

To my family

Contents

Chapter 1

Price Protection and Welfare*

1.1 Introduction

Conlisk, Gerstner, and Sobel [1984] (hereby CGS) consider the problem of a price setting monopolist offering a good to a constant flow of potential buyers with differing valuations for the good. They find that prices follow repeating cycles, inside which they are monotonically decreasing: at the beginning of these cycles, prices start at a high level; then, they drop monotonically until all consumers in the market make a purchase. When that happens, prices rise and the cycle repeats itself. The length of such cycles is calculated through the monopolist's problem, who is unable to commit to future prices, but is able to predict his future selves' behavior.

What happens to this dynamics if the monopolist offers a price protection policy? CGS do not account for this situation and yet it is not uncommon for retailers in the United States to offer such policies. This paper contributes to extant theory by filling this gap. Namely, this paper introduces price protection to the CGS model. Considering a simple adaptation of price protection, where consumers get the price difference back with some probability if the price drops the period following

*I am indebted to my advisor, Stephen Morris. I also thank Faruk Gul, Jakub Kastl and Can Urgun for helpful comments and discussions.

a purchase, this paper looks at how pricing decisions change, and how welfare is affected by the introduction of this policy.

The general outline of our model parallels that of the CGS model: a monopolist offers a good it can produce at zero cost to two types of consumers, a mix of which arrive at every period. These types of consumers are called the High and the Low type, based on their willingness to pay for the product. Selling only to the higher type may be profit maximizing, but the firm is unable to commit to it and so it cannot occur in equilibrium: if it did, the lower type consumers would keep on accumulating, and the incentive to sell at a lower price would be overwhelming. Therefore, at some period, prices must be such that the lower valuation consumers are willing to make a purchase. At such a time, every consumer in the market makes a purchase, and the economy goes back to the initial state. Realizing this, the firm can never charge as much as what the high valuation consumer is willing to pay, when facing rational consumers who do not completely discount the future. So far, this model is no different than that of CGS. But now, we add price protection.

The addition of price protection gives the consumer an exogenous probability of receiving a rebate equal to the price difference, whenever the price drops the period following his purchase. We show that there is a Pareto worsening with the increase of this probability, where both types of consumers' welfare and the firm's profits decrease in the relevant parameter range. While the policy leads to more stable prices, it also leads to delays in sales and generally longer cycles. This means that a participant entering a random period is, on average, further away from a period where the price is low, so increasing the effective price high type consumers pay, and delaying the time until low type consumers are able to make a purchase.

In this paper we isolate and focus specifically on the effects of price protection on the dynamics of price setting, and tests the welfare implications that arise from this direction. As described below, other mechanisms and behavioral preferences could lead to different conclusions. In particular, the role of monopoly here is important, as the promise of the policy can attract consumers independently of their actual benefit to the consumer.

In the following subsection, we discuss the concept and modalities of price protection in theoretical terms. Next, we review extant literature which is related to our research question. We then proceed to developing our model. This is done in sections 2 to 5, where we describe it, propose an equilibrium for a large subset of the relevant parameter range and estimate welfare properties. Our findings are discussed throughout these sections. Finally, we conclude our paper by discussing implications of our model as well as suggestions for future research.

1.1.1 Price Protection in General

Retailers for consumer products often offer price protection policies. Under such policies, whenever the seller makes a price reduction on one of his sold goods, all consumers who recently purchased that item have the right to receive the price difference in return. These policies vary in length and in the details of when they apply and how to obtain the refund, but their main mechanism is generally the same.

Price protection policies attract consumers, in part, by assuaging consumer's fears of regretting the purchase if prices drop in the future. Price protection might also entice potential buyers to make an immediate purchase, rather than wait for a potential future promotion.

At the same time, they can be profitable for firms as well. Sellers may decide to offer these policies to avoid the costs of dealing with product returns, as generous return policies can be used by consumers to mimic price protection. Price protection can also be used by the seller as a commitment device, as they make it less profitable for the firm to lower its prices in the future.

However, there are situations where the agent who offers the price protection policy is distinct than the agent who prices the product. This is the case of price protection policies offered by credit card companies, for example, which is not the focus of this paper.

As a final remark concerning the definition of price protection, notice the distinction between price protection policies and price matching policies. In the latter, sellers offer to match the price of their competitors. These policies are also offered as a benefit to consumers, while they may help firms sustain cartels. In such a situation, price matching drops consumer welfare, while increasing the surplus of the sellers.

With these general features of price protection in mind, we now review previous literature which is related to the model we develop in this paper.

1.1.2 Related Literature

The durable goods literature with a single cohort of long lived consumers has explained why prices start at a high level when new products are launched and drop with time. Coase [1972], Stokey [1981], Bulow [1982], Gul et al. [1986] analyze these markets when no new consumer enters the market for a durable good. The monopolist price discriminates across time periods, taking advantage of the impatience of the high valuation consumers to extract a higher surplus from the market, while competing with his future selves, which affects his monopoly rent.

The folk theorem in Sobel [1991] shows that organized buyers can sustain any feasible payoff, in subgame perfect equilibria, when there is a high enough discount rate. That result depends on the monopolist and the buyers participating in punishing phases, in case the monopolist does not follow with the expected pricing scheme, or in case the consumers do not make purchases as expected. These punishment phases allow the market to sustain pricing cycles of arbitrary length.

Cooper [1986] looks into duopolies with infinite periods, but nondurable goods and the option of offering a price protection policy. He assumes that at each period the demand function is the same, and concludes that there is a clear benefit of (at least) one party offering a price protection (in his nomenclature - a most-favored-customer pricing) policy: a Stackelberg equilibrium can be achieved in each period, with the party offering the policy being the leader, as his previous period's price makes it unprofitable to lower the price to his Cournot best response level. Alternatively, Cooper [1984] and Cooper and Fries [1991] consider monopolists facing two consumers, and show how the monopolist can increase his welfare by either signaling high costs, or by increasing his bargaining power against the second purchaser.

Lee et al. [2000] consider the case of price protection on a non-durable good that suffers technological depreciation on a 3 period model. In their model, they evaluate how price protection between a manufacturer and a reseller can be used as a coordination mechanism. That helps shift sales forward in time, and increase total sales. Neither is observed in the current paper.

Butz [1990] and Lai et al. [2010] consider two period models, with the value of the good dropping between the first and second periods.

The next five sections are the heart of this paper. In section 2, we describe our model conceptually and introduce notation. Next, in section 3, we develop our model mathematically and deduce its fundamental equations. Following the description of the general case, we consider in section 4 the specific but quite relevant case of an equilibrium with low incidence of price protection — roughly, where fewer than half the population takes advantage of it. Subsequently, in section 5, we look at the welfare effects of increasing price protection incidence in this scenario. The last section concludes the paper.

We thus begin by describing our model conceptually in the following section.

1.2 Model

We consider the repeated sale of a continuum of a durable homogeneous good by a price setting monopolistic firm. At each period, a continuum with mass 1 of potential buyers enters the market, each with a valuation for the good $\theta \in \{\theta_L, \theta_H\} \in \mathbb{R}^2_{++}$. Buyers have a probability $\alpha \in (0,1)$ of being of type $\theta_H > \theta_L$, are risk-neutral, and stay in the market until they make a purchase, at which point they leave the market forever.

At each period $t \in \mathbb{N}$, the monopolistic firm observes the history of prices, $\mathcal{P}_{t-1} = \{p_{-1}, \ldots, p_{t-1}\}$, of purchases, $\mathcal{X}_{t-1} = \{x_{-1}, ..., x_{t-1}\}$,and of the high- and low-valuation populations of available consumers, $\mathcal{M}_t = \{M_0, ..., M_t\}$, where $M_\tau = \left\{M^H_\tau, M^L_\tau\right\}$ for $\tau \in \{0, ..., t\}$. The firm then chooses a price for the current period $p_t \in \mathbb{R}_+$. The firm can instantaneously produce the good at zero cost, so that the market can always be satisfied. After the firm's choice, consumers observe the history of prices $\mathcal{P}_{t-1} \cup \{p_t\} = \mathcal{P}_t$, as well as $\mathcal{X}_{t-1}$ and $\mathcal{M}_t$, and make a decision on whether to purchase immediately or wait.

Consumers who decide to purchase at period t will pay price p_t and receive the good. If, however, the price drops in the following period, $t+1$, the consumers who bought the good at t have a chance ϵ, taken as exogenous, of benefiting from price protection: they receive back the price difference between the two periods. This probability can be interpreted as the chance that a consumer notices the price drop in the future, or is able to fulfill the required bureaucracy to receive the price difference back. It could be that ϵ depends on the magnitude of the price drop, which isn't considered here.

Both consumers and the firm have a discount rate of δ. Since the consumers receive the price protection rebate only at the period after their purchase, that rebate is appropriately discounted. For simplicity and to facilitate the equilibrium format proposed, a tie-breaking condition is used: consumers will make a purchase earlier, rather than later, if either action is expected to give them the same utility.

This paper focuses on finding a Markov Equilibrium in this model, which is defined as a set of (i) correct expectations on future prices, (ii) a price function: $p\left(p_{t-1}, x_{t-1}, M_t^H, M_t^L\right) = p_t \in \mathbb{R}_+$ which maximizes the firm's profits subject to consumer's decisions and expectations, and (iii) a purchasing decision function $d\left(p_{t-1}, x_{t-1}, M_t^H, M_t^L; \theta\right) \in \{0,1\}$ which maximizes type θ consumer's utility, based on their expectations of future prices. The first lemma at the next section will allow us to simplify the relevant history of the game.

1.3 Partial Description of Equilibrium

In this section, we develop the fundamental properties and equations of our model. We begin by proving that, if a Markov equilibrium in pure strategies exist — which will be dealt with in Section 1.4—, then prices exhibit a cyclical behavior. We next employ this result to derive the fundamental equation from which the dynamics of pricing in equilibrium will be obtained.

The following Lemma 1will help us prove the existence of cycles in Markov equilibria for our model. By itself, it shows that from any initial state of the world we will eventually arrive at a state when prices are the lowest possible — which we prove to be equal to the lowest valuation—, and all past entrants have made a purchase. The existence of cycles is therefore a corollary of this lemma, since after such states the past history is not directly relevant for profits and utilities.

Lemma 1. *In any Markov equilibrium, for any starting state* $S_0 = \left(p_{-1}, x_{-1}, M_0^H, M_0^L\right)$, *there exists a* $\bar{t} \geq 0$ *such that:*

$$0 \leq t < \bar{t} \Rightarrow p_t > \theta_L$$

$$p_{\bar{t}} = \theta_L$$

Proof. First, note that prices below θ_L are not rationalizable, and thus $p(S) \geq \theta_L$ for any state $S = \left(p, x, M^H, M^L\right)$. To see this, note that $p(S)$ is always non-negative and so it has a lower bound $\underline{p} \geq 0$. Given this, a type θ consumer will always make a purchase immediately whenever the firm sets a price p such that:

$$\theta - p + \epsilon\delta\left(p - \underline{p}\right) \geq \delta\left(\theta - \underline{p}\right)$$

$$p \leq \underline{p} + \frac{1-\delta}{1-\epsilon\delta}\left(\theta - \underline{p}\right)$$

Given $\theta_L < \theta_H$, both types of consumers will make a purchase immediately whenever $p \leq \underline{p} + \frac{1-\delta}{1-\epsilon\delta}\left(\theta_L - \underline{p}\right)$, which is greater than $\underline{p}$ when this lower bound ($\underline{p}$) is below θ_L. In other words, for any lower bound $\underline{p} < \theta_L$, the firm can sell to all available customers by setting a price above this lower bound, and therefore setting the price to be $\underline{p} < \theta_L$ is not rationalizable. Thus we know that $\underline{p} \geq \theta_L$.

Second, note that (i) if $p_t > \theta_L$ for all $0 \leq t \leq T$, then $M_T^L = M_0^L + T(1-\alpha)$; and (ii) for every $\tau \geq 0$, there is always a $\tau_h \geq \tau$ such that all $M_{\tau_h}^H = \alpha$, that is, in which all H-consumers decided to make a purchase the previous period[1]. These two facts together mean that, if $p_t > \theta_L \forall t \geq 0$, then for any number $K \in \mathbb{R}$ there exists a period t_K such that H-types made a purchase at t_{k-2}, so that: $M_{t_K}^L = M_0^L + t_K(1-\alpha) > K$ and either of the two cases is true:

1. H-types also made a purchase at t_{k-1}:$M_{t_K}^H = \alpha$, $x_{t_k-1} = \alpha$ and $p_{t_K-1} \leq \theta_H$ (otherwise H-consumers wouldn't have bought it at $t_K - 1$)

[1]Note that, because of the tie breaking condition, if one H-consumer makes a purchase, all H-consumers will do as well. The alternative would be for the company never to sell to any consumer - but that cannot be an equilibrium: the firm could get better off by just setting price at θ_L always.

2. H-types did not make a purchase at t_{k-1}: $M^H_{t_K} = 2\alpha$, $x_{t_k - 1} = 0$ and $p_{t_{K-1}}$ is irrelevant, since no buyers the previous period means there's no liability.

In either case, the firm's liability is bounded, at most it will have to pay $\epsilon\alpha\left(\theta_H - \theta_L\right)$ if it wishes to sell to the low types at t_k, while it can obtain $M^L_{t_K}\left(1-\alpha\right)\theta_L > K\left(1-\alpha\right)\theta_L$ by setting $p_{t_K} = \theta_L$ (giving up at most $\alpha\theta_H$ in revenue if it were to sell just to H-consumers). Therefore, the hypothesis that $p_t > \theta_L \forall t$ is absurd. □

Corollary. *Markov Equilibria in this model generate cycles: since at some point they set $p = \theta_L$, any Markov Equilibrium will eventually take the state of the economy to $(p_{-1}, x_{-1}, M_H, M_L) = (\theta_L, x_{-1}, \alpha, 1-\alpha)$[2]. As in CGS, we call the groups of periods from any such state until the period preceding the next instance of such a state as a cycle.*

Now we can look for the behavior of prices within such cycles. In particular, as consumers take the sequence of prices as given, we can describe a set of inequalities that limit the upper bound of such prices at each period, when taking as given the end of the cycle. Given that the firm is maximizing profits, it will choose prices that make the inequalities binding, as increasing prices up to them will not affect the quantity sold. The rest of this section describes the general price path.

With Lemma 1 and its corollary in mind, we will look for equilibria in which the history before the current cycle, that is, before the last time low valuation consumers made a purchase, is irrelevant. Note that if there are multiple such sub-cycle equilibria, we can combine them in any way to form equilibria for the infinite game.

In particular, we can rename the time periods to be the periods inside a cycle: Let n be the length of the current cycle and let $j \in \{n-1, ..., 0\}$ be an index which labels the periods periods within each cycle. For convenience of notation, we label the

[2]Note that the value of x_{-1} is irrelevant as $p_{-1} = \theta_L$ and the fact that the price is never below θ_L means that effectively, the firm will never pay for the price protection in the next period.

periods backwards, so that $j = 0$ denotes the last period of the cycle and $j = n - 1$ denotes the start of the cycle. When using this notation, we will be using a superscript so that, for example, $p_{n-1} = p^0$.

With this, we can state two remarks:

Remark 1. In any equilibrium, high-valuation consumers must expect to obtain at least utility equal to $\delta^j (\theta_H - p^0)$ to be willing to purchase at period j.

This remark follows by a reasoning of backwards induction. Namely, if an H-type customer waits until the last period in the cycle to buy the product by p^0, it will receive a utility of $\theta_H - p^0$. Anticipating a utility of $\theta_H - p^0$ by j periods yields a utility of $\delta^j(\theta_H - p^0)$. In order to anticipate this sale by j periods, the customer must be at least indifferent between the utility it obtains by anticipating and the utility it obtains by waiting, which leads to the current Remark.

Remark 2. The firm will adjust p^j so that the high-valuation consumers get exactly the above utility in every period.

Indeed, for the monopolist, increasing the price today is beneficial, but only as long as consumers don't postpone their purchases. If the monopolist increases the price too much, the utility of high-valuation consumers will decrease below the threshold specified in the previous remark. This is turn means high-valuation consumers will not be willing to purchase at period j but would rather opt to postpone purchases. Hence, the monopolist will adjust its prices so that it is as high as possible but without leading to the postponement of sales. This happens precisely when the utility obtained by high-valuation consumers equals the utility presented in the previous remark.

We are now in a position to derive the fundamental equation which describes the pricing dynamics adopted by the monopolist in equilibrium. Taken together, these remarks teach us that, at any period j, an H-type customer will only make a purchase if the utility obtained by such purchase equals $\delta^j(\theta_H - p^0)$.

In the absence of price protection, the utility obtained by an H-type consumer in a purchase at period j is simply $\theta_H - p^j$. In the presence of price protection policies, we must further add the utility derived from this policy. Such a utility equals the difference between prices at period j and prices at the following period, if the former prices exceed the latter, or zero otherwise. Hence, the utility derived from obtaining the price rebate, if there is a price purchase policy, is $\mathbf{1}_{\{p^j > p^{j-1}\}} \left(p^j - p^{j-1}\right)$, where we recall that $j-1$ denotes the period following j, since we are labeling the periods in decreasing order.

Considering that price protection is offered with a probability ϵ, the total utility obtained by an H-type consumer from making a purchase at period j is $\theta_H - p^j + \epsilon\delta\mathbf{1}_{\{p^j > p^{j-1}\}} \left(p^j - p^{j-1}\right)$. It thus follows from the above remarks that a sale will only be made in period j if

$$\theta_H - p^j + \epsilon\delta\mathbf{1}_{\{p^j > p^{j-1}\}} \left(p^j - p^{j-1}\right) = \delta^j \left(\theta_H - p_0\right)$$

Lemma 1 shows, however, that in the last period of the cycle, the monopolist will be willing to sell his products at θ_L in order to profit from selling the good to L-type consumers who have been increasing in numbers over the previous periods. It follows that $p^0 = \theta_L$. By making this change in the above equation, we obtain the fundamental equation of price dynamics for our model:

$$\theta_H - p^j + \epsilon\delta\mathbf{1}_{\{p^j > p^{j-1}\}} \left(p^j - p^{j-1}\right) = \delta^j \left(\theta_H - \theta_L\right)$$

The equation above enables us to derive the prices the monopolist will set for its product at each period within the cycle.

More specifically, we show in Appendix A that the equation above enables two different regimes of price evolution that must hold in any pure strategy equilibria, if they exist. In the first regime, prices decrease monotonically within each cycle, and this paper shows their existence. In the second, which is not the focus of this paper, prices oscillate, dropping in even periods and raising in odd periods. Interestingly, whether prices shall follow one regime or the other is uniquely determined by parameters ϵ and δ, though this equilibrium's existence question isn't dealt with here. Namely, there exists a threshold value

$$\epsilon_\delta \equiv \frac{1}{1+\delta}$$

such that if $\epsilon < \epsilon_\delta$, then prices follow a monotonically decreasing path within each cycle, while if $\epsilon > \epsilon_\delta$, then prices oscillate in the fashion described above. Throughout the remainder of this paper, we will focus on the case of $\epsilon < \epsilon_\delta$. Notice that this includes the case studied in CGS, where $\epsilon = 0$. For any discount rate δ between 0 and 1, it also includes any rate of price protection lower than 50% — which is the case of δ=1, and that range widens as δ decreases.

The next section derives price dynamics inside cycles when $\epsilon < \epsilon_\delta$.

1.4 Decreasing Prices Equilibrium

In the previous section, we have derived the equation of price evolution under price protection policies. We analyze this expression in depth in Appendix A and show that it allows for two markedly different behaviors. Whether prices will exhibit one behavior or the other depends exclusively on a single parameter, namely, the probability (ϵ) that any given consumer will claim his right for a refund should prices drop. We are particularly interested in cases where it is small and by this we mean

$$\epsilon < \frac{1}{1+\delta}$$

As shown in Appendix A, this is a sufficient and necessary condition for prices to exhibit a cyclic behavior, decreasing monotonically within each cycle. This is precisely the behavior found by CGS, although price protection imposes marked deviations from their predictions. It is these deviations which we seek to study in this section. Thus, in this section, we shall start from the general equation we arrived at in the previous section and develop it under the assumption that ϵ is small.

Under this assumption, this section shows that an equilibrium exists. We shall see that, as time progresses with the product being offered at a price above θ_L, the number of L-type consumers progressively accumulate. Eventually, the number of L-type consumers willing to buy the product for a price of θ_L grows sufficiently large for the monopolist to be willing to sell its product to these clients at a price θ_L even if that means having to incur on the cost of rebate for customers of type H who had bought in the previous period.

More specifically, we shall see that the price offered by the monopolist follows cycles composed of a constant number of periods. Each cycle begins with the monopolist setting a high price for its goods. Since H-type consumers know that the monopolist will eventually lower its price in an attempt to sell to L-type customers, these consumers will only buy the product if their utility from buying the product at a higher price exceeds their utility of waiting to buy the product at a lower price in the future. This leads the monopolist to progressively reduce the product price until the number of L-type consumers willing to buy the product is sufficiently high to justify selling the product at a price of θ_L. Subsequently, when all L-type consumers have bought the product, the monopolist raises the price once again.

It follows from this dynamics that the product price follows a cyclic pattern, continually decreasing within each cycle as the number of L-type consumers progressively enter the market. Mathematically, we shall see that, during the n periods which compose each cycle, the number of L-type consumers, M_L, progressively accumulates and exceeds a series of thresholds $M_L^{n-2}, ..., M_L^0$. As each threshold is exceeded, the monopolist drops the price successively to $p^{n-1}, ..., p^1$ and, finally, to $p^0 = \theta_L$.

In order to demonstrate this dynamics, let us first consider the problem faced by consumers and subsequently the problem faced by the monopolist.

1.4.1 The Consumer's Problem

We now consider the problem faced by consumers in the presence of price protection. As stated at the beginning of this section, we will focus on a scenario where ϵ is small and, therefore, where prices are monotonically decreasing within each cycle. This assumption leads to prices which are, in each period, given explicitly by the equation

$$p^j = \theta_H - \delta^j (\theta_H - \theta_L) \frac{1 + \epsilon (1 - \delta) \left(\frac{-\epsilon}{1-\epsilon\delta}\right)^j}{1 + \epsilon (1 - \delta)} \tag{1.1}$$

To see this, note that, since prices are monotonically decreasing, $\mathbf{1}_{p_j > p_{j-1}} = 1$. This allows us to rewrite the fundamental equation as

$$\theta_H - p^j + \epsilon\delta \left(p^j - p^{j-1}\right) = \delta^j (\theta_H - \theta_L)$$

Or, rearranging the terms to make p_j explicit,

$$p^j = \frac{\theta_H - \delta^j (\theta_H - \theta_L) - \epsilon\delta p^{j-1}}{1 - \epsilon\delta}$$

We now proceed by induction. In the last period, $p^0 = \theta_L$, which satisfies equation 1.1 when $j = 0$. So, assuming the equation is valid for $j - 1$, let's show that it is valid for j as well:

$$\begin{aligned}
p^j &= \frac{\theta_H - \delta^j(\theta_H - \theta_L) - \epsilon\delta p^{j-1}}{1-\epsilon\delta} \\
&= \frac{\theta_H - \delta^j(\theta_H - \theta_L) - \epsilon\delta\left[\theta_H - \delta^{j-1}(\theta_H - \theta_L)\frac{1+\epsilon(1-\delta)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j-1}}{1+\epsilon(1-\delta)}\right]}{1-\epsilon\delta} \\
&= \frac{\theta_H(1-\epsilon\delta) - \delta^j(\theta_H - \theta_L)\left(1-\epsilon\frac{1+\epsilon(1-\delta)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j-1}}{1+\epsilon(1-\delta)}\right)}{1-\epsilon\delta} \\
&= \theta_H - \delta^j(\theta_H - \theta_L)\frac{\frac{1+\epsilon(1-\delta)-\epsilon-\epsilon^2(1-\delta)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j-1}}{1+\epsilon(1-\delta)}}{1-\epsilon\delta} \\
&= \theta_H - \delta^j(\theta_H - \theta_L)\frac{\frac{1-\epsilon\delta-\epsilon(1-\delta)(1-\epsilon\delta)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j}}{1+\epsilon(1-\delta)}}{1-\epsilon\delta} \\
&= \theta_H - \delta^j(\theta_H - \theta_L)\frac{1-\epsilon(1-\delta)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j}}{1+\epsilon(1-\delta)}
\end{aligned}$$

as suggested.

Interestingly, Equation 1.1 may be rewritten to express the price in each period, p^j as a weighted average between the valuations of high- and low-type consumers. To show this, let us temporarily use the abbreviation $A_j = \delta^j\frac{1+\epsilon(1-\delta)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j}}{1+\epsilon(1-\delta)}$, so that Equation 1.1 may be more succinctly written as $p^j = \theta_H - (\theta_H - \theta_L)A_j$. By rearranging the terms on this equation we obtain $p^j = (1-A_j)\theta_H + A_j\theta_L$. Thus, we see that the price in each period j is a weighted average between the high and low valuations where the normalized weight of θ_L is $\delta^j\frac{1+\epsilon(1-\delta)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j}}{1+\epsilon(1-\delta)}$. This weight monotonically increases over time (when our assumption $\epsilon < \frac{1}{1+\delta}$ holds) until it reaches unity by the end of the cycle, where $j = 0$. Such an increase on the relative weight of θ_L means prices decrease monotonically within the cycle.

From a consumer's standpoint, prices ought to be at most the value given by Equation 1.1. Since the monopolist wants to charge consumers as much as possible this equation gives the actual prices in each period.

We now move on to consider the monopolist's problem.

1.4.2 The Monopolist's Problem

On the Equilibrium Path

Taking Sections 1.3 and 1.4.1 together, we have shown that when $\epsilon < \frac{1}{1+\delta}$, the price of the good evolves in each period according to Equation 1.1. That equation tells us that, in the last period of a cycle, prices eventually reach θ_L. However, we do not yet know how many periods it takes for the price to reach its minimum value. In other words, though we know from Lemma 1 that prices evolve in cycles, we remain ignorant on how long these cycles are. In order to answer this question, we must consider the monopolist's side of the problem, which is the focus of this subsection.

The goal of this subsection is, therefore, to find a pricing function $p(p_{-1}, \alpha, \alpha, M_L)$ that takes the states of the world along the equilibrium path and leads to a price decision for the monopolist. This function depends on (i) the previous period's price, p_{-1}, (ii) the number of buyers in the previous period — which is α along the equilibrium path[3]—, (iii) the number of high type consumers in the market — which is also always α as high type consumers are always making a purchase in equilibrium —, and (iv) the number of low type consumers, M_L— which in equilibrium is increasing every period by $1 - \alpha$ as long as $p_{-1} > \theta_L$.

This pricing function is described in Theorem 1, where thresholds for M_L are given, as deduced from this subsection, through backwards induction from the firm's and consumer's problems. We will now recall the monopolist's problem to find these thresholds.

The monopolist's problem is to decide when to lower its prices, passing from p^{n-1} to p^{n-2} and so on, until it eventually decides to lower its price to $p^0 = \theta_L$. A monopolist decides to lower prices because it wants to reap the benefits of selling the good to L-type clients who have been increasing in numbers, waiting to buy the good.

[3]Except at the first period of each cycle, but in that case $p_{-1} = \theta_L$, the lowest price the firm will ever set, and so the number of previous buyers is irrelevant.

Figure 1.1: Price thresholds

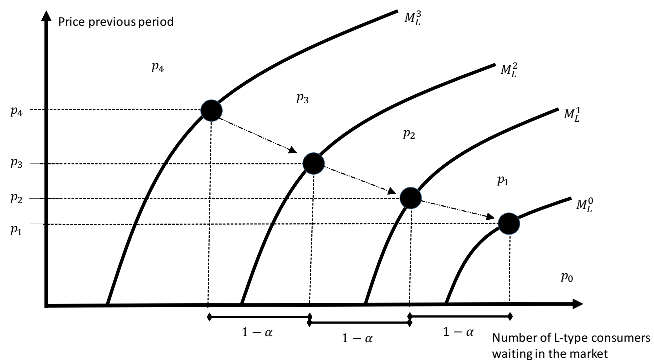

Thus, a monopolist decides to lower prices based on the number o L-type consumers in each period. We call this number M_L. If M_L is sufficiently high, a monopolist may choose to offer its product at price $p^0 = \theta_L$ and sell its good to the M_L L-type customers who have been waiting to buy the good. By sufficiently large, we mean that M_L should exceed a threshold M_L^0.

Likewise, as M_L steadily increases, the monopolist progressively reduces its price from $p^{n-1}$to p^{n-2} and so on. The monopolist decides to reduce its price from p^j to p^{j-1} whenever it sees that the number of L-type consumers has sufficiently increased, that is, whenever he sees that M_L has risen over some threshold value, M_L^{j-1}. The monopolist problem consists of identifying the values of these thresholds. Figure 1.1 illustrates these thresholds, mapping regions of the previous price and the number of low type consumers into the price to set for today. The shape of the curves is illustrative only.

Strictly speaking, the threshold values are a function of the price previously set by the monopolist. The reason for this is that the price set yesterday by the monopolist affects his cost to reduce the price today since the higher the previous price, the higher the cost of the rebate when the price drops. So, the higher the previous price, the larger M_L must be, so that the revenues of selling to M_L L-type consumers compensate the cost of rebate. Therefore, strictly speaking, M_L^j is a function of the previous price, p_{-1}, and we should therefore write $M_L^j(p_{-1})$.

Before we write the solution to this problem, however, a number of comments are useful in order to make us better understand this problem.

Note that the monopolist sees price as depending on the number M_L of L-type consumers in the market. The consumer, however, sees the price as a function evolving through time, changing in every period according to Equation 1.1. In order to better understand how these two formulas relate, it is useful to note that there is a direct relationship between M_L in each period and the period itself. Indeed, recall our assumption that, in every period, $1-\alpha$ L-type consumers enter the market. It follows that, in each period, the number of L-type consumers in the market will be $1-\alpha$ times the number of periods since the beginning of the cycle, that is, $M_L = (1-\alpha)(n-1-j)$. Therefore, M_L may be thought of both as a measure of number of L-type consumers and as a measure of time since the beginning of the cycle. This equivalence establishes a relationship between how the consumer expects prices to change (namely, depending on the number of periods until the price is lowered to its minimum value) and how the monopolist decides to set its prices (namely, depending on the number of L-type consumers in the market).

In addition, note that prices evolve sequentially from the initial price all the way to p^0, with every price value being assumed once and only once in each path to the low price period. To see this, suppose the initial price is p^2. With this in mind, let us begin by considering what would happen if the monopolist were to lower the price from p^2 to p^0, without ever pricing the good at p^1. Since consumers are assumed to have correct beliefs on the monopolist's behavior, they will refuse to buy at a price p^2 and postpone their purchases until the price is lowered to p^0. Therefore, the monopolist will never jump any price in the sequence from p^{n-2} to p^0. Alternatively, consider the case where the monopolist repeats the price p^1 two periods in a row before decreasing it to p^0. If consumers are willing to pay p^1 for the good in the second of these two periods, then they are also willing to pay p^2 in the first period in order to have the product earlier. Therefore, it is not in the interest of the monopolist to repeat any price for successive periods. It follows from both these considerations that prices change every period, following the sequence from p^{n-1} to p^0 without any jumps or repetitions.

These discussions lead us to an interesting result concerning the threshold values which M_L must overcome in order for the monopolist to decide to lower the price of the good. Since prices evolve sequentially in each period with no repetitions or omissions, it follows that, at each period, M_L should overcome one (and only one) threshold. But M_L increases every period by $1-\alpha$. It follows that the thresholds of M_L must be separated by a distance of $1-\alpha$. This in an important finding, because it allows us to reduce the monopolist's problem to finding a single threshold, for instance, $M_L^0(p_{-1})$. All other thresholds may be found simply by sequentially adding $1-\alpha$. In particular, along the equilibrium path,

$$M_L^j(p^{j+1}) - M_L^{j-1}(p^j) = 1 - \alpha$$

Let us therefore consider the problem of finding $M_L^0(p_{-1})$.

In equilibrium, a cycle has $n \in \mathbb{N}^*$ periods. In each period, $1-\alpha$ L-type consumers enter the market. It is the accumulation of these consumers which motivates the monopolist to drop the price to $p^0 = \theta_L$ in the last period and rip the benefits of selling to L-type consumers. This happens in the last (namely, the n-th) period of the cycle, when the number of L-type consumers equals $(1-\alpha)n$. It follows that $M_L^0(p^1) \leq (1-\alpha)n$ since in the n-th period the price is already set to p^0. In the previous (namely, the $(n-1)$-th) period, the price is p^1, meaning the number of L-type consumers is still less than $M_L^0(p^1)$. The number of L-type consumers in this period is $(1-\alpha)(n-1)$, so we have $M_L^0(p^1) \geq (1-\alpha)(n-1)$. Taken together, these conditions give lower and upper bounds for $M_L^0(p^1)$:

$$(1-\alpha)(n-1) \leq M_L^0(p^1) \leq (1-\alpha)n$$

Now let us imagine a more general case, where a monopolist is faced with the previous arbitrary price p_{-1} and has to decide on a new price. This price will lead to a new state of the world such that one will get to price p^0 in a finite number j of periods. Given that, we have both that the consumer will not be willing to pay more than p^j today and that the monopolist is able to extract p^j from the consumer. Therefore, the price should be exactly p^j, one of the prices in the sequence $p^{n-1}, ..., p^0$ given by equation 1.1 above.

If the number of L-type consumers is sufficiently large, the monopolist may choose to lower its price to p^0. We will denote by $m_L^0(p_{-1}, p_j)$ the minimum amount of L-type consumers necessary in the market to make the monopolist prefer to set the price to p^0 over setting the price to p^j, $j \neq 0$. This value must obey the condition that the monopolist's earnings by lowering the price to p^0 must be at least equal to setting the price to p^j. Since there are α H-type consumers and m_L^0 L-type consumers in the market, selling the product at p^0 gives the monopolist revenues of αp^0 and $m_L^0 p^0$ from each group respectively. The α H-type consumers who bought the product in the previous period at a price p_{-1} will give rise to a rebate cost of $\alpha\epsilon\,(p_{-1} - p^0)$. Finally, if we denote the present value of future earnings from the beginning of a cycle (given the equilibrium repeated cycle path) as Π, this means the total earnings obtained by the monopolist by selling the good at price p^0 equals $\alpha p^0 + m_L^0 p^0 - \alpha\epsilon\,(p_{-1} - p^0) + \delta\Pi$.

Alternatively, if the monopolist sets the price at some other value, p^j, prices will subsequently fall for j periods until they finally reach p^0. The number of L-type consumers when that happens will be $m_L^0 + (1-\alpha)\,j$. These consumers therefore represent an earning to the monopolist of $\delta^j\,[m_L^0 + (1-\alpha)j]\,p^0$. As always, the monopolist also sells to α H-type consumers, earning αp^j in the first period and $\alpha \sum_{i=0}^{j-1} \delta^{j-i} p^j$ in the subsequent periods. Consumers who had previously bought the product at p_{-1} however, may become a cost of $\alpha\epsilon\left(p_{-1} - p^j\right)$ if $p^j < p_{-1}$. In the following $j-1$ periods, where the prices are sure to fall, H-type consumers will also represent, at each period i, a gain of αp^i and a cost of rebate of $\delta^{j-i}\alpha\epsilon\left(p^{i+1} - p^i\right)$. Finally, future earnings of the monopolist will be postponed by j periods, representing a present value of $\delta^{j+1}\Pi$. Combining these gains and losses, the monopolist's total gain by changing the price from p_{-1} to p^j equals $\delta^j\left[m_L^0 + (1-\alpha)\, j\right] p^0 + \alpha p^j + \mathbf{1}_{p^j < p_{-1}} \alpha\epsilon(p_{-1} - p^j) + \alpha \sum_{i=0}^{j-1} \delta^{j-i}\left[p^i - \epsilon\left(p^{i+1} - p^i\right)\right]$.

Since m_L^0 is the minimum number of L-type consumers in the market for which a monopolist will prefer to lower its prices to p^0 rather than setting it at another value p^j it follows that m_L^0 must satisfy

$$
\begin{aligned}
\alpha p^0 + m_L^0 p^0 - \alpha\epsilon(p_{-1} - p^0) + \delta\Pi = \delta^j [m_L^0 + (1-\alpha)j] p^0 + \alpha p^j + \\
+ \mathbf{1}_{p^j < p_{-1}} \alpha\epsilon(p_{-1} - p^j) + \\
+ \alpha \sum_{i=j-1}^{0} \delta^{j-i}\left[p^i - \epsilon\left(p^{i+1} - p^i\right)\right]
\end{aligned}
$$

Rearranging the terms of this equation and being explicit on the fact that m_L^0 is a function of both p_{-1} and p^j, one obtains

$$
\begin{aligned}
m_L^0\left(p_{-1}, p^j\right) = \frac{\delta^j j}{(1-\delta^j)} - \frac{\alpha}{1-\alpha} - \delta \frac{\Pi}{\theta_L(1-\alpha)} \\
+ \alpha \frac{p^j - \mathbf{1}_{\{p_{-1} > p^j\}} \epsilon\left(p_{-1} - p^j\right) + \sum_{i=0}^{j-1} \delta^{j-i}\left[p^i - \epsilon\left(p^{i+1} - p^i\right)\right] + \epsilon\left(p_{-1} - \theta_L\right)}{\theta_L(1-\alpha)(1-\delta^j)}
\end{aligned}
$$

A monopolist will decide to lower its value to p^0 rather then setting it to another value p^j when the number of L-type consumers exceeds the value $m_L^0(p_{-1}, p^j)$ given by equation above. Now recall that our goal is to find $M_L^0(p_{-1})$, namely, the number of L-type consumers which must be on the market in order to make the monopolist lower the price to p^0. In order for the monopolist to decide to set the price to p^0, setting the price to p^0 must be better than any other option. This means that number of L-type consumers which must be in the market in order to make the monopolist to set the price to p^0 must equal the maximum value of $m_L^0(p_{-1}, p^j)$ when all possible values of p^j are considered. Therefore:

$$M_L^0\left(p_{-1}\right) = \max_j m_L^0\left(p_{-1}, p^j\right)$$

This leads us to the following Theorem:

Theorem 1. *An equilibrium exists such that the monopolist offers the good each period at a price equal to:*

$$p\left(p_{-1}, \alpha, \alpha, M_L\right) = \begin{cases} p^0 & \text{if } M_L \geq M_L^0\left(p_{-1}\right) \\ p^1 & \text{if } M_L^0\left(p_{-1}\right) > M_L \geq M_L^1\left(p_{-1}\right) \\ \vdots & \\ p^{n-1} & \text{if } M_L^{n-2}\left(p_{-1}\right) > M_L \end{cases}$$

where $p^i = \theta_H - \delta^i\left(\theta_H - \theta_L\right)\frac{1+\epsilon(1-\delta)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^i}{1+\epsilon(1-\delta)}$ *and* M_L^i, $i \in \{1, 2, ...n\}$ *is any sequence that satisfies*

$$M_L^j\left(p_{-1}\right) - M_L^{j-1}\left(p\left(p_{-1}, \alpha, \alpha, M_L^j\left(p_{-1}\right)\right)\right) = 1 - \alpha$$

with M_L^0 being implicitly defined by the conditions below for an integer n:

$$
\begin{aligned}
(1-\alpha)(n-1) \leq M_L^0(p_1) \leq (1-\alpha)\, n \\
M_L^0(p_{-1}) = \max_j m_L^0\left(p_{-1}, p^j\right) \\
m_L^0(p_{-1}, p_j) = \frac{\delta^j j}{(1-\delta^j)} - \frac{\alpha}{1-\alpha} - \delta \frac{\Pi}{\theta_L(1-\alpha)} \\
+\alpha \frac{p^j - \mathbf{1}_{\{p_{-1} > p^j\}} \epsilon (p_{-1} - p^j) + \sum_{i=0}^{j-1} \delta^{j-i} \left[p^i - \epsilon (p^{i+1} - p^i)\right] + \epsilon (p_{-1} - \theta_L)}{\theta_L (1-\alpha)(1-\delta^j)}
\end{aligned}
$$

Also, in such an equilibrium, high type consumers will make a purchase in state (M_H, M_L) if the price at that state is, in fact, at most $p(p_{-1}, \alpha, \alpha, M_L)$ and low type consumers will make a purchase if the price is θ_L.

Off the Equilibrium Path

The above theorem describes the equilibrium path, but we should take notice of what happens outside of the equilibrium as well. For that, we must consider a general state:

$$S_0 = (p_{-1}, x_{-1}, M_H, M_L)$$

which happens at a period we will, without loss of generality, denote as period 0. If $(x_{-1}, M_H) = (\alpha, \alpha)$, then we can locate the solution in Theorem 1. Alternatively, we have to look at the monopolist's options, and in some way, we must go back to the equilibrium path (we can see this from the proof of Lemma 1, as at some point, the object's price must go to θ_L and everyone makes a purchase, as keeping the price above θ_L forever is unsustainable).

The proof that when facing a state S_0 that is not a part of the equilibrium path, the actors will trend to move back to the equilibrium is shown in Appendix B.

In the next section, we consider the impact of price protection on welfare.

1.5 Welfare Analysis

In this section we study the effects of price protection policies on welfare. As we shall argue in this section, offering price protection policies damage social welfare. This may possibly explain why price protection is not a more universal phenomenon. It may also explain why stores have been stopping to offer such policies.

There are multiple ways to compare the effect of offering price protection on welfare in our model. First, note that the equilibrium share of consumers that take advantage of price protection, ϵ, is exogenous. Therefore, we could compare welfare when $\epsilon = 0$ and $\epsilon = \epsilon^* > 0$, the equilibrium share of consumers who do take advantage of price protection in equilibrium. Alternatively, we can consider how a marginal increase in ϵ affects utilities in our model. Second, notice that we should carefully pick a comparison date: comparing a consumer's continuation utility on the day that prices are lowest in one model, versus on a day where prices are highest in the other would not be a fair comparison.

Our strategy will be to first note the effect of ϵ on the monopolist's profits, and then on consumers' utilities.

Regarding profits, in Subsection 1.4.2 we used a variable Π that represented the present value of a monopolist's profits, starting from the initial period of a cycle. We did not need to expand on it for the purposes of finding the equilibrium, but its expression is straightforward from the model and the equations in that subsection:

$$\Pi = \frac{1}{1-\delta^n}\left[\underbrace{\alpha \sum_{j=0}^{n-1} \delta^j p^{n-1-j}}_{\text{sum of sales each period}} - \underbrace{\epsilon\alpha \sum_{j=1}^{n-1} \delta^j \left(p^{n-j} - p^{n-j-1}\right)}_{\text{sum of refunds each period}} + \underbrace{n\left(1-\alpha\right)\delta^{n-1}p^0}_{\text{sales to L-types}}\right]$$

It can be shown that Π decreases on ϵ (note that n and $p_{\cdot}$ vary with ϵ as well, and thus this is not straightforward), and therefore offering a price protection policy in our model decreases profits.

We now show that consumer's utilities also drop with price protection, in two steps. For the first, that an increase in the length of cycles decreases consumer's expected utility, both for high- and low-valuation types. The intuition behind this is that (i) an arriving consumer's utility is higher the closer he is to the next discount period; and that (ii) longer cycles means that a random consumer will be on average farther away from the next discount period. In the second step, we will show that an increase in ϵ will (weakly) increase the length of cycles. This is intuitive, as a higher ϵ slows down the price decrease, as it makes price drops costlier. Together, we get that an increase in ϵ weakly increases n which itself decreases consumer welfare.

As can be gathered from the demonstration of Equation 1.1, for a given number of periods i before the promotion period, the consumer must expect to pay the same amount (after receiving the rebate) regardless of the probability of receiving the rebate ϵ. Therefore, finding how cycle lengths are related to the rebate probability will give us insight into the average price high type consumers pay, besides telling us the average waiting time for low types. Because our discount rate δ is strictly below 1, both high and low types strictly prefer shorter cycle lengths, as high types would end up paying less on average, while low types would manage to make a purchase earlier.

As for the effect of ϵ on cycle lengths, we must return to Theorem 1, which states that the length of the cycles must satisfy:

$$(1-\alpha)(n-1) \leq M_L^0\left(p^1\right) \leq (1-\alpha)\, n \tag{1.2}$$

$$M_L^0\left(p_{-1}\right) = \max_j m_L^0\left(p_{-1}, p^j\right) \tag{1.3}$$

$$m_L^0\left(p_{-1}, p^j\right) = \frac{\delta^j j}{\left(1-\delta^j\right)} - \frac{\alpha}{1-\alpha} - \delta\frac{\Pi}{\theta_L\left(1-\alpha\right)} \tag{1.4}$$

$$+\alpha\frac{p^j - \mathbf{1}_{\{p_{-1}>p^j\}}\epsilon\left(p_{-1}-p^j\right) + \sum_{i=0}^{j-1}\delta^{j-i}\left[p^i - \epsilon\left(p^{i+1}-p^i\right)\right] + \epsilon\left(p_{-1}-\theta_L\right)}{\theta_L\left(1-\alpha\right)\left(1-\delta^j\right)}$$

From equation 1.2, we see that n is weakly increasing with $M_L^0(p^1)$. In addition, from equation 1.3 and the envelope theorem, it follows that if $m_L^0(p_{-1}, p^j)$ is increasing on ϵ, then so is $M_L^0(p_{-1})$. Therefore, we need to see how m_L^0 behaves with ϵ[4], which we can obtain from equation 1.4. It can be shown that m_L^0 is increasing on ϵ. This means that cycle lengths are weakly increasing on the probability of consumers obtaining price protection.

Since cycle lengths are weakly increasing with ϵ, consumer welfare drops with price protection, as low type consumers have to wait longer, on average, to make their purchases, while high type consumers are paying more on average as their benefit of waiting drops.

1.6 Conclusion

In this paper, we studied the effects of price protection for a durable good sold by a monopolist to rational consumers. Our model can thus be conceptualized as a generalization of the CGS model, where we introduce an additional parameter, ϵ: the fraction of consumers which will exercise their rights for a rebate if prices actually drop. In the presence of price protection policies, the evolution of prices exhibit a behavior which is somewhat different from that predicted by CGS in the absence of such policies. Indeed, we have shown that prices may evolve in two very different regimes, depending on whether the parameter ϵ exceeds or fails to exceed a certain threshold which depends solely on the discount rate, namely

$\epsilon_\delta = \frac{1}{1+\delta}$

This paper focus on the case where the fraction of consumers claiming their right to rebate is low ($\epsilon < \epsilon_\delta$), during which prices evolve in cycles and decrease monotonically within each cycle. This behavior is qualitatively similar to that described by CGS.

[4]Note that ϵ appears implicitly in both prices p_i and profit Π. Also, Π includes n, which is itself a function of ϵ as well.

We show, however, that cycles are longer than those predicted by CGS in the absence of price protection policies.

Indeed, an additional finding of this paper is that social welfare decreases as ϵ grows. In particular, this means that the state of maximum welfare is achieved in the absence of price protection policies ($\epsilon = 0$). This might suggest the need for a regulator to discourage the adoption of such policies. Such need, however, is not of serious concern, since price protection policies decrease the welfare of both the consumers and the monopolist — the one who would decide on such policies.

It is left for future research how such policies fare under competition, and why sellers decide to offer them in practice. One real world confirmation of our welfare result is that Amazon's offering of price protection policies have dropped as its market share increased. In particular, they now offer it for just one product category, television sets. Though this paper does not attempt to explain such a decision, this might be an attempt to avoid consumers returning such products when prices drop, as the costs associated to the returns of these fragile and cumbersome products might outweigh the welfare costs of price protection.

Finally, there is also the issue of credit card companies offering price protection for purchases made using their cards. These companies take care of paying the rebate out of their own earnings. Since both the firm and the consumers can observe that this policy exists, the pricing functions in those markets would be the same as those studied in this paper - the pricing equations would be unchanged. However, the CGS equilibrium cycle length under such policies would differ. The process for solving this problem would be similar to the ones seen in this paper, by excluding the rebates from the firm's profit function.

Appendix

1.A Price Protection Probability Threshold

The majority of this paper considers a pure strategy equilibrium where prices are decreasing. As mentioned in Section 1.4, for that to happen it is sufficient that the following condition holds:

$$\epsilon < \frac{1}{1+\delta}$$

To see that, note that we have shown in Section 1.4.1 that if prices are decreasing, then they must follow the following structure:

$$p_j = \theta_H - \delta^j \left(\theta_H - \theta_L\right) \frac{1 + \epsilon\left(1-\delta\right)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^j}{1+\epsilon\left(1-\delta\right)}$$

Where we know that, for every j:

$$p_j \geq p_{j-1}$$

This implies that:

$$\theta_H - \delta^j \left(\theta_H - \theta_L\right) \frac{1 + \epsilon\left(1-\delta\right)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^j}{1+\epsilon\left(1-\delta\right)} \geq \theta_H - \delta^{j-1} \left(\theta_H - \theta_L\right) \frac{1 + \epsilon\left(1-\delta\right)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j-1}}{1+\epsilon\left(1-\delta\right)}$$
$$\delta\left[1 + \epsilon\left(1-\delta\right)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^j\right] \leq \left[1 + \epsilon\left(1-\delta\right)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j-1}\right]$$

$$
\begin{aligned}
\delta\left[\epsilon\left(1-\delta\right)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j}\right]-\epsilon\left(1-\delta\right)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j-1} &\leq 1-\delta\\
\epsilon\left(1-\delta\right)\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j-1}\left[\delta\left(\frac{-\epsilon}{1-\epsilon\delta}\right)-1\right] &\leq 1-\delta\\
\epsilon\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j-1}\left(\frac{-\epsilon\delta-1+\epsilon\delta}{1-\epsilon\delta}\right) &\leq 1\\
\epsilon\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j-1}\left(\frac{-1}{1-\epsilon\delta}\right) &\leq 1\\
\left(\frac{-\epsilon}{1-\epsilon\delta}\right)^{j} &\leq 1
\end{aligned}
$$

If j is odd, then the left hand side is negative, and the inequality is immediately satisfied. If j is even, then it must be that:

$$
\begin{aligned}
\left(\frac{\epsilon}{1-\epsilon\delta}\right)^{j} &\leq 1\\
\left(\frac{\epsilon}{1-\epsilon\delta}\right) &\leq 1\\
\epsilon &\leq 1-\epsilon\delta\\
\epsilon &\leq \frac{1}{1+\delta}
\end{aligned}
$$

1.B Out of Equilibrium Path

We shall note that there are 3[5] possible classes of response for the monopolist when faced with an arbitrary state, $S_0 = (p_{-1}, x_{-1}, M_H, M_L)$, which are chosen by backwards induction:

- Low price
- High price
- Intermediate price

[5] One could include the first class in the third, but I believe this division is more intuitive.

Figure 1.2 illustrates the following discussion. Starting from the middle, the possible paths always end in one of three possibilities. Either blue, where the monopolists charges $p = \theta_L$ and therefore the cycle restarts in the next period; or green, where the state of the world becomes similar to the one seen in Theorem 1; or finally orange, where there is no liability due to no one buying the previous period, there are 2α high type consumers on the market and an arbitrary number of low type consumers in the market.

In the first, the monopolist sets the price to be θ_L, which clears the market (of both high and low types). Intuitively, this happens when p_{-1}, x_{-1} and M_H are relatively low, and M_L is relatively high, as this would lead to a low liability for reducing prices, together with a high ratio of low to high types in the market, making it profitable to bring the purchase of the low type consumers to the present. In this case, at period 1 we are back to the equilibrium path, in the beginning of a cycle.

The second class of responses is to set some price $\overline{p}$ high enough such that no consumer makes a purchase at period 0, so that the next state of the world is

$$S_1^{high} = (\overline{p}, 0, M_H + \alpha, M_L + (1 - \alpha))$$

At this point, there is no incentive not to sell, given that there's no rebate to pay, and so let's call the next price, at period 1, $\overline{p}_a$[6]. All high type consumers make a purchase now, so that the new state is

$$S_2^{high,inter} = (\overline{p}_a, M_H + \alpha, \alpha, M_L + 2(1 - \alpha))$$

Note that now the cost of reducing prices may be high again, so that the monopolist may be reluctant to sell to anyone this period. So we are back to three possibilities, low price, high price and an intermediate price. The first leads to ending the cycle,

[6] If M_L is high enough compared to M_H, then this price would be θ_L and the cycle restarts

while the third leads to a Theorem 1 state. The second is the orange state with no liability and 2α high type consumers in the market, which we will consider later in this section more generally:

$$S_3^{high,inter,high} = (\overline{p}, 0, 2\alpha, M_L + 3(1-\alpha))$$

Finally, the third class of responses is to set an intermediate price, where only high type consumers will make a purchase. After setting that price, the state becomes:

$$S_1^{inter} = (p, M_H, \alpha, M_L + (1-\alpha))$$

That state can lead to an arbitrarily large liability, so we are back to the three possible actions of a low price, a high price and an intermediate price. Again, if the first or third paths are taken, the cycle restarts and the state becomes a Theorem 1 state, respectively. In the second path, however, there is now no liability and an intermediate price would lead us to a state equivalent to $S_2^{high,inter}$, which we already saw (though it remains to look at what happens in the orange state, where the next price is high).

So now the only remaining case to consider is when there is no liability, 2α high type consumers in the market and an arbitrary number of low types: $S = (p, 0, 2\alpha, M_L)$. Of course, this case fits in the general S_0 state, except that $x_{-1} = 0$ and $M_H = 2\alpha$, so we can redo the same logic, with this simplifying restriction. The $x_{-1} = 0$ restriction means there is never a case in which the monopolist would choose a high price in such a class of states, removing one of the orange states. The other orange state, on $S_3^{inter,high,inter} = (p, M_H + \alpha, \alpha, M_L + 3(1-\alpha))$ can be seen to not occur infinitely as on $S_1^{inter} = (p, M_H, \alpha, M_L + 1 - \alpha)$ it cannot be optimum to choose a high price as M_H increases.

Figure 1.2: Out of Equilibrium

Start at the center of the figure, with generic state $S_0 = (p_{-1}, x_{-1}, M_H, M_L)$. Blue states represent the end of the cycle, as prices go to θ_L. Green states represent converging to the standard cycle from Section 1.4.

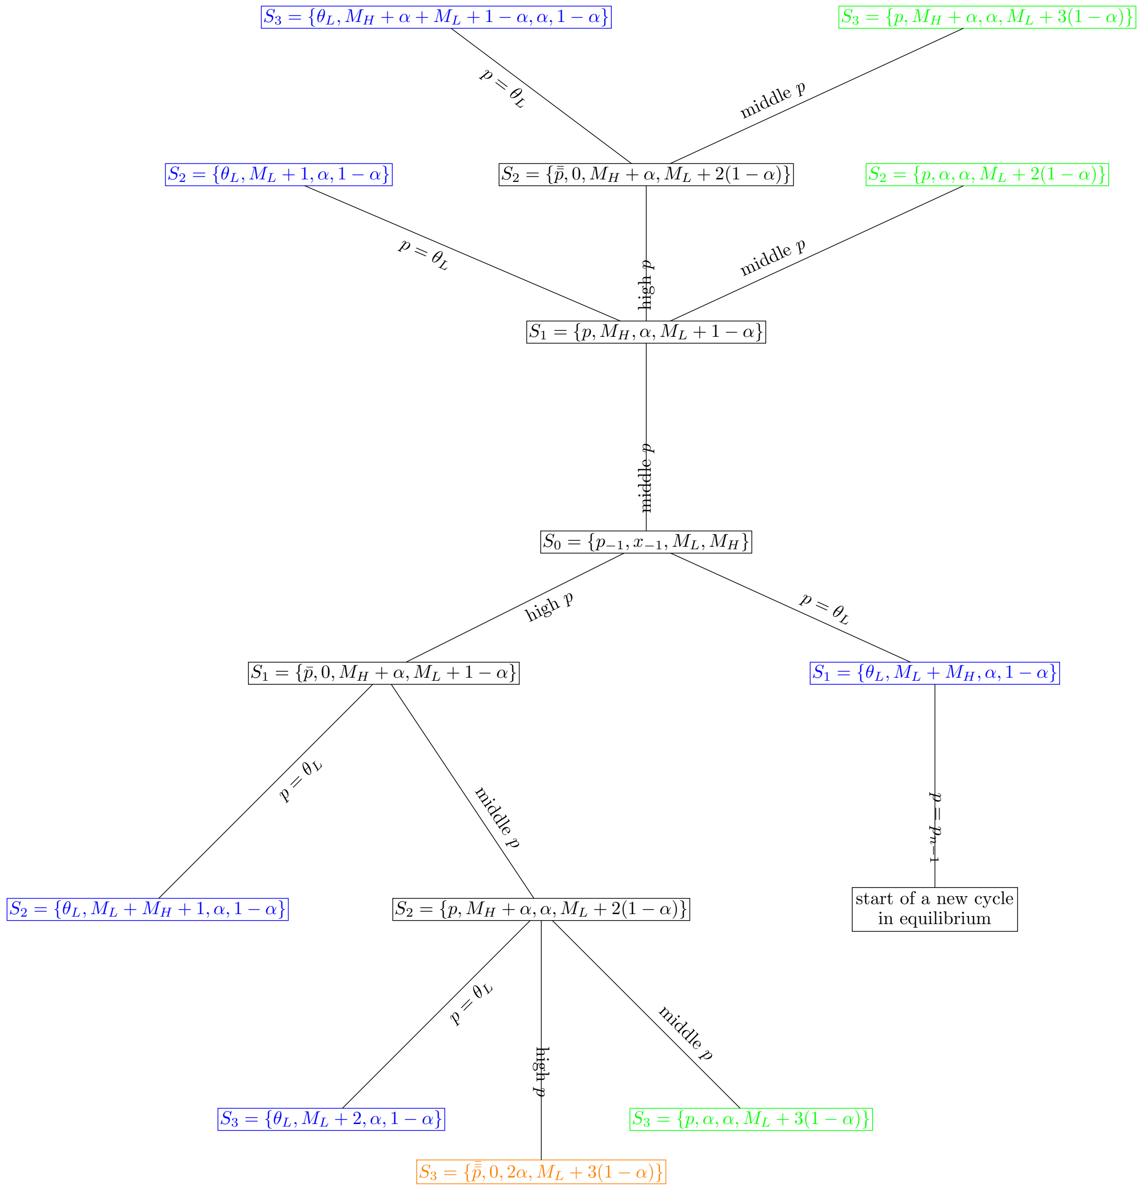

Chapter 2

Local Measures of Investor Attention Using Google Searches*

2.1 Introduction

Traditional asset pricing models like the CAPM assume that investors freely obtain all relevant information for their investment decisions. However, that assumption requires investors to be able to pay attention to a very high amount of data, which seems unrealistic. In fact, retail investors hold stocks from only a handful of firms (see, for example, Barber and Odean [2000], Gargano and Rossi [2016]). It is therefore useful to understand whether this sparsity occurs due to market structure, such as transaction costs, or because investors actually prefer to focus on a small number of shares. There is a large and growing literature suggesting that attention is limited (e.g., Merton [1987], Sims [2003], Peng and Xiong [2006], Gennaioli and Shleifer [2010])[1].

*I am grateful to my advisor, Stephen Morris, as well as Wei Xiong and Jakub Kastl for support and comments. I also thank Ioannis Branikas and Dmitry Mukhin for helpful insight.

[1]That holds true even for mutual funds, as seen in Chen et al. [2017] where web visitor data from EDGAR is used to study investor attention.

It is of interest, therefore, to find a measure of attention in order to separate the two potential reasons for the low diversification in portfolios. This paper considers a popular measure of investor attention and tests it against a major stylized fact from the literature.

The recent literature on investor attention, more specifically, has used Google searches for companies' ticker symbols as a measure of investor attention. As explained in Da et al. [2011], the search volume index (SVI) for tickers provided by Google Trends is correlated with other usual measures of investor attention (e.g., abnormal returns and news coverage), while being available in a more timely fashion. Drake et al. [2012] use the SVI for tickers to estimate the evolution of investor attention around earnings announcements. Chen [2017] uses global searches for market index tickers to test for home bias in international investment. Reyes [2018] uses ticker searches in the United States as a proxy for investor interest in Merger and Acquisition performance. In general, Stephens-Davidowitz and Varian [2014] provide a guide to using Google Trends[2], along with other resources, for using Google search data for research.

However, any good measure of investor attention should also follow common investor behavioral biases. One such longstanding bias is that investors tend to put more weight of their resources in companies that are close to home, following what is called *local bias*[3]. We should therefore expect that people are more likely to search

[2] trends.google.com

[3] This leads to a divergence from the CAPM and a sub-optimal portfolio allocation, as households consider distance between their location and firms' headquarters as a factor in their decisions. This phenomenon has been observed both between and inside countries (e.g.,Coval and Moskowitz [1999], Zhu [2002]).

There are many potential explanations for such a bias. First, spacial closeness can be related to ease of information acquisition, and therefore households may be better at picking stocks that are close to home. However, recent studies such as Seasholes and Zhu [2010] point to no extraordinary gains from such investments. Demarzo et al. [2004], Hong et al. [2014] also suggest that local bias may arise as a form of hedging against keeping-up-with-the-Joneses fears, as investing in local stocks is an easy way to coordinate portfolios and share the same wealth shocks as nearby households. Familiarity heuristics with the company (e.g., Huberman [2001]) and latent subjective expectations about the prospects of a city (as recently pointed out by Branikas et al. [2018]) seem to be a major factor in generating local bias.

for the tickers of companies whose headquarters are near them[4]. However, that is not what I find. There's little evidence that ticker searches follow this bias, and while there may be a few plausible explanations for that, other regressions presented in this paper seem to reject them.

My methodology uses metropolitan Google Trends data to test whether Google searches for stock tickers for companies satisfy local bias. I focus on the S&P 500 stocks and remove the ones whose tickers have ambiguous meaning, as searches for them may not represent investor interest in the company. Individuals are grouped by metropolitan region, as made available by Google, and 65 major metropolitan regions are used in this research. The main idea of the regressions performed in this paper is to test whether distance between company A's headquarters and individual's address is related this individual's search for A's stock ticker.

I find a negative but weak evidence of ticker searches being affected by local bias. Performing the regression monthly, results are ambiguous, with some months even suggesting that investors are more likely to search firms farther away. While one possible explanation would be that investors near a company do not need to research the company online (as they can gather information with local means), constraining the regression to households farther than a minimum distance seems to reject this hypothesis. This possibility is also rejected by an alternative measure of investor interest which I propose, and which does conform to local bias.

A more natural explanation for this conflict is that many investors are unaware of the company's ticker, and therefore search for its popular name when looking to invest in it. This cannot be considered a precise measure of investor interest, however, since regular consumers may be looking for the company's name to research its service or products, or perhaps consumer support, and search for it for that reason. As a solution, those searches can be filtered out by adding the word "stock" to the

[4]Naturally, employees might be also driving some of the searches for the companies they work for.

name of the company. Google Correlate[5] suggests that, for popular stocks, this measure is highly correlated with the ticker search, and therefore seems to be a natural replacement. I find that distance becomes a highly significant factor in explaining changes in investor interest, and regressions with it have a much improved Akaike Information Criterion (Akaike [1974]).

Using this "name plus stock" (henceforth referred as NPS) measure, I estimate that doubling the distance between a household and a company's headquarters may decrease interest in the company's stock by around 20%. As a visualization exercise, this suggests that Boston, MA households would be around 20% less interested in investing in Citigroup (headquartered in Manhattan, NY) if the firm were to be located in Washington, DC.

To test for bias arising from consumer's knowledge of the company's ticker, I also performed the main regressions on different subsets of the sample, considering size and industry. It is interesting to note that, when looking at larger firms, or at firms in industries that have large advertisement expenditure, both ticker and "name plus stock" searches are being affected by distance. However, for small firms, or firms not in those industries, only the NPS measure is relevant. That suggests that not enough retail investors know the ticker of a company to overcome noise from other search users for those terms.

The results of this paper, together with the procedure described in it, may be useful in testing other local investor interest stylized facts, such as whether local investors react quicker to news about a company. Similarly, they can be used to test the relation between product familiarity and investment, as in Keloharju et al. [2012]. Additionally, it could be used to estimate the effects of advertisement on investor interest on a local setting, such as is done by Branikas [2019]. This data

[5] google.com/trends/correlate

may also be used together with EDGAR and Bloomberg visitor data, to compare the attentions of retail versus institutional investors[6].

The organization of the paper is as follows: The next section details the data acquisition process, and how it was used. Section 3 provides the regression estimates for the regression on ticker searches and some robustness checks and the alternative search formula, while section 4 tests the regression on selected subsamples as a robustness check. Section 5 concludes.

2.2 Data

2.2.1 Google Trends

In this section, I describe how Google transforms search volumes into the search indices it provides, and what this paper does to the index numbers to obtain a measure of each stock's search popularity. As explained by Stephens-Davidowitz and Varian [2014], Google has made normalized information on search volumes available on its Google Trends website. While they do not provide the actual number of searches and limit the ways in which users can obtain that data, for example, by limiting the comparison to at most 5 search terms at a time, it is possible to work with the data in order to obtain a measure that is proportional to search volumes per location.

While there are different forms of requests that can be done to the Trends website, this paper focuses on making requests for small numbers of terms on monthly ranges, and saving the returned SVI breakdown by metropolitan regions. The SVI that is returned by Trends is in the form of integers between 0 and 100 denoting the relative popularity of each of the terms inside each region.

[6] EDGAR and Bloomberg visitor statistics are used as proxies for institutional investor attention; it's likely that both types of investors use both Google and EDGAR/Bloomberg for information, although in different proportions.

To transform the universe of searches into the SVI, Google takes a number of steps. To start with, it takes a sample of all searches made[7]. From this sample, let the number of searches for the term i[8] in location j[9] during time period t[10] be denoted by $x_{i,j,t}$. If $x_{i,j,t}$ is below some unpublished threshold, Trends considers it to be zero.

While obtaining this $x_{i,j,t}$ would be ideal, when returning information SVI as a time series, Trends normalizes this number in two different ways: (i) with respect to all searches in the same period, in the same location, and (ii) with respect to the most searched triple (i, j, t) (after the first normalization) submitted in the request. However, Trends also allows us to make comparison searches inside regions, between terms. For example, on a comparison search for the terms "New York" and "California" for the year 2017, Trends returns results by state, as shown in Figure 2.2a, which suggests a strong local interest in searches for the names of states. Similarly, Figure 2.2b also shows a local interest effect, in this case for searches of "New Mexico" and "California", but suggests that the effect is quickly overwhelmed by other factors as distances grow larger, suggesting the use of a decreasing slope function, such as log, transforming the distance. In this setting, the numbers provided for each state represents the fractions, for each search term i in the set of searched terms in I, the SVI $y_{i,j,t,I} = \frac{x_{i,j,t}}{\sum_{\tilde{i} \in I} x_{\tilde{i},j,t}}$.

Since we are focusing on hundreds of stock tickers, the framework provided by Trends seems insufficient for a relevant comparison, as the denominator changes with each group of stocks. As a workaround, however, we can choose a benchmark search term, i_{bench}, in order to obtain search volumes relative to this benchmark. We can

[7]This sample seems to change regularly, and so we are able to test the accuracy of the index by querying the same search parameters multiple times.

[8]For example, a stock ticker such as "AAPL".

[9]For example, a state or a metropolitan region.

[10]A date range.

therefore calculate a new standardized index for investor attention:

$$A_{i,j,t;i_{bench}} = \frac{y_{i,j,t,I}}{y_{i_{bench},j,t,I}} = \frac{x_{i,j,t}}{x_{i_{bench},j,t}}$$

which is independent of the set of search terms chosen. However, it is now relative to the attention directed at the arbitrarily chosen benchmark.

In order to keep the interpretation of the index as a fraction and to make it independent of the benchmark, we normalize it with respect to the sum of all indices $A_{i,j,t;i_{bench}}$[11]:

$$a_{i,j,t} = \frac{A_{i,j,t;i_{bench}}}{\sum_{\tilde{i}} A_{\tilde{i},j,t;i_{bench}}} = \frac{\frac{x_{i,j,t}}{x_{i_{bench},j,t}}}{\sum_{\tilde{i}} \frac{x_{\tilde{i},j,t}}{x_{i_{bench},j,t}}} = \frac{x_{i,j,t}}{\sum_{\tilde{i}} x_{\tilde{i},j,t}}$$

One problem that we face is that Trends replaces the true $y_{i,j,t,I}$ with zero when the number of searches is below some threshold. This choice makes it is somewhat common, especially in smaller regions and for less popular search terms, that these fractions may be indefinite. To solve this problem, we add a small number (0.1) to the observations which return zero values. The choice of a popular benchmark stock, such as Apple's (whose ticker is AAPL), alleviates, but does not eliminate, such occurrences. Additionally, the focus on major metropolitan areas reduces the incidence of zero values.

Finally, 30 samples are drawn for each request, to reduce the problem of Trends' sampling on its provided data. While we can observe differences between individual samples, the results seem to converge.

The following summarizes the process for obtaining and arranging the Google Trends data:

[11] Note that, except in how Google Trends rounds numbers in the data collection part, the resulting $a_{i,j,t}$ will be independent of the choice of the benchmark search term

Picking the relevant stocks

From more than 800 companies listed in the S&P 500 index between 2004 and 2016, 385 remained after removing problematic tickers.

The tickers manually chosen to be removed fall on two broad categories. The first one includes tickers for which searches for their strings would probably not be mainly due to interest in the firm (for example, searches for “A”, “ACT” or “TOY”), which would probably bias our estimates towards zero. The second one includes tickers that resemble too much a popular name of the firm, and therefore we would not be capturing investor interest (for example, searches for “AOL” or “NYT”), but mainly consumer interest, which would probably bias our estimates away from zero. The existence and lack of clear solution towards dealing with these problematic tickers is also a reason against using SVI for tickers as a measure of investor interest, as this measure cannot be used for gauging investor interest for a large number of firms.

Splitting stocks into groups

Since Trends limits the number of search terms per query to 5, we must choose a benchmark (for our exercise, “AAPL”) and include it in every single group[12]. Since there are 384 remaining companies after removing the benchmark, it is convenient to separate the companies into 96 groups of 5 (4 random stocks plus the benchmark). For lack of a clear choice, this was done alphabetically.

Downloading and adjusting the indices from Google Trends

A Python script was used to request the search indices for each of the 96 group of shares, at each of the 144 months, for each of the 30 draws, generating 414,720 tables. The samples were averaged along the draws and increased by 0.1 to remove

[12] An alternative, for the case in which the sampling is overwhelmed with zeroes, would be to add multiple benchmarks, at different thresholds of stock popularity. This is planned as a future step for this paper as a robustness check.

the remaining zeroes. From the resulting numbers, Apple's index was standardized to 1 and every other stock's index was considered in proportion to the benchmark, allowing us to join together the 96 groups into one single file for each month.

At the end of this process, we are left with relative search popularity indices for stock tickers in each month, which according to the current literature would be a proxy for investor interest.

2.2.2 Stock Characteristics

Most stock characteristics are as used in Branikas et al. [2018].

Monthly stock prices and calculated returns were drawn from CRSP, while firm accounting variables were obtained from Compustat, quarterly. Financial characteristics are calculated at a monthly frequency, and comprise of stock price, market capitalization, book-to-market ratio, turnover ratio (defined as volume over number of shares outstanding), momentum (defined as past annual return), volatility (defined as the standard deviation of monthly returns over the past year), profitability (defined as in Novy-Marx [2013] as the ratio of past annual gross profits to assets), investment (defined as the past annual growth rate of assets) and past annual sales. Observations with missing data are dropped.

The model assumes investors (search users) may be influenced in their searching decisions by stock prices in that month and risk factors in the previous month[13]. Firm's balance sheet information is lagged so we can assume investors are aware of them. In particular, accounting variables from fiscal year $t-1$ are matched with searches from July of year t until June of year $t+1$ as in Fama and French [1992].

Finally, firms are separated into 17 categories, based on Kenneth R. French's classification[14] to control for visibility to households of different industries.

[13] An argument could be made that households are making the searches specifically to learn that information, and therefore it could not affect their decisions. However, we find that these factors are significantly correlated with searches, and it suggests that searchers already have some knowledge.

[14] http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/det_17_ind_port.html

2.2.3 Metropolitan Demographics

This paper focuses on 65 major metropolitan regions. While Google Trends seems to use Nielsen's Designated Market Areas (DMAs), the exact demarcations of such regions is proprietary, and therefore precise demographic information is unavailable for research. As a workaround, DMAs were matched with Metropolitan Statistical Areas as defined by the US Census Bureau, which allows us to proxy demographics from government agencies. Unemployment and population numbers were taken from the Bureau of Labor Statistics, while income was drawn from the Bureau of Economic Analysis.

2.3 Empirical Results

2.3.1 Regression

In this section, we test the existence of local bias effects on the standard literature measure of investor attention, the SVI for company tickers. The main set of regressions tests the relation between search interest on the ticker of a company and the distance between the investor and the company's headquarter. That can be seen by the following equation:

$$w_{i,j,t} = \alpha + \beta dist_{i,j} + \gamma \cdot \mathbf{X}_{j,t} + \delta \cdot D_{i,t} + \theta_t + \epsilon_{i,j,t}$$

where $w_{i,j,t}$ is the interest measure calculated from Google Trends data, $dist_{i,j}$ is the distance between region i and firm j (either linear or log of distance), $\mathbf{X}^{fin}_{j,t}$ is the set of financial characteristics for firm j in time t (following the lagging procedure described previously), $D_{i,t}$ is the set of demographics characteristics for region i in time t, and θ_t is a time fixed effect.

The following subsection presents results of this regression (with both linear and log of distance, with subsets of the controls above). The next one uses my alternative search formula.

2.3.2 Local Bias Effects on Ticker Searches

As can be seen in panel A of table 2.2, the relation between ticker searches and interest is negative, but not significant, especially when controls are added. Numbers are small, as interest is represented as proportional to interest in Apple's stock (our numeraire). Since Apple's ticker is about 50 times as popular as the average ticker in our sample, that leads to average interest being around 2%. The regressions vary according to whether the log of distance or distance itself is being used, and whether controls for demographics or stock characteristics are included.

There are a few factors that could be lowering this estimate, and for which we can test. First, perhaps some cities with a large number of professional investors may be affecting these numbers. Removing New York from this sample does not make the coefficients of interest any more significant, and in fact, it makes it easier to reject that they are negative. This can be seen in panel B of table 2.2, which replicates the previous regressions, while removing the metropolitan region that contains New York City.

Second, it could be that investors extremely close to the target company are not searching for it online, as they obtain private information from other sources. As seen in panel C of table table 2.2 this is also not a major factor, as restricting the regression to people far from each target stock, for distances of one hundred miles does not increase the t-stats of the relevant estimates. Additionally, given that local bias would lead investors to being more familiar with nearby companies, and therefore being more likely to know their stock's tickers, the next section helps reject this possibility.

2.3.3 Alternative Search Formula Using NPS

This section proposes an alternative search formula for potential retail investors which is consistent with local bias. Namely, it considers investors searching for the popular name of the stock, followed by the word "stock" ("name plus stock", hereby NPS). There are two main potential faults with this measure. The first one is that it might be difficult at some point to ascertain what is the company's popular name, especially from the perspective of an investor. For example, the Coca Cola Company could be represented as Coca-Cola or Coke. Still, the weekly correlation between the two options ("Coca-Cola stock" against "Coke stock") is close to 92% according to Google correlate, and so it likely makes little difference. However, the ticker measure isn't useful in measuring interest for this company, as searches for KO has multiple other meanings that are unrelated to the soda company. A similar problem occurs for investor interest in The New York Times. Its stock ticker, NYT, is also used by the readers of the newspaper, and therefore investor interest is being confounded with consumer interest.

In following NPS, we change the benchmark from ticker searches towards "Apple stock", from the previously used "AAPL". According to Google correlate, the weekly correlation between the two search terms since 2004 is over 60%, which may seem low, but is many times higher than the correlation between ticker searches and any of the other measures of investor attention in Da et al. [2011], namely abnormal returns, turnover, media coverage and measures of investor sentiment. This correlation, however, drops for less popular companies, as their tickers may not be as well known[15].

Table 2.3 repeats the regressions but using interest in NPS rather than stock tickers. As can be seen, except in the case of linear distance without any controls, the t-statistics of the main regressors are vastly improved. Figure 2.2 represents the

[15] In fact, search volumes for "AAPL" are higher than search volumes for "Apple stock", but this seems to be part of the exception, and NPS is more popular than ticker searches in general.

monthly t-statistics for the coefficients of the log of distance, for the regressions with full controls.

Similarly, Figure 2.3 represents the monthly marginal effect for a change in distance, for the same regression. We can see that log of distance is significant for some periods, in explaining interest in tickers, but not in most, while it is significant in explaining interest in NPS throughout the sample period, and consistently with a much larger (negative) effect.

While the coefficients in tables above are not easy to interpret, we provide in table 2.4, for the same regressions, the average effect on interest of a change of one standard deviation in the distance of a company, with relation to a household, while keeping everything else constant, for comparison between ticker and NPS effects. As can be seen, a change in one standard deviation of distance, or around 750 miles, decreases ticker searches by somewhat over 1%, while it decreases NPS searches by around 10%. Asterisks represent the significance of the underlying coefficient, as in the previous tables.

2.4 Local Interest in Selected Subsamples of Stocks

This section considers separating the sample in different forms, and studying the local bias effects on searches in different company groupings. The main results can be seen in table 2.5, where the regressions for company interest on log of distance between the company's headquarters and the search location are provided, for eight different subsamples.

Panel A on Table 2.5 shows the regressions for the sample slit according to size. As expected, for the top 25% of companies, both searches for tickers and searches for

NPS exhibit local bias, while for the bottom 25%, only searches for NPS do so. This can be explained by three reasons, as explained below.

First, the tickers of larger companies are probably more well known and meaningful, and therefore retail investors are more likely to use them for searches. This explanation is also helpful in strengthening the other two, below, as the size of the company increases both the number of users interested in it, and the fraction of users that knows its ticker.

Second, larger companies are searched more in general, and therefore their Google Trends index is less likely to be rounded or truncated to zero. We can see that local bias is also much less significant for NPS in smaller firms, the t-statistic of its standard error dropping from -5.34 to -1.76. However, this is still almost significant and shows that users are searching more for companies close to home.

Thirdly, for all companies' search terms, there are searches made from both users interested in the company, and from users interested in something else. For smaller companies, the number of users in the second group may overwhelm the first, and therefore the searches for tickers is not meaningful. The fact that searches for NPS do exhibit local bias, however, shows that this measure is not suffering from this problem.

Panel B on Table 2.5 has the sample slit according to the company's sector. Considering French's 17 industry sectors (as explained in the link at footnote 14), the companies in the five industries with highest advertisement expenditure are assumed to be more well known to home investors, while the remaining companies are less well known. The first group includes companies in the food, clothes and drug industries, together with automobile and some other services and products.

As expected, the results are similar to before, and ticker searches are significantly dependent on distance for the firms in high advertisement expenditure sectors, while not for the remaining firms. That suggests that not enough searches are being done for tickers of the remaining companies, and therefore the resulting search index is being truncated too often. However, searches according to the NPS formula are still showing a significant local bias effect.

2.5 Conclusion

Measuring investor interest can be useful both for predicting stock movements and for understanding investor behavior. Google Trends seems to provide user data on a timely and fine manner, and thus could prove to be a useful resource in household investment research. However, for each information gathering objective, there are multiple ways an investor can formulate a search query. While a company's ticker is a straightforward solution, it is far from perfect since tickers are chosen to be unique among the universe of tickers, but not in the universe of potential searches, and therefore many stock tickers must be dropped from analysis.

This paper presents an additional problem with searches for tickers, as they do not present the local bias effects that we have come to expect from decades of household finance research. Therefore, at least on a local level, ticker searches do not seem to reflect investor interest. While a simple explanation for that would be that local investors do not need to search online for stock information, the fact that this paper shows that they are searching for the company plus the word "stock" suggests the opposite. In fact, it seems likely that investors, being more focused on local stocks, would be more familiar with their tickers than distant investors, and therefore more prone to perform ticker searches. Since that does not occur, this suggest the NPS formula is more appropriate as a measure of investor interest.

Figure 2.1: Google Trends Figures

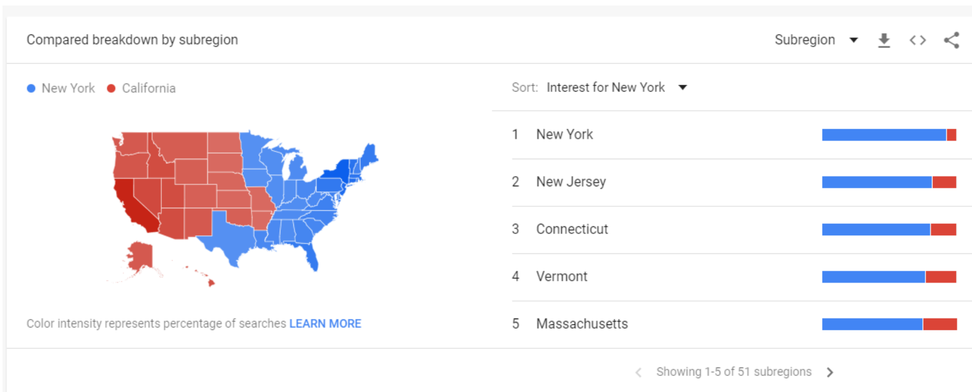

(a) Local bias in searches for New York and California

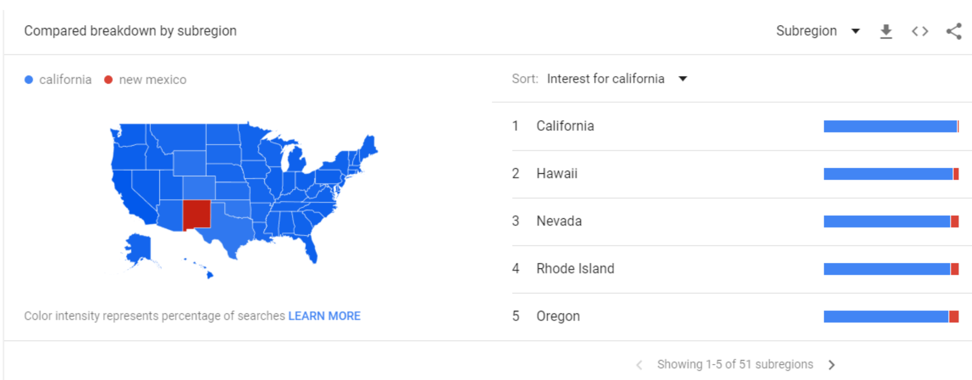

(b) Local bias in searches for New Mexico and California

Table 2.1: Summary Statistics

This table presents summary statistics for the variables used in the analysis. Monthly stock prices and calculated returns were drawn from CRSP, while firm accounting variables were obtained from Compustat, quarterly. Financial characteristics are calculated at a monthly frequency, and comprise of stock price, market capitalization, book-to-market ratio, turnover ratio (defined as volume over number of shares outstanding), momentum (defined as past annual return), volatility (defined as the standard deviation of monthly returns over the past year), profitability (defined as in Novy-Marx [2013] as the ratio of past annual gross profits to assets), investment (defined as the past annual growth rate of assets) and past annual sales. Observations with missing data are dropped.

Variable	Mean	Std. Dev.	Min	Max
		Panel A: Demographics		
Income per capita (USD/capita)	42.6	7.4	27.4	79.2
Population (million)	2.6	3.1	0.1	2.0
Unemployment (%)	6.4	2.1	2.4	15.4
		Panel B: Stock Characteristics		
return	0.01	0.1	-0.8	2.6
book-to-market	0.5	2.4	-234.3	53.1
turnover	0.3	0.2	0.00004	4.9
momentum	0.1	1.0	-1.0	137.7
volatility	0.08	0.05	0.02	0.8
profitability	0.29	0.23	-2.5	1.5
investment	0.1	0.4	-0.8	10.1
size (billion USD)	25	49	0.001	750
		Panel C: Interest and Distance		
interest (for tickers, %)	0.29	1.9	0	91
distance (in miles)	1100	760	11	5100

Table 2.2: Interest in stock tickers

This table presents the regression results of investor interest in log of distance and distance. The dependent variable is *search index for stock tickers*. The independent variables are *distance* and *log of distance*. The measure of investor interest here is the search index for stock tickers, which is 100 for AAPL and all other tickers are represented in proportion to that. Columns (1)-(2) focus on estimating the effect of linear distance, while (3)-(4) consider log of distance. (1) and (3) provide controls for stock characteristics and industry fixed effects, while (2) and (4) include as well controls for demographic characteristics of the searcher's location. AIC is the Akaike information criterion. Panel A includes the full sample, panel B removes searches made inside the New York City metropolitan area, and panel C removes searches made for companies within 100 miles of the searcher, to remove consumers who get extremely local information from the company. In parentheses, t-statistics based on standard errors clustered at each of the 65 metropolitan levels are provided.

	(1)	(2)	(3)	(4)
Panel A: Full sample				
distance	-0.00044	-0.00042		
	(-0.64)	(-0.60)		
log_dist			-0.018	-0.019
			(-1.81)	(-1.80)
Demographics controls	No	Yes	No	Yes
Stock characteristics	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	2,887,285	2,887,285	2,887,285	2,887,285
R-squared	0.161	0.161	0.161	0.161
AIC	1.10e+07	1.10e+07	1.10e+07	1.10e+07

The above controls' listing is applicable to all panels in this table

	(1)	(2)	(3)	(4)
Panel B: Excluding NYC				
distance	-0.00035	-0.00033		
	(-0.51)	(-0.47)		
log_dist			-0.018	-0.018
			(-1.70)	(-1.69)
Observations	2,842,646	2,842,646	2,842,646	2,842,646
R-squared	0.161	0.161	0.161	0.161
AIC	1.10e+07	1.10e+07	1.10e+07	1.10e+07

	(1)	(2)	(3)	(4)
Panel C: Searches further than 100 miles				
distance	0.00023	0.00027		
	(0.34)	(0.39)		
log_dist			0.0054	0.0057
			(0.47)	(0.48)
Observations	2,801,308	2,801,308	2,801,308	2,801,308
R-squared	0.158	0.158	0.158	0.158
AIC	1.10e+07	1.10e+07	1.10e+07	1.10e+07

Table 2.3: Interest in name plus "stock" (NPS)

This table presents the regression results of investor interest in log of distance and distance. The dependent variable is *search index for name of company and "stock" (NPS)*. The independent variables are *distance* and *log of distance*. In parentheses, t-statistics based on standard errors clustered at each of the 65 metropolitan levels are provided. The measure of investor interest here is the search index for the popular name of the company plus the word "stock" (NPS), which is 100 for AAPL and all other tickers are represented in proportion to that. Columns (1)-(2) focus on estimating the effect of log of distance, while (3)-(4) consider linear distance. Columns (1) and (3) control for stock characteristics and industry fixed effects, while (2) and (4) include as well controls for demographic characteristics of the searcher's location. AIC is the Akaike information criterion.

	(1)	(2)	(3)	(4)
distance	-0.0020	-0.0022		
	(-3.62)	(-3.84)		
log_dist			-0.061	-0.063
			(-4.96)	(-5.02)
log_size	0.20	0.20	0.20	0.20
	(34.1)	(34.2)	(34.6)	(34.6)
bk2mkt	0.010	0.010	0.010	0.010
	(6.36)	(6.32)	(6.13)	(6.10)
turnover	0.41	0.41	0.41	0.41
	(37.0)	(37.1)	(38.6)	(38.7)
mom12	0.010	0.010	0.010	0.010
	(8.70)	(8.69)	(8.69)	(8.69)
vol12	1.31	1.31	1.33	1.33
	(28.2)	(28.3)	(28.5)	(28.5)
invest	-0.053	-0.053	-0.053	-0.053
	(-19.0)	(-18.9)	(-18.8)	(-18.7)
profit	-0.11	-0.11	-0.098	-0.098
	(-4.81)	(-4.79)	(-4.55)	(-4.53)
Demographics controls	No	Yes	No	Yes
Month FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
Observations	2,698,230	2,698,230	2,698,230	2,698,230
R-squared	0.053	0.053	0.055	0.055
AIC	4.90E+06	4.90E+06	4.90E+06	4.90E+06

Table 2.4: Estimated effect from change of distance in interest

The table below provides interpretable coefficients for the effect of a change in distance on the change in interest. For linear distance, this is calculated as the coefficient multiplied by the standard deviation of distance, divided by the average of the interest, and therefore represents the interest change caused by a increase in distance by one standard deviation. For log of distance, it is calculated as the coefficient divided by the average of the interest, and represents the elasticity from an increase in distance.

		Average effect on ticker	Average effect on NPS
log_dist	stock characteristics	-5.7%	-26.6%
	stock characteristics and demographics	-5.8%	-28.0%
distance	stock characteristics	-1.5%	-9.8%
	stock characteristics and demographics	-1.2%	-10.9%

Table 2.5: Splitting the sample

This table presents the regression results of investor interest in log of distance, with full controls. In parentheses, t-statistics based on standard errors clustered at each of the 65 metropolitan levels are provided. The measure of investor interest here is standardized to be 100 for Apple Co., and for other companies it proportionally represents searches for tickers in the odd columns, while the even columns consider searches for company name plus stock (NPS). On panel A, columns (1) and (2) focus on the top 25% largest companies by revenue, while (3) and (4) represent the smallest. On Panel B, columns (5) and (6) represent the companies in the top 5 sectors - from French's classification, see footnote 14) based on advertisement expenditure (consumer services (17), food (1), other consumer products (7), cars (12) and retail (15)), while columns (7) and (8) represent the remaining companies.

	Panel A: Size			
	Large		Small	
	Ticker	NPS	Ticker	NPS
	(1)	(2)	(3)	(4)
log_dist	-0.046	-0.15	-0.011	-0.016
	(-5.67)	(-5.34)	(-0.47)	(-1.76)
Observations	721,148	674,093	711,724	666,048
R-squared	0.204	0.130	0.211	0.106
Industry FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
AIC	1.30e+06	1.80e+06	3.20e+06	862663
elasticity	-16.5%	-33.1%	-3.0%	-12.5%

	Panel B: Advertisement expenditure			
	Large		Small	
	Ticker	NPS	Ticker	NPS
	(5)	(6)	(7)	(8)
log_dist	-0.036	-0.044	-0.0018	-0.081
	(-4.40)	(-3.69)	(-0.10)	(-4.57)
Observations	1,231,314	1,219,432	1,655,971	1,478,798
R-squared	0.030	0.074	0.179	0.048
Industry FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
AIC	3.30e+06	1.90e+06	7.30e+06	3.00e+06

Figure 2.2: T-stats for the coefficient of log of distance on the monthly regressions with all controls, for ticker and for NPS search

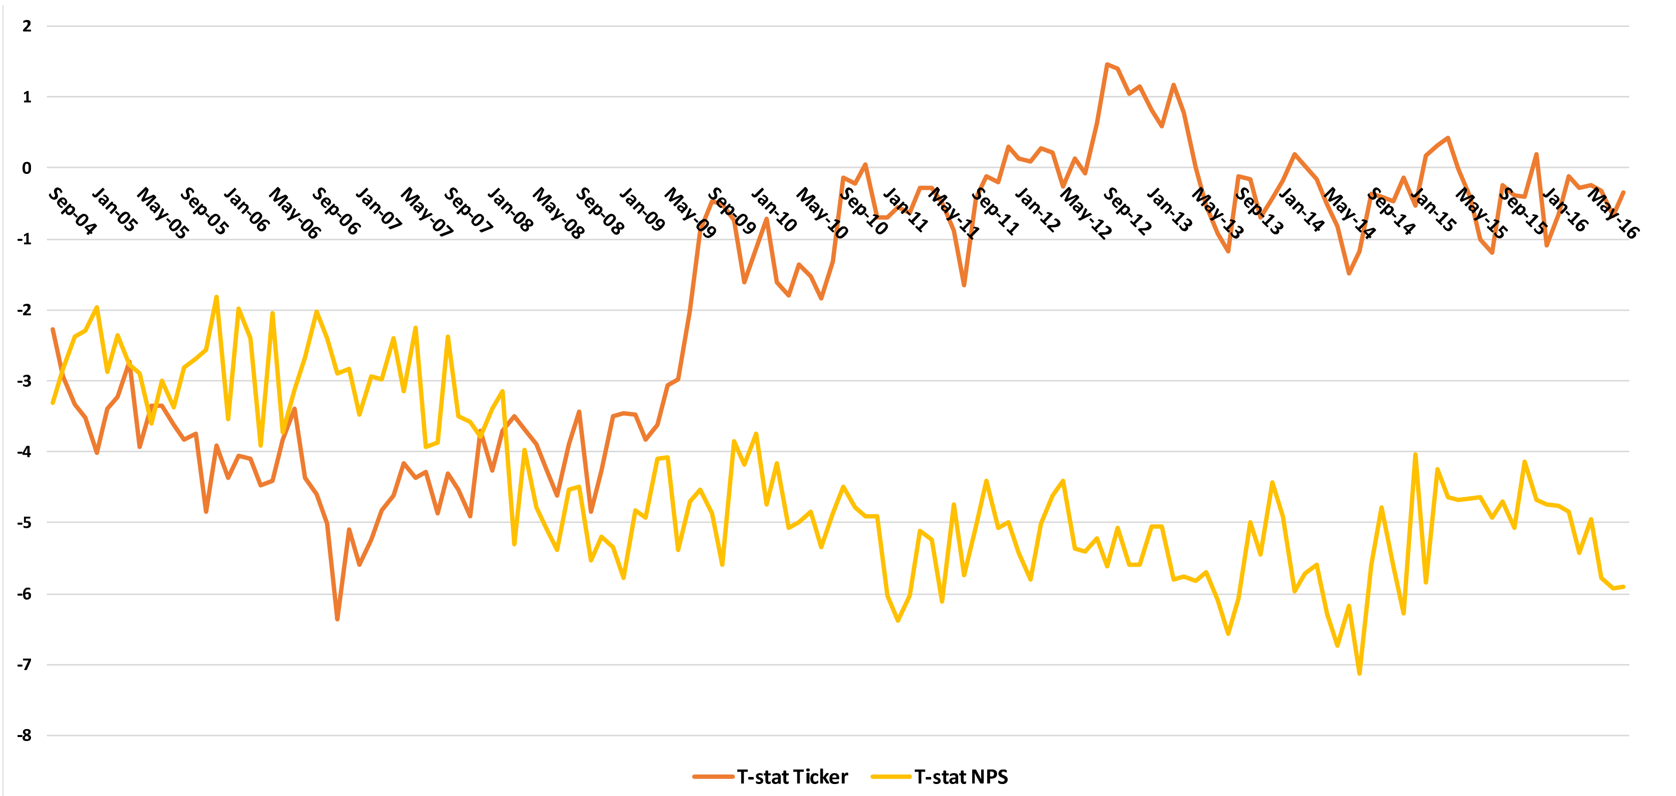

Figure 2.3: Marginal effect of log of distance on the monthly regressions with all controls, for ticker and for NPS search

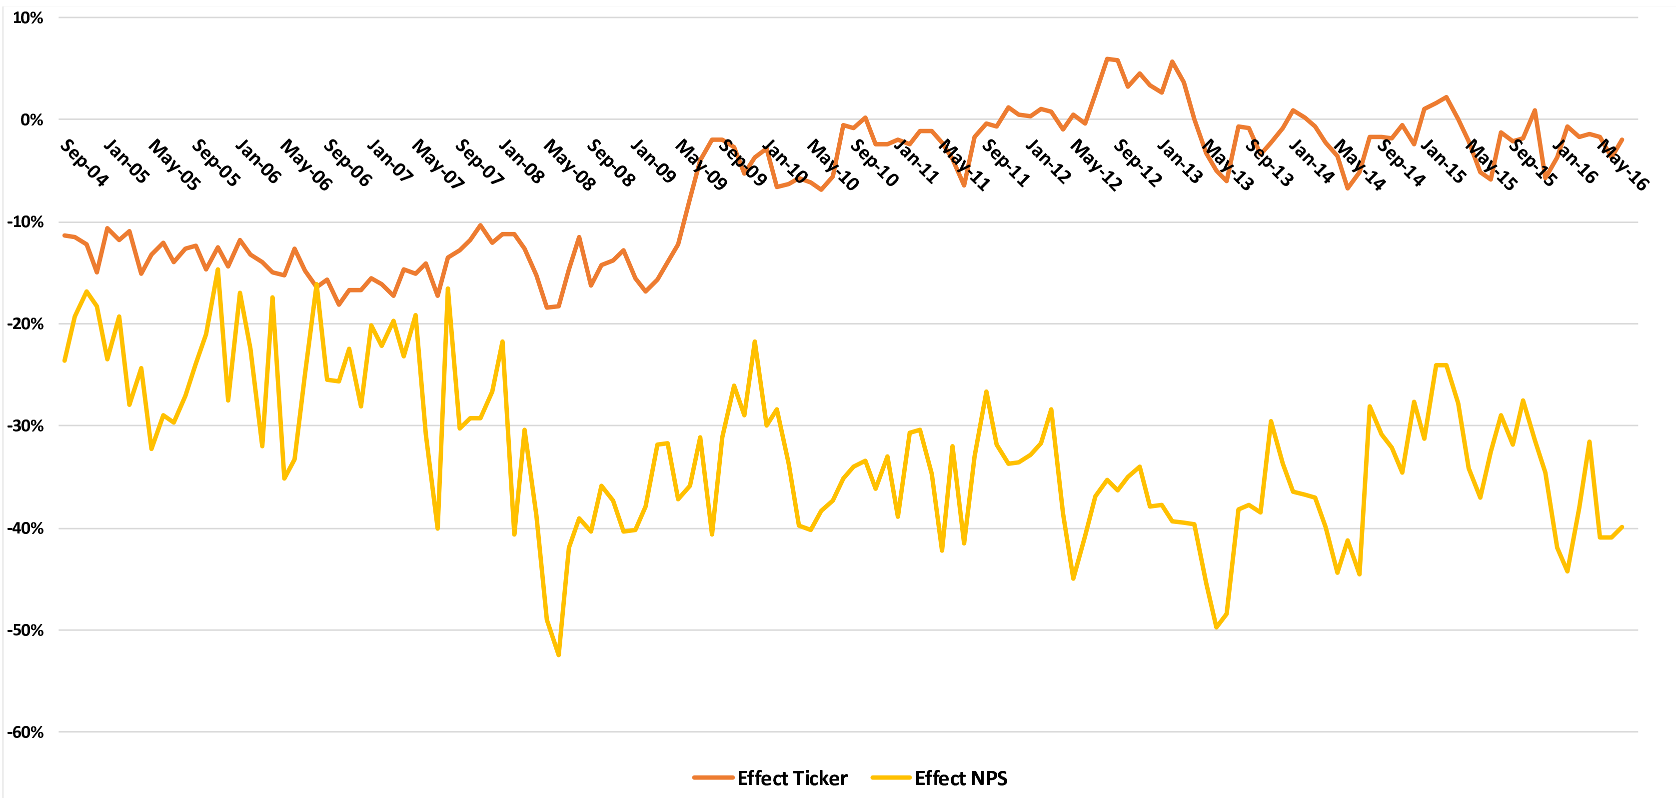

Chapter 3

Advertising Exposure and Investor Attention: Evidence from Super Bowl Commercials*

3.1 Introduction

Publicly listed companies in the US spend roughly 200 billion dollars every year on product advertising. Advertising is a signal not only for the quality of the products that the company provides (Milgrom and Roberts [1986]), but also for its future financial prospects. Under this premise, economists have recently documented a positive relationship between advertising expenditure and measures of stock investment, both at the aggregate (Grullon et al. [2004], Lou [2014]) as well as at the micro level (Branikas [2019]). Between the stages of being exposed to an ad and investing in the company behind it there is the intermediary step of paying more attention to that company's stock. That is the focus of this paper. In particular, we

*Co-authored with Ioannis Branikas, Assistant Professor of Finance at the Lundquist College of Business at the University of Oregon. We thank Stephen Morris, Wei Xiong, Harrison Hong and Jakub Kastl for helpful advice. We thank Gretchen Gamrat for excellent research assistantship.

use the Super Bowl as an experiment to study the effect of high profile commercials on the attention of retail investors.

As of 2019, prices for a 30-second advertisement between Super Bowl quarters start at around 5 million dollars. Even in the beginning of our sample period, in 2011, prices were close to 3 million dollars. On average, around 30 companies every year advertise on the Super Bowl. Figure 3.1 shows in more details the number of advertisers and the costs of advertising for every year in the sample period.

As the match consistently attracts the most viewers of any other program or event on national television, it is the highest prized spot for companies to post ads to attract customers' attention (see, for example, Hartmann and Klapper [2017] and Stephens-Davidowitz et al. [2017][1]). However, Super Bowl commercials also impact the stock market. Fehle et al. [2005] find that Super Bowl ads provide on average a 2% increase in the 20 day post-event cumulative abnormal returns of their companies' stocks. They also show that particularly small stock trade orders experience an increased abnormal net buying activity, pointing to an increased interest from retail investors.

To study the impact of commercials on retail investor attention, we use the novel measure of local investor attention proposed by Buchbinder [2019] to test for the effect of Super Bowl advertisements. This measure is a reformulation of the Google Trend index on searches for the name of each company, followed by the word "stock" — which we call NPS, short for "Name Plus Stock". We can therefore obtain an estimate of the number of searches for one NPS, relative to all other NPSs in the sample - which we interpret as the relative investment interest in that firm. This measure is strictly preferable to using the ticker of a stock, since (i) it conforms with local bias, as shown in Buchbinder [2019], and (ii) it allows us to keep most firms in our sample without removing those with problematic tickers that have double

[1] Like us, Stephens-Davidowitz et al. [2017] also use Google Trends, though for searches for movies. They show that online searches for a movie go up after a movie trailer is shown in the Super Bowl.

meanings (such as "ACT"), or whose tickers may be confused with their names (such as "AOL").

Besides the Google Trends data, we collect three additional sets of information: First, financial characteristics on all the historical members of the S&P500 firms in our sample period. Second, local viewership ratings from Nielsen and demographics in 56 large Designated Market Areas (DMAs). Finally, data on the Super Bowl commercials, including their length and appeal to consumers, captured by the USA Today Ad Meter.

We combine data on viewership and on the Super Bowl commercials to construct our exposure measure of customers to a given firm. Our main exposure measure indicates the fraction of the local population that watches an ad. We also use the commercials' characteristics to experiment with different exposure metrics.

Our first hypothesis is that exposure to advertisement increases investor interest in the stock of the advertising company. We start our analysis with OLS regressions of our investor interest variable on the different advertisement exposure measures, treating them as exogenous and controlling for the distance between the firm's headquarters and the searcher's location, several stock characteristics, and DMA demographics. Our findings suggest that advertisements are highly important in leading investor interest.

Next, we deal with endogeneity concerns that could affect the regional exposure to Super Bowl commercials: (i) viewers may choose to watch the match for the advertisement; (ii) advertisers might be fundamentally different from non-advertisers; or (iii) firms may choose to advertise based on who they expect to be watching the game. We address the first issue by using the participation of a local team in the Super Bowl as an instrument for viewership, similar to Stephens-Davidowitz et al. [2017]. To address the second issue, we select a subsample of firms using a propensity score matching process, based on Roberts and Whited [2013]. We finally explain why the third issue is unlikely to affect our results due to market characteristics of the advertisement slots on the Super Bowl.

The effects of Super Bowl commercials on investor attention that we estimate are both statistically and economically significant. Our main specification finds that a 1 standard deviation increase in the advertisement's exposure measure leads to a increase of over six times in the relative interest for the advertiser's stock.

Our second hypothesis is that firm recognizability in an advertisement is essential in driving this increase in interest. We consider a firm to be recognizable from an advertisement if (i) the advertisement is about the firm, or (ii) the advertisement is about a product whose name overlaps with the company's name. We test this hypothesis by regressing our interest measure on both recognizable advertisements and non recognizable advertisements, and show that the effect of the latter is negligible.

Finally, we conjecture a third hypothesis: that advertisement has a stronger effect on investor interest than being near the firm's headquarters. This hypothesis is in line with Branikas [2019], who shows that the effect of advertising on the household portfolio choice can overcome the effect of local bias.

To test our last hypothesis, we distinguish firms as local versus distant (based on thresholds for the distance between their headquarters and the searcher's DMA), as well as with high versus low advertising exposure. Consistent with our previous results, we find that high regional advertising exposure attracts investors' attention both for local as well as for distant firms. More interestingly, we find that high exposure has a stronger effect on attention than being local to the consumer. In fact, we estimate the advertisement effect to be at least twice as large as the local bias effect.

Our paper contributes to the literature by quantifying the effects of big advertising events, such as the Super Bowl, on investor attention. The estimates that we provide are particularly important due to both the general unavailability of more recent investor portfolio data and to the small number of stocks that active retail investors hold on average (about two or three, according to Barber and Odean [2000], Gargano and Rossi [2016]). Our study also complements the recent findings of Liaukonyte and Zaldokas [2019], who show that EDGAR and Google searches for a company at the state level increase shortly after the airing of a relevant commercial on TV.

This paper is structured as follows. The next section describes the data that is used. Following that, the empirical results are described. Finally, we conclude.

3.2 Hypotheses

Every year, firms invest billions of dollars in advertisements. Whereas the primary goal of such commercials is to foster an increase in product sales (e.g., as in Hartmann and Klapper [2017]), researches have noticed that advertising is positively related to a number of measures of stock investment (Grullon et al. [2004], Branikas [2019]). The question of how exposure to advertising translates into investment behavior, however, remains unaddressed by extant literature. This paper seeks to contribute to current research by helping to bridge this gap.

We begin by noting that while papers such as Grullon et al. [2004], Lou [2014], Branikas [2019] studied a behavior of the investor with regards to a company, most advertisements are not of the companies themselves, but rather of the products whose sales they hope to increase. While such advertisement are usually designed to directly trigger an increase in product consumption, we expect investors will usually procure some more information about a company before making a decision.

Only after researching the firm's stock can the ad viewer exhibit the behaviors which are captured by the measures of stock investment studied by previous scholars. It follows from such reasoning that advertisement exposure ought to be positively related to investor attention in an advertising firm. This leads us to our first hypothesis:

H1: Exposure to advertisement is increases investor interest in the stock of the advertising company.

We may visualize this hypothesis by imagining that the person watching the advertisement during the Super Bowl is not an investor. It is just a person watching a game. But he is also a consumer of products and companies know that. Therefore, they advertise their products, seeking to increase sales. When people watch these ads, however, they are reminded of the existence of the firm behind these products and become interested in these firms. This, in turn, leads them to manifest this increased interest by searching for these firms' stocks online and, eventually, actually investing in them.

This dynamics has, however, an implicit requirement. Namely, that the person watching the ad is able to recognize the firm behind the product. This may be fairly simple for some products (such as Pepsi, offered by PepsiCo) but much more difficult to other products (such as Doritos, also offered by PepsiCo). Arguably, if consumers fail to immediately recognize the firm behind the product, they are less prone to becoming interested in the advertising firm. This leads us to our second hypothesis:

H2: The association of exposure to advertisement and investor interest is contingent on ad recognizability.

By ad recognizability we mean the ease with which consumers are able to identify the firm which is behind it. An ad of a product such as Pepsi would therefore be highly recognizable, whereas an ad of a product such as Doritos would be much less recognizable. In other words, Hypothesis 2 states that an important prerequisite for ad exposure to translate into investor attention is that the advertising firm be readily identifiable from the ad itself. One would therefore expect that an ad for Pepsi would be associated with higher investor interest than an ad for Doritos.

That a firm be easily recognizable from its ad is, therefore, in our framework, an important requirement for ad exposure to translate into investor interest on a company's stock.

Established research has acknowledged another factor which is relevant for investor's choice on which stocks to invest in: local bias. Such research suggests that investors are more eager to invest in firms whose headquarters are close to their own location than to invest in firms whose headquarters are far away. While most research on local bias has considered actual investment behavior, rather than sheer investment interest, Buchbinder [2019] has shown local bias to be important in predicting investor interest itself. In other words, they have found evidence for the existence of local bias in attention. This paper confirms that analysis, and a negative effect from distance on attention is seen throughout the paper, but it is not its focus.

What this paper does consider is whether advertisement exposure is able to overcome local bias in attention, as we see in Branikas [2019]. Can a distant firm become more interesting to investors than a local firm through advertising? We argue that the answer to this question is affirmative and hence put forth our final hypothesis:

H3: The effect of advertisement exposure on investor interest is stronger than local bias.

3.3 Data

3.3.1 Google Trends

The collection of the data on investor interest from Google Trends is based on the NPS - name of the company plus the word "stock" - measure developed in Buchbinder [2019]. Yet, here, we focus only on a single day each year, instead of every month of each year - in particular, the Monday after the Super Bowl. To control for the general level of interest in a stock in a given DMA, we have also downloaded data for the Mondays one month and even one week before the Super Bowl. However, since many Super Bowl advertisements are teased in the days and weeks preceding the game, the interest right before the match might be already contaminated by the ads.

The data collected from Google Trends is rearranged to generate a relative interest, which can be interpreted as the fraction (in percentage points) of searches for a given stock, out of the number of searches for any stock in the sample. So, for example, an interest value of 5 for "Apple stock" in New York means that out of all searches in New York for all the firms in our sample (in the given period), 5% of those searches were related to Apple. To address sampling error in the Google Trends data, we repeat the downloading process 30 times and average our results.

We go over the procedure to obtain this interest variable in detail on Construction of Interest Variable.

3.3.2 Metropolitan Demographics

This paper focuses on 65 major metropolitan regions. While Google Trends seems to use Nielsen's Designated Market Areas (DMAs), the exact demarcations of such regions is proprietary, and therefore precise demographic information is unavailable for research. As a workaround, DMAs were matched with Metropolitan Statistical Areas as defined by the US Census Bureau, which allows us to proxy demographics from government agencies. Unemployment and population numbers were taken from the Bureau of Labor Statistics, while income was drawn from the Bureau of Economic Analysis.

Summary statistics of demographic characteristics are offered in Panel A of Table 3.1.

3.3.3 Stock Characteristics

Since Google Trends rounds down to zero the search results for lesser used terms — as is the case for smaller and relatively unknown stocks — we define the investment universe to be the historical members of the S&P500 between 2007 and 2018.

Monthly stock prices and calculated returns were drawn from CRSP, while firm accounting variables were obtained from Compustat, quarterly. Financial characteristics are calculated at a monthly frequency, and comprise of stock price, market capitalization, book-to-market ratio, turnover ratio (defined as volume over number of shares outstanding), momentum (defined as past annual return), volatility (defined as the standard deviation of monthly returns over the past year), profitability (defined as in Novy-Marx [2013] as the ratio of past annual gross profits to assets), investment (defined as the past annual growth rate of assets) and past annual sales.

The construction of these variables is as in Fama and French [1992]. When we include these characteristics in our regressions as controls, we lag them by one month to ensure that they known to investors. Stocks with missing observations are dropped.

Finally, firms are separated into 17 categories, based on Kenneth R. French's classification to control for visibility to households of different industries[2].

We also calculate the average distance between each DMA in our sample and the companies' headquarters, using their zip code information.

Excluding the companies with missing data leads us to a sample of 571 stocks. Summary statistics of stock characteristics are presented in Panel B of Table 3.1.

3.3.4 Super Bowl - Advertisements, Teams and Viewers

Within the sample period, we gather Super Bowl advertising data from the game itself as well as past editions of USA Today and Advertising Age. For each ad spot, we obtain the length of the segment, a brief description of what was advertised, and

[2] http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/det_17_ind_port.html

the name of the company which purchased the ad. On average, each segment is 30 seconds long, with the vast majority of advertisers appearing multiple times within the sample. Additionally, we collect the relative appeal of each segment using USA Today Ad Meter ratings which compile votes from the public into scores for each ad[3].

We aggregate all advertisement characteristics at the stock level, and focus on companies in our S&P500 investment universe.

We also collect local viewership ratings of the Super Bowl from Nielsen in order to measure the advertisement exposure of the commercials. These are usually reported by local newspapers, on the Monday after the game.

Summary statistics of Super Bowl viewership and advertisement are detailed in Panel C of Table 3.1. 2.8% of the firms in the sample each year have an advertisement on the Super Bowl. The firms in our sample that do advertise buys on average 67 seconds of advertisements, while some firms have up to 4 minutes and 15 seconds of commercials in a single year.

Finally, we also gather the teams that have played on the Super Bowl at each year in our sample, together with their host cities.

Stock Recognizability from Commercials

Whereas we are considering investor interest in companies' stocks, Super Bowl commercials often focus on products rather than on the firms themselves. From watching the ad of a given product to searching for it on Google, the consumer must be able to identify which company is behind the advertised product. In other words, the firm must be recognizable from that ad.

[3] USA Today assembles volunteers to measure their reactions to ads that run during the Super Bowl. The volunteers use handheld meters to register how much they like or dislike each ad. For each ad, they start with a neutral score on their meter and then change the dial as their opinions of the ad change. Each participant's reactions are fed into a computer and averaged on a second-by-second basis. The score given to each ad is based on the point during the commercial when it achieves its highest average score.

Note that a firm may be recognizable from one ad and not from another. To show this, consider, for example, two ads, one for Pepsi and one for Doritos. Though both products are offered by the same company (namely, PepsiCo), consumers are arguably more likely to think of PepsiCo upon watching the ad from Pepsi than upon watching the ad from Doritos. Hence, the property of a firm being recognizable from an ad is a property of the advertisement, not of the firm. We will therefore use the term "recognizable ad" to indicate that the firm behind the ad is recognizable from that ad.

To measure recognizability, we look at (i) whether the ad is about the company or one of its products; and (ii) if it is about a product, whether there is an overlap between the name of the product being promoted and the company's name. Each ad in our sample is individually assessed and a dummy variable created with unit value if the ad was deemed recognizable and zero otherwise. Firms with multiple advertisements on a single year will be considered to have a transparency value of 1 if at least one of their advertisements is deemed recognizable.

From panel C in Table 3.1, we can see that almost 80% of advertisements in our sample are recognizable.

3.4 Estimation

Our empirical specification consists of the following equation:

$$a_{i,j,t} = \alpha + \beta \cdot Exposure_{i,j,t} + Controls_{i,j,t} + \epsilon_{i,j,t} \tag{3.1}$$

where $a_{i,j,t}$ is the attention measure for firm i in region j on date t, and our exposure measure is defined as

$$Exposure_{i,j,t} = View_{j,t} \times Ad_{i,t} \tag{3.2}$$

$View_{j,t}$ is the viewership for region j at year t and $Ad_{i,t}$ is an indicator variable which equals 1 when firm i makes a recognizable advertisement in year t. This measure is essentially the fraction of region j's population that has watched firm i's advertisement during the Super Bowl.

Our controls include: (i) local viewership of the game ($View_{j,t}$) and DMA demographics; (ii) firm $i's$ (nationwide) advertisement in the Super Bowl ($Ad_{i,t}$) and stock characteristics; and (iii) the distance between firm i's headquarters and region j. Following Stephens-Davidowitz et al. [2017], we also control for local interest of a stock before the event and include DMA and year fixed effects. We always use conservative two-way clustered standard errors at the DMA and year levels.

To test Hypothesis 1, we first estimate the above framework treating $Exposure_{i,j,t}$ as exogenous. We then discuss and address endogeneity concerns in Subsection 3.4.2.

3.4.1 OLS Regressions

In Table 3.2, we present results from the OLS regression of Equation 3.1. In Column 1, we control for regional viewership of Super Bowl ($View_{j,t}$), firm $i's$ recognizable advertising ($Ad_{i,t}$), past attention, and year and DMA fixed effects. In Column 2, we add stock characteristics as controls, while in Column (3) we include also DMA demographics.

Across all columns, the results for our exposure variable are similar. In particular, in Column 3 with full controls, the estimated coefficient for exposure is 2.368 and is statistically significant, with a t-statistic of 2.27. The economic effect of a 1 standard deviation increase in our exposure measure is calculated increase interest by over six times, compared to the average interest for a given stock, in a given DMA.

That figure indicates a strong initial support in favor of our first hypothesis. In other words, evidence suggests that there is a strong correlation between viewing a recognizable ad and becoming interested in the stocks of the advertising company.

Advertisement Characteristics

In this subsection we experiment with two alternative definitions for recognizable advertisement exposure, based on the characteristics of the commercials. Specifically, we redefine exposure based on the length or ratings of the commercial. For companies that have multiple advertisements in a single year, we produce these measures by adding their lengths and averaging their ratings.

We present the results of these OLS regressions in Table 3.3.

In Column 1, the estimated coefficient of exposure, defined as $View_{j,t} \times Ad\_Length_{i,t}$, is found to be 0.035, and is marginally statistically significant with a t-statistic of 2.00. The implied economic effect of a 1 standard deviation increase in this measure is to almost double the interest, relative to the average regional interest for a stock. This figure suggests that the longer the exposure to a company's commercial, the higher the attention to that company's stock.

Alternatively, in Column 2, where we define exposure as $View_{j,t} \times Ad\_Rating_{i,t}$, the coefficient is estimated as 0.405, with a t-statistic of 2.02. The economic effect is now to increase by about three quarters, relative to the average regional interest for a stock, suggesting that better commercials bring higher investor attention.

The regressions presented in Table 3.3 provide additional support in favor of Hypothesis 1. They also suggest that the advertising characteristics are important in determining investor attention.

Recognizable and Not Recognizable Ads

We now focus on Hypothesis 2, and investigate whether the advertising effect depends on the recognizability of a company's stock from its commercial.

As stated earlier, we measure the recognizability of an advertisement with an indicator variable that equals one if the name of the promoted product overlaps with the name of the company. This measure is designed to capture whether the advertisement is easily mentally related to the firm behind it. For example, Kraft Foods has produced Super Bowl commercials for some of its brands during our sample period, such as Mio Fit (a brand of drink concentrate) and Planters (nuts), without displaying its name in the video. Such ads are therefore considered to have zero recognizability. However, in 2018, the company shot a commercial to promote its corporate image, and thus its ad is deemed recognizable.

To test Hypothesis 2, we first run an OLS regression of Equation 3.1, where in our exposure measure we replace $Ad_{i,t}$ with a new indicator variable that equals 1 if firm i advertises during the Super Bowl, regardless of whether the advertisement is recognizable. We present the results of this regression in Column 1 of Table 3.4. The estimated coefficient is now 1.705 which is 28% lower than before. The corresponding t-statistic is 1.63, so that the advertisement effect is statistically significant only at the 10% level. This shows that waiving the recognizability requirement decreases the estimated impact of advertising.

Next, to show this more clearly, we keep the original exposure measure intact (i.e., $View_{j,t} \times Ad_{i,t}$, which assumes recognizability) and introduce another exposure measure for the stocks that cannot be recognized from their commercials (i.e., $View_{j,t} \times Ad\_Unrec_{it}$). We present the results of this OLS regression in Column 2 of Table 3.4. The estimated coefficient of our exposure main measure for recognizable ads remains close to its previous level, at 2.370, with a identical t-statistic of 2.27. On the other hand, the coefficient of the exposure of stocks with non recognizable ads is estimated to be 1.133, i.e., roughly 50% smaller than the main one. Its t-statistic is 0.60, making it irrelevant at any reasonable level of statistical significance.

Therefore we are able to accept Hypothesis 2, according to which stock recognizability from the advertisement is necessary in attracting investor interest.

3.4.2 Accounting for Endogeneity

In this section, we deal with endogeneity concerns for our main exposure measure, which could affect our conclusions about Hypothesis 1. As our exposure measure ($View_{j,t} \times Ad_{i,t}$) is the product of two variables, regional Super Bowl viewership and the existence of a firm's recognizable advertisement, we will deal with endogeneity from each of those inputs. The next subsection deals with the simplest of the two, viewership, by using an instrumental variable approach. Afterwards we also use a propensity score matching procedure to obtain a subsample in which we can ignore any endogeneity concerns about what makes a firm decide to make an advertisement.

Regional Viewership

The first concern is that individuals in a region may choose to watch the game (and thus its commercials) because they know an advertisement for a company they are already interested in investing in is being shown. If these people have a sufficient mass, that would explain why cities with high viewership are also cities that search a lot for the advertisers.

We follow Stephens-Davidowitz et al. [2017], and use the participation of a DMA's team in the Super Bowl to instrument regional viewership in that DMA. For every DMA, in each year, we define an indicator variable, $Team_{j,t}$ that equals 1 if DMA j has a local team playing in the match on year t. Our instrument is valid if (i) it is relevant, in the sense that it highly predicts viewership in a given DMA; and (ii) it satisfies the exclusion restriction. Specifically, the second requirement is that our instrument affects local investor's interest in a stock only through the viewership channel. That is, there is no differential effect on the local investor's interest for stocks when their local team is in the Super Bowl, for any reason other than their inclination to watch the game in support of their team.

In Table 3.5 we present the regressions of local viewership on our instrument and regional and year controls: DMA demographics as well as fixed effects for DMA and year. In Column 1, where our instrument is absent, the R^2 is estimated to be 57.6%. In Column 2, where our instrument is included, the R^2 becomes 64.9%, which represents an increase of about 13%. Specifically, the estimated coefficient of the instrument is estimated to be 7.67 with a t-statistic of 8.38. The implied first stage F statistic is 70.1, making our instrument a strong predictor for regional viewership.

Next, in Table 3.6, we present the IV 2SLS and reduced form regressions. Column 1 displays the 2SLS where we instrument exposure, $View_{j,t} \times Ad_{i,t}$, and viewership, $View_{j,t}$, with $Team_{j,t} \times Ad_{i,t}$ and $Team_{j,t}$. The estimated IV coefficient on the exposure is 9.831, i.e., about four times higher than the OLS estimate, and is statistically significant, with a 2.09 t-statistic. In column 2, we show the reduced form of this regression, where we replace $View_{j,t} \times Ad_{i,t}$, and $View_{j,t}$, with $Team_{j,t} \times Ad_{i,t}$ and $Team_{j,t}$, respectively, The estimated coefficient of $Team_{j,t} \times Ad_{i,t}$ is 0.604 and is highly statistically significant with a t-statistic of 2.42.

Our exclusion restriction is violated if participation of a local team in the Super Bowl changes risk preferences in that DMA. For instance, households in such a DMA would be overconfident after their team makes it into the Super Bowl, and begin to consider stock investment. However, under that setup, one would expect that on the Monday after the game, the investment behavior of these individuals would be different, based on whether they have won the game or not.

To examine this possibility we focus on the subsample of DMAs whose team made it to the game and run an OLS regression of investor interest on $Winner_{j,t} \times Ad_{i,t}$, and $Winner_{j,t}$ and full controls, where $Winner_{j,t}$ equals 1 if DMA j's team won the Super Bowl in year t. We report the estimates of this regression in Table table 3.7. The estimated coefficients of both aforementioned variables are very small and statistically insignificant at any reasonable level of statistically significance.

We conclude therefore that winning in the Super Bowl does not shift investor attention in (the high viewership) DMAs that have made it to the game. This suggests that most, if not all, of the effect of local team participation in the Super Bowl should go through viewership.

Super Bowl Advertisement of Firms

A second concern would be that firms that air Super Bowl ads have stocks for which investor interest may already be high. If true, that possibility makes the selection of the Super Bowl advertisers in our investment universe not random. We therefore follow Roberts and Whited [2013] and conduct a propensity score matching procedure to construct a subsample of stocks where at least observable financial characteristics do not predict the probability of airing a Super Bowl commercial.

In Table 3.8, we present linear probability regressions of the indicator variable $Ad_{i,t}$ (which equals one if firm i airs a commercial in year t) on stock characteristics and year fixed effects. The columns are identical in their controls, but differ in the observations that are used in the sampling.

Column 1 displays the estimation results using the whole sample. Indeed, we see that certain financial characteristics are highly predictive of Super Bowl commercials. That is most evident for size, which is expected since Super Bowl commercials are quite expensive.

We therefore focus on companies on the top size quartile and re-run the linear probability regression in Column (2). There, we see that although size is not a predictor of Super Bowl commercials anymore, a few other characteristics such as the book to market ratio continue to predict such advertisements in a statistically significant way.

To address this issue, we focus on a finer subsample of stocks. Specifically, we use the predicted advertising probabilities from the previous regression, to match the Super Bowl advertisers with at most four non advertisers that were sufficiently similar in their estimated advertising probabilities. In Column (3) where we repeat the linear probability regression in this finer subsample, we can see that none of our financial characteristics controls are relevant in explaining why a firm would make an advertisement, with a large drop in the R^2, from 9% and 24% in Columns 1 and 2 respectively to 1% in Column 3.

For the above subsample of stocks, we replicate the IV 2SLS and reduced form regressions of Table 3.6 in Table 3.9. The 2SLS estimate for the our exposure coefficient is found to be 9.831, similar to the estimate in the whole sample (which is 8.835), with a t-statistic of 2.09. In the same spirit the reduced form coefficient estimate is 0.604, i.e., again similar to the coefficient in the whole sample (0.552), with a t-statistic of 2.42.

Since our coefficient estimates for exposure in the subsample are remarkably close to the estimates in the whole sample, this particular endogeneity concern does not seem to be an issue.

Super Bowl Advertisement of Firms in High Viewership Regions

A last possible concern for endogeneity is that companies choose to become Super Bowl advertisers in order to be exposed to certain regions, particularly the ones that are anticipated to have high viewership for the match.

For example: suppose, in the extreme, that The Coca Cola Company, headquartered in Atlanta, chose to air a Super Bowl commercial in 2016 — when the Atlanta Falcons made it to the match — just to promote its image in their local area. However, data from Advertising Age indicates that, usually, more than 90% of the ad slots are sold off at least one month before the game (and thus before the match participants are determined). Additionally, these commercials are typically expensive and take time to produce. Therefore the above endogeneity concern seems unlikely.

3.5 Advertisement Effect: Local and Distant Stocks

Our previous section focused on identifying the effect of advertising on the Super Bowl on investor attention. However, as shown in Buchbinder [2019] as well as in the regressions here, distance ($Log\_Distance_{i,j}$) is also an important factor for investor attention. Therefore, in this section, we investigate how the Super Bowl advertising effect interacts with local bias. In particular, in line with our Hypothesis 3, we attempt here to investigate when advertising is most powerful.

In particular, for every DMA, we distinguish stocks as follows:

1. Local versus distant based on the distance of a stock's headquarters from the DMA i.e., lower versus higher than 100 miles or 250 miles. Both these thresholds have been extensively used in the literature of local bias in the US;

2. With high versus low Super Bowl advertising exposure, based on whether the stock has a Super Bowl ad and the local Super Bowl viewership rating being high in the region. Of course, if a firm does not have any Super Bowl advertisement, its exposure is zero, and therefore considered "low".

Based on the above two criteria, we end up having stocks that are (i) distant with low advertisement exposure, (ii) distant with high advertisement exposure, (iii) local with low advertisement exposure and (iv) local with high advertisement exposure. We pick the first group as the base group and express our regression framework as follows:

$$
\begin{aligned}
a_{i,j,t} =& \alpha + \beta_1 \cdot Away_{i,j,t} \times (Ad_{i,t} \times High\_View_{j,t}) \\
&+ \beta_2 \cdot Local_{i,j,t} \times (Ad_{i,t} \times High\_View_{j,t}) \\
&+ \beta_3 \cdot Local_{i,j,t} \times (1 - Ad_{i,t} \times High\_View_{j,t}) \\
&+ Controls_{i,j,t} + \epsilon_{i,j,t}
\end{aligned} \tag{3.3}
$$

We present the results from our regression in Tables 3.10 and 3.11, each one representing the different definition of locality: 100 and 250 miles, respectively.

In each table, Columns 1, 2 and 3 represent regressions using the whole sample of stocks, while columns 4, 5 and 6 refer to the finer subsample of stocks whose predicted advertising probabilities in Super Bowl are comparable (see subsection 3.4.2). Columns 1 and 4 always show the OLS regression, 2 and 5 the 2SLS (using the participation of a DMA's team in the match as an instrument for regional viewership), and Columns 3 and 6 the reduced form regression (where regional viewership is replaced with the instrument).

In each table, the presented results are robust across all columns. We comment on three findings:

Finding 1. Exposure to Super Bowl has a positive effect on investor attention.

First, independently of whether a firm is local or distant, high regional exposure to a Super Bowl commercial has a positive effect on investor attention. Analytically, the coefficient on $Away_{i,j,t} \times (Ad_{i,t} \times High\_View_{j,t})$ is positive and highly statistically significant. Taking into account that our base group is distant stocks with low regional advertising exposure, this means that distant stocks that advertise can capture more of investor's attention (relative to distant stocks that do not).

In the same spirit, the coefficient on $Local_{i,j,t} \times (Ad_{i,t} \times High\_View_{j,t})$ is estimated to be higher than the coefficient on $Local_{i,j,t} \times (1 - Ad_{i,t} \times High\_View_{j,t})$, indicating that advertising also has a positive effect on local stocks.

Therefore, these results are indeed aligned with our results from the previous section.

Finding 2. Super Bowl ads have a stronger pull on investor attention than local bias.

Second, distant stocks that get high regional advertising exposure through Super Bowl have a stronger effect on investor attention — at least twice as large — than local stocks that do not. In each regression, the estimated coefficient on "away" firms with high exposure (i.e., $Away_{i,j,t} \times (Ad_{i,t} \times High\_View_{j,t})$) is much higher than the coefficient estimate for local firms with low exposure (i.e., $Local_{i,j,t} \times (1 - Ad_{i,t} \times High\_View_{j,t})$).

Therefore we can accept Hypothesis 3, according to which the advertising effect is stronger than local bias effects.

Finding 3. Very distant firms get a Super Bowl advertisement interest "bonus".

Finally, there is one notable distinction in the estimation results of Tables 3.10 versus 3.11, based on whether the distance threshold for a stock's locality is 100 or 250 miles. Specifically, in Table 3.10, where we use a 100 mile radius to define firms as local, local firms with a high regional advertisement exposure are the ones which attract the most investor interest. That is, in every column of that table, the estimated coefficient on $Local_{i,j,t} \times (Ad_{i,t} \times High\_View_{j,t})$ is higher than the estimated coefficient on $Away_{i,j,t} \times (Ad_{i,t} \times High\_View_{j,t})$.

However, in Table 3.11, where we use the 250 mile threshold, we find the reverse: farther firms that advertise attract a higher share of attention than local firms that also advertise. Perhaps these distant firms are not previously familiar to the investors and attract their attention for the very first time after the game, generating a boost in the advertisement effect.

We conjecture that these distant firms may be getting a "novelty bonus", and thus attracting more attention than local firms who also advertise.

3.6 Conclusion

Advertising signals the quality of a product but also of a company more generally. A number of scholars have therefore endeavored to study the relationship between advertising expenditure and measures of stock investment. We argue, however, that there is an intermediary stage between the advertisement and the actual investment which has not yet been addressed in the literature, namely, an increase in investor attention for the firm itself. Our paper contributes to extant research in the field by addressing this gap.

We theorize that advertising exposure translates into increased investor attention. In this view, exposure to product advertising drives investors to become increasingly interested in the stocks of the firms behind these products. This should lead to more individuals researching the fundamentals of these companies.

In order to empirically assess the validity of our hypothesis, we gather data from Super Bowl commercials and Google Trends information on stock searches, as well as control variables, such as financial information. We then attempt to confirm our hypothesis that exposure to advertisement positively impacts investor attention. An OLS regression on our measure of attention on the share of a population that watches an advertisement suggests we are on the correct track.

As a first check on the robustness of our findings, we employ two alternative measures for exposure. These measures take into account that longer or better advertisements might have a stronger effect on an investor's attention. In all such tests, the effect of advertisement exposure on investor attention has both statistical and economic significance. Interestingly, though, we note that the major part of this effect is due to the sheer existence of an exposure. In other words, our findings suggest that the quality of the ad or its length are of positive but lesser importance.

Moreover, we show that the effect of advertising in investor interest is modulated by the consumer's ability to recognize the firm behind the ad. Namely the effect of advertising in investor interest is weakened — statistically insignificant — when the consumer cannot easily associate the company's name to the product being advertised.

We then address potential endogeneity issues that may arise in our main exposure measure. Our results are very robust to using an instrument for viewership and selecting the sample of firms with a propensity score matching procedure which deal with these concerns.

Finally, we show that the effect of advertising in investor interests for a company's stock is sufficiently large do overcome local bias, being at least twice as strong as the latter.

Figure 3.1: Super Bowl Advertisers: Number of Firms and Cost per 30 Seconds

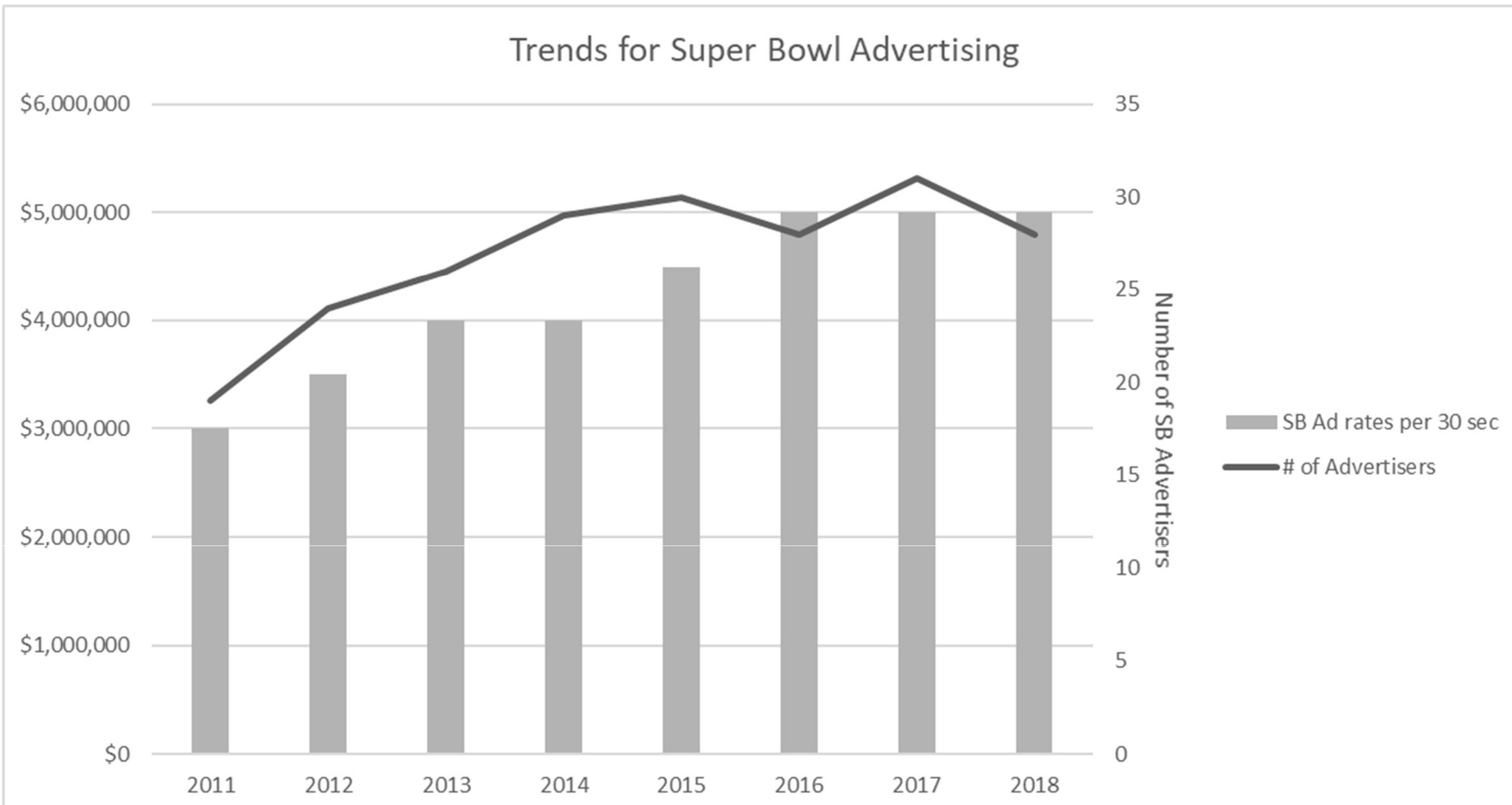

Table 3.1: Summary Statistics

This table presents summary statistics for the variables used in the analysis. Observations with missing data are dropped.

Interest is the share of searches for a particular stock, and thus its average is fixed as the inverse of the (yearly average) number of firms remaining in the sample. Distance is measured approximately, as between the zip code of a firm and the center of a DMA. Monthly stock prices and calculated returns were drawn from CRSP, while firm accounting variables were obtained from Compustat, quarterly. Financial characteristics are calculated at a monthly frequency, and comprise of stock price, market capitalization, book-to-market ratio, turnover ratio (defined as volume over number of shares outstanding), momentum (defined as past annual return), volatility (defined as the standard deviation of monthly returns over the past year), profitability (defined as in Novy-Marx [2013] as the ratio of past annual gross profits to assets), investment (defined as the past annual growth rate of assets) and past annual sales.

Super Bowl and Advertisement characteristics comprise the viewership of the game (estimated by Nielsen), the rating (surveyed by USA Today) and, from various sources and evaluated by the authors: the sum of the lengths of ads by a company in a year, a dummy that measures whether a company has presented a commercial during the match and the recognizability of such ads.

Variable	Mean	Std. Dev.	Min	Max
	Panel A: Demographics			
Income per capita (USD/capita)	45.6	7.5	33.5	79.2
Population (million)	3.0	3.2	0.4	20.0
Unemployment (%)	6.4	2.2	2.5	15.0
	Panel B: Stock Characteristics			
interest (%)	0.26	2.3	0	98
distance (in miles)	1100	700	11	2700
book-to-market	0.5	0.5	-0.9	11.0
turnover	0.2	0.2	0.0	4.8
momentum	0.1	0.4	-0.9	9.9
volatility	0.1	0.04	0.2	0.8
profitability	0.3	0.2	-1.1	1.4
investment	0.1	0.4	-0.7	6.8
size (billion USD)	30	55	0.3	850
	Panel C: Super Bowl and Advertisement			
Local viewership (%)	50	4	38	61
Rating (grade between 0 and 10)	5.3	1.1	2.4	7.8
Length (seconds)	67	49	15	255
Ad (%)	2.8	16	0	100
Recognizability (%)	78	42	0	100

Table 3.2: OLS Regression

This table presents the regression results of investor interest in our main measure of the reach of recognizable advertisements towards Super Bowl spectators: the viewership of a recognizable Ad (*Ad × Viewership*). The dependent variable is *search index for name of company and "stock" (NPS)* on the Monday after each Super Bowl which is normalized to add up to 100 among all companies in the sample. In column (1), we control just for the interest one month before the game, with DMA and year fixed effects. In column (2), we include controls for firm characteristics, while in (3) we also control for local demographics. In parentheses, t-statistics based on two-way clustered standard errors at the level of the year and at the level of the DMA are provided.

	(1)	(2)	(3)
Ad_Recog × Viewership	2.418	2.368	2.368
	(2.262)	(2.268)	(2.268)
Ad_Recog	-0.169	-0.500	-0.500
	(-0.302)	(-0.932)	(-0.932)
Viewership	-0.054	-0.060	-0.060
	(-2.308)	(-2.831)	(-2.820)
Interest_1m_earlier	0.203	0.173	0.173
	(6.595)	(6.632)	(6.632)
Log_Distance		-0.058	-0.058
		(-4.480)	(-4.480)
Log_Size		0.336	0.336
		(13.098)	(13.098)
Book_to_Market		0.079	0.079
		(2.936)	(2.936)
Turnover		0.927	0.927
		(3.135)	(3.135)
Momentum		0.076	0.076
		(0.882)	(0.882)
Volume		0.263	0.263
		(0.520)	(0.520)
Investment		0.023	0.023
		(0.558)	(0.558)
Profitability		-0.007	-0.007
		(-0.114)	(-0.114)
Industry FE	NO	YES	YES
DMA Demographics	NO	NO	YES
DMA FE	YES	YES	YES
Year FE	YES	YES	YES
Observations	129,601	128,629	128,629

Table 3.3: OLS Regression: Length and Ratings

This table presents the regression results of investor interest in alternative measures of the reach of recognizable advertisements towards Super Bowl spectators. The dependent variable is *search index for name of company and "stock" (NPS)* on the Monday after each Super Bowl which is normalized to add up to 100 among all companies in the sample. In column (1), the measure is the product of a recognizable ad's length and the viewership, as we can expect the effect to increase on both the number of eyes watching the commercial and on the time each eye watches it. In column (2), it is the ad rating times viewership - as the effect could increase on both the number of viewers and on the quality of the advertisement. Controls for firm characteristics and demographics are included in both columns. In parentheses, *t*-statistics based on two-way clustered standard errors at the level of the year and at the level of the DMA are provided.

	(1)	(2)
Ad_Length × Viewership	0.035	
	(1.996)	
Ad_Rating × Viewership		0.405
		(2.018)
Ad_Length	-0.013	
	(-1.577)	
Ad_Rating		-0.116
		(-1.149)
Viewership	-0.075	-0.069
	(-2.771)	(-2.409)
Interest_1m_earlier	0.175	0.174
	(6.661)	(6.616)
Log_Distance	-0.057	-0.058
	(-4.462)	(-4.525)
Stock Characteristics	YES	YES
Industry FE	YES	YES
DMA Demographics	YES	YES
DMA FE	YES	YES
Year FE	YES	YES
Observations	128,629	128,629

Table 3.4: Main Specification: Recognizable and Not Recognizable (RnNR) Ads

This table presents the regression results of investor interest in measures of the reach of recognizable advertisements towards Super Bowl spectators. The dependent variable is *search index for name of company and "stock" (NPS)* on the Monday after each Super Bowl which is normalized to add up to 100 among all companies in the sample. In column (1), the main explanatory variable is the interaction between a company having an ad and the metropolitan viewership for it, independent of the ad being recognizable or not. In columns (2), we distinguish between the exposure to recognizable ads and to non recognizable ads. Both columns include controls for having an advertisement, viewership, previous interest, distance, financial and demographics characteristics, and fixed effects for MSA, year and industry. In parentheses, t-statistics based on two-way clustered standard errors at the level of the year and at the level of the DMA are provided.

	(1)	(2)
Ad_RecUnrec × Viewership	1.705	
	(1.628)	
Ad × Viewership		2.370
		(2.266)
Ad_Unrecog × Viewership		1.133
		(0.600)
Ad_RecUnrec	-0.364	
	(-0.694)	
Ad		-0.506
		(-0.940)
Ad_Unrecog		-0.757
		(-0.900)
Viewership	-0.055	-0.067
	(-2.018)	(-2.597)
Interest_1m_earlier	0.174	0.173
	(6.587)	(6.643)
Stock Characteristics	YES	YES
Industry FE	YES	YES
DMA Demographics	YES	YES
DMA FE	YES	YES
Year FE	YES	YES
Observations	128,629	128,629

Table 3.5: Viewership and Local Teams Participation

This table shows the relation between Super Bowl viewership and having a participating team on the match. Year and regional fixed effects are included in both regressions. Dependent variable is Nielsen viewership for the game, with mean 46.4 and standard deviation 5.4. Variable *Team* equals 1 if a local team is participating in the Super Bowl, and 0 otherwise. In parenthesis, *t*-statistics are presented.

	(1)	(2)
Team		7.667
		(8.379)
Constant	38.452	38.415
	(49.601)	(56.966)
DMA Demographics	YES	YES
DMA FE	YES	YES
Year FE	YES	YES
Observations	997	997
R-squared	0.576	0.649

Table 3.6: Two Stages Least Squares — Local Team in SB

This table presents the regression results of investor interest in the instruments for our exposure measures. The dependent variable is *search index for name of company and "stock" (NPS)* on the Monday after each Super Bowl which is normalized to add up to 100 among all companies in the sample. In column (1), we present the instrumental variable regression where we instrument viewership with having the local team in the match (*Team*). In column (2), we run the reduced form, substituting viewership for having a local team in the match. We control for previous interest, local bias (Log_Distance), stock characteristics, demographics and fixed effects for year, industry and region. In parentheses, *t*-statistics based on two-way clustered standard errors at the level of the year and at the level of the DMA are provided.

	(1)	(2)
	2SLS	Reduced form
Ad × Viewership	8.835	
	(2.322)	
Ad × Team		0.552
		(2.261)
Viewership	-0.201	
	(-2.523)	
Team		-0.013
		(-2.399)
Ad	-3.724	0.621
	(-1.890)	(4.610)
Interest_1m_earlier	0.173	0.186
	(6.640)	(6.890)
Log_Distance	-0.058	-0.058
	(-4.472)	(-4.916)
Stock Characteristics	Yes	Yes
Industry FE	Yes	Yes
DMA Demographics	Yes	Yes
DMA FE	Yes	Yes
Year FE	Yes	Yes
First Stage F-Stat	32.203	
Observations	128,629	165,032

Table 3.7: Winning and Searching — Team in Game Subsample

This table presents the regression results of investor interest in the winner viewership of a recognizable Ad ($Ad \times Winner$) . The dependent variable is *search index for name of company and "stock" (NPS)* on the Monday after each Super Bowl which is normalized to add up to 100 among all companies in the sample. This regression uses the subsample of DMAs with a team in the game. Column (1) considers all firms in the sample, while column (2) considers the PSM subsample constructed in Section 3.4.2. Since the results are all insignificant at any reasonable level, we tend to reject the hypothesis that there is a difference between winning and losing on investor interest. In parentheses, t-statistics based on two-way clustered standard errors at the level of the year and at the level of the DMA are provided.

	(1)	(2)
Ad × Winner	0.018	0.26
	(0.07)	(0.51)
Winner	-0.006	-0.250
	(-0.45)	(-0.73)
Ad	1.101	1.016
	(3.03)	(2.39)
Interest_1m_earlier	0.317	0.653
	(2.14)	(5.97)
Log_Distance	-0.073	-0.047
	(-2.03)	(-0.42)
Stock Characteristics	Yes	Yes
Industry FE	Yes	Yes
DMA Demographics	Yes	Yes
DMA FE	Yes	Yes
Year FE	Yes	Yes
Observations	5,165	355

Table 3.8: Predicting Super Bowl Commercials — Subsamples

This table presents the regression results of "having a recognizable advertisement" on our financial controls. The dependent variable is the dummy *Ad_recog*, which equals one if the firm in question has a recognizable ad in that year. Column (1) uses the whole sample of firms in our database. Column (2) excludes small stocks — as size seems to be the most important variable in predicting Super Bowl advertisement. Column (3) uses propensity score matching to include only the stocks that do advertise, plus the four most similar firms to each advertiser — in the propensity score scale — that did not. In parentheses, *t*-statistics based on clustered standard errors at the level of the year are provided.

	(1)	(2)	(3)
	Whole sample	Large firms	PSM subsample
Log_Size	0.033	-0.005	0.001
	(13.766)	(-0.533)	(0.033)
Book_to_Market	0.025	0.089	-0.029
	(2.829)	(2.478)	(-0.395)
Turnover	0.008	0.171	-0.104
	(0.484)	(1.147)	(-0.223)
Momentum	0.014	0.049	0.023
	(0.651)	(1.581)	(0.553)
Volume	0.222	1.029	2.787
	(2.137)	(2.001)	(1.773)
Investment	-0.015	-0.073	-0.278
	(-4.886)	(-6.194)	(-1.766)
Profitability	-0.001	0.059	-0.070
	(-0.048)	(0.554)	(-0.281)
Industry FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
R-squared	0.09	0.24	0.01
Observations	2,947	782	198

Table 3.9: Two Stages Least Squares — Local Team in SB — PSM Subsample

This table presents the regression results of investor interest in the instruments for our exposure measures on a subsample selected through propensity score matching. The dependent variable is *search index for name of company and "stock" (NPS)* on the Monday after each Super Bowl which is normalized to add up to 100 among all companies in the sample. In column (1), we present the instrumental variable regression where we instrument viewership with having the local team in the match. In column (2), we run the reduced form, substituting viewership for having a local team in the match. We control for previous interest, local bias (*Log_Distance*), stock characteristics, demographics and fixed effects for year, industry and region. In parentheses, *t*-statistics based on two-way clustered standard errors at the level of the year and at the level of the DMA are provided.

	(1)	(2)
	2SLS	Reduced form
Ad × Viewership	9.831	
	(2.086)	
Ad × Team		0.604
		(2.415)
Viewership	-1.923	
	(-0.629)	
Team		-0.068
		(-0.629)
Ad	0.722	-4.154
	(3.079)	(-1.675)
Interest_1m_earlier	0.155	0.172
	(6.870)	(6.725)
Log_Distance	-0.117	-0.122
	(-1.738)	(-2.357)
Stock Characteristics	Yes	Yes
Industry FE	Yes	Yes
DMA Demographics	Yes	Yes
DMA FE	Yes	Yes
Year FE	Yes	Yes
First Stage F-Stat	32.777	
Observations	11,088	11,088

Table 3.10: Effect of Super Bowl vs Effect of Distance — 100 miles threshold

This table presents the regression results of investor interest in measures of the reach of recognizable advertisements towards Super Bowl spectators. The dependent variable is *search index for name of company and "stock" (NPS)* on the Monday after each Super Bowl which is normalized to add up to 100 among all companies in the sample. Columns (1-3) use the whole sample, while columns (4-6) use the subsample derived from the PSM procedure from subsection 3.4.2. Columns (1) and (4) represent an OLS regression to compare the effects of having an advertisement and local bias in interest, where local is defined as a firm having its headquarters within 100 miles of the investor. Columns (2) and (5) use having a team in the match as an instrument for viewership, while (3) and (6) represent a reduced form regression substituting viewership for having a team in the game. We control for past interest, local bias, stock characteristics, demographics and fixed effects for year and industry. In parentheses, t-statistics based on two-way clustered standard errors at the level of the year and at the level of the DMA are provided.

	(1)	(2)	(3)	(4)	(5)	(6)
	Whole sample			PSM subsample		
	OLS	2SLS	Reduced form	OLS	2SLS	Reduced form
Away × (Ad × High_View)	0.573	1.167		0.679	1.411	
	(4.481)	(3.130)		(3.139)	(2.679)	
Local × (Ad × High_View)	1.144	1.193		1.151	1.418	
	(2.570)	(2.753)		(1.810)	(2.341)	
Local × (1 - Ad × High_View)	0.397	0.409		0.258	0.447	
	(5.503)	(5.344)		(0.964)	(1.378)	
Away × (Ad × Team)			1.052			0.963
			(3.089)			(2.747)
Local × (Ad × Team)			1.837			1.749
			(2.163)			(2.533)
Local × (1 - Ad × Team)			0.401			0.278
			(5.371)			(0.889)
Past interest	Yes	Yes	Yes	Yes	Yes	Yes
Stock Characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
DMA Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Observations	165,032	165,032	165,032	11,088	11,088	11,088

Table 3.11: Effect of Super Bowl vs Effect of Distance — 250 miles threshold

This table presents the regression results of investor interest in measures of the reach of recognizable advertisements towards Super Bowl spectators. The dependent variable is *search index for name of company and "stock" (NPS)* on the Monday after each Super Bowl which is normalized to add up to 100 among all companies in the sample. Columns (1-3) use the whole sample, while columns (4-6) use the subsample derived from the PSM procedure from subsection 3.4.2. Columns (1) and (4) represent an OLS regression to compare the effects of having an advertisement and local bias in interest, where local is defined as a firm having its headquarters within 250 miles of the investor. Columns (2) and (5) use having a team in the match as an instrument for viewership, while (3) and (6) represent a reduced form regression substituting viewership for having a team in the game. We control for past interest, local bias, stock characteristics, demographics and fixed effects for year and industry. In parentheses, t-statistics based on two-way clustered standard errors at the level of the year and at the level of the DMA are provided.

	(1)	(2)	(3)	(4)	(5)	(6)
	Whole sample			PSM subsample		
	OLS	2SLS	Reduced form	OLS	2SLS	Reduced form
Away × (Ad × High_View)	0.606	1.451		0.727	1.674	
	(4.841)	(2.756)		(3.400)	(2.494)	
Local × (Ad × High_View)	0.428	0.503		0.562	0.903	
	(2.031)	(2.722)		(1.540)	(2.054)	
Local × (1 - Ad × High_View)	0.110	0.125		0.320	0.546	
	(4.758)	(4.553)		(1.766)	(2.110)	
Away × (Ad × Team)			1.308			1.184
			(2.720)			(2.495)
Local × (Ad × Team)			0.673			0.788
			(1.834)			(3.558)
Local × (1 - Ad × Team)			0.106			0.195
			(4.415)			(1.020)
Past interest	Yes	Yes	Yes	Yes	Yes	Yes
Stock Characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
DMA Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Observations	165,032	165,032	165,032	11,088	11,088	11,088

Appendix

3.A Construction of Interest Variable

Since Google Trends limits downloads to contain at most five terms, we first separate the stocks in our sample into groups, each including 4 firms from the sample, in addition to a benchmark firm, which we chose to be Apple. The benchmark will allow us to make comparisons between the groups. We then request from Google Trends the relative search indices for each group (I, j, t), where I is a set of terms (each linked to a firm, i), j is a location and t is the time period. Google in turn picks a random sample of searches satisfying the time and location criteria, and returns a set of Search Volume Indices (SVI) $y_{i,t,j,I} = \frac{x_{i,j,t}}{\sum_{\tilde{i} \in I} x_{i,j,t}}$ for each $i \in I$, where $x_{i,j,t}$ is the number of searches in the sample for firm i, made in location j, during time range t. The SVI is therefore a number between 0 and 100%, representing the estimated proportion of searches made for a single term i, out of a set of terms I.

Letting i_{bench} be our benchmark firm, included in all searches, we calculate a relative attention index, for each firm i where $i \in I$, at each location j and year t:

$$A_{i,j,t,i_{bench}} = \frac{y_{i,t,j,I}}{y_{i_{bench},t,j,I}} = \frac{x_{i,t,j}}{x_{i_{bench},t,j}}$$

which is independent of the group, though dependent on the chosen benchmark (note that $A_{i_{bench},j,t,i_{bench}} = 1$). We therefore renormalize it to obtain:

$$a_{i,t,j} = \frac{a_{i,t,j}}{\sum_i a_{i,t,j}} = \frac{\frac{x_{i,t,j}}{x_{i_{bench},t,j}}}{\sum_i \frac{x_{i,t,j}}{x_{i_{bench},t,j}}} = \frac{x_{i,t,j}}{\sum_i x_{i,t,j}}$$

so that we can interpret the interest variable as a relative interest, $a_{i,t,j}$, the fraction of searches for firm i, relative to all the searches for the firms in our sample. As explained, this process is repeated 30 times, and the average relative interest is kept.

Bibliography

H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, December 1974.

Brad M. Barber and Terrance Odean. Trading is hazardous to your wealth: the common stock investment performance of individual investors. *The Journal of Finance*, 55(2):773–806, April 2000.

Ioannis Branikas. Advertising exposure and portfolio choice: estimates based on sports sponsorships. *Working Paper*, 2019.

Ioannis Branikas, Harrison G. Hong, and Jiangmin Xu. Location choice, portfolio choice. *SSRN Electronic Journal*, 2018.

Gabriel Buchbinder. Local measures of investor attention using google searches. *Princeton PhD Thesis*, 2019.

Jeremy I. Bulow. Durable-goods monopolists. *Journal of Political Economy*, 90(2): 314–332, 1982.

David Butz. Durable-good monopoly and best-price provisions. *American Economic Review*, 80:1062–76, February 1990.

Huaizhi Chen, Lauren Cohen, Umit G. Gurun, Dong Lou, and Christopher J. Malloy. IQ from IP: Simplifying search in portfolio choice. *SSRN Electronic Journal*, 2017.

Tao Chen. Investor attention and global stock returns. *Journal of Behavioral Finance*, 18(3):358–372, June 2017.

Ronald H. Coase. Durability and monopoly. *The Journal of Law and Economics*, 15 (1):143–149, 1972.

John Conlisk, Eitan Gerstner, and Joel Sobel. Cyclic pricing by a durable goods monopolist. *The Quarterly Journal of Economics*, 99(3):489–505, 1984.

Thomas E. Cooper. *Facilitating Practices and Most-favored-customer Pricing*. PhD thesis, Princeton University, 1984. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2016-05-11.

Thomas E. Cooper. Most-favored-customer pricing and tacit collusion. *The RAND Journal of Economics*, 17(3):377, 1986. doi: 10.2307/2555718.

Thomas E. Cooper and Timothy L. Fries. The most-favored-nation pricing policy and negotiated prices. *International Journal of Industrial Organization*, 9(2):209–223, January 1991. doi: 10.1016/s0167-7187(05)80003-7.

Joshua D. Coval and Tobias J. Moskowitz. Home bias at home: Local equity preference in domestic portfolios. *The Journal of Finance*, 54(6):2045–2073, 1999.

Zhi Da, Joseph Engelberg, and Pengjie Gao. In search of attention. *The Journal of Finance*, 66(5):1461–1499, September 2011.

Peter M. Demarzo, Ron Kaniel, and Ilan Kremer. Diversification as a public good: Community effects in portfolio choice. *The Journal of Finance*, 59(4):1677–1716, August 2004.

Michael S. Drake, Darren T. Roulstone, and Jacob R. Thornock. Investor information demand: Evidence from google searches around earnings announcements. *Journal of Accounting Research*, 50(4):1001–1040, April 2012.

Eugene F. Fama and Kenneth R. French. The cross-section of expected stock returns. *The Journal of Finance*, 47(2):427–465, June 1992.

Frank Fehle, Sergey Tsyplakov, and Vladimir Zdorovtsov. Can companies influence investor behaviour through advertising? Super Bowl commercials and stock returns. *European Financial Management*, 11(5):625–647, 2005. ISSN 1468-036X. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1354-7798.2005.00301.x.

Antonio Gargano and Alberto G. Rossi. Does it pay to pay attention? *SSRN Electronic Journal*, 2016.

N. Gennaioli and A. Shleifer. What comes to mind. *The Quarterly Journal of Economics*, 125(4):1399–1433, November 2010.

Gustavo Grullon, George Kanatas, and James P. Weston. Advertising, breadth of ownership, and liquidity. *The Review of Financial Studies*, 17(2):439–461, April 2004. ISSN 0893-9454. URL https://academic.oup.com/rfs/article/17/2/439/1576966.

Faruk Gul, Hugo Sonnenschein, and Robert Wilson. Foundations of dynamic monopoly and the coase conjecture. *Journal of Economic Theory*, 39(1):155–190, 1986.

Wesley R. Hartmann and Daniel Klapper. Super Bowl ads. *Marketing Science*, October 2017. URL https://pubsonline.informs.org/doi/abs/10.1287/mksc.2017.1055.

Harrison Hong, Wenxi Jiang, Na Wang, and Bin Zhao. Trading for status. *The Review of Financial Studies*, 27(11):3171–3212, 2014.

Gur Huberman. Familiarity breeds investment. *Review of Financial Studies*, 14(3): 659–680, July 2001.

Matti Keloharju, Samuli KnÃŒpfer, and Juhani Linnainmaa. Do investors buy what they know? product market choices and investment decisions. *Review of Financial Studies*, 25(10):2921–2958, August 2012.

Guoming Lai, Laurens G. Debo, and Katia Sycara. Buy now and match later: Impact of posterior price matching on profit with strategic consumers. *Manufacturing & Service Operations Management*, 12(1):33–55, January 2010. doi: 10.1287/msom.1080.0248.

Hau L. Lee, V. Padmanabhan, Terry A. Taylor, and Seungjin Whang. Price protection in the personal computer industry. *Management science*, 46(4):467–482, 2000.

Jura Liaukonyte and Alminas Zaldokas. Background noise? tv advertising affects real time investor behavior. *Working Paper*, (ID 3336176), February 2019. URL https://papers.ssrn.com/abstract=3336176.

Dong Lou. Attracting investor attention through advertising. *The Review of Financial Studies*, 27(6):1797–1829, June 2014. ISSN 0893-9454. URL https://academic.oup.com/rfs/article/27/6/1797/1596985.

Robert C. Merton. A simple model of capital market equilibrium with incomplete information. *The Journal of Finance*, 42(3):483–510, July 1987.

Paul Milgrom and John Roberts. Price and advertising signals of product quality. *Journal of Political Economy*, 94(4):796–821, 1986. ISSN 0022-3808. URL https://www.jstor.org/stable/1833203.

Robert Novy-Marx. The other side of value: The gross profitability premium. *Journal of Financial Economics*, 108(1):1–28, April 2013.

Lin Peng and Wei Xiong. Investor attention, overconfidence and category learning. *Journal of Financial Economics*, 80(3):563–602, June 2006.

Tomas Reyes. Limited attention and m&a announcements. *Journal of Empirical Finance*, 2018.

Michael R. Roberts and Toni M. Whited. Endogeneity in empirical corporate finance. In *Handbook of the Economics of Finance*, pages 493–572. Elsevier, 2013.

Mark S. Seasholes and Ning Zhu. Individual investors and local bias. *The Journal of Finance*, 65(5):1987–2010, 2010.

Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, April 2003.

Joel Sobel. Durable goods monopoly with entry of new consumers. *Econometrica: Journal of the Econometric Society*, pages 1455–1485, 1991.

Seth Stephens-Davidowitz and Hal Varian. A hands-on guide to google data. *Tech. Rep.*, 2014.

Seth Stephens-Davidowitz, Hal Varian, and Michael D. Smith. Super returns to Super Bowl ads? *Quantitative Marketing and Economics (QME)*, 15(1): 1–28, 2017. URL https://ideas.repec.org/a/kap/qmktec/v15y2017i1d10.1007_s11129-016-9179-0.html.

Nancy L. Stokey. Rational expectations and durable goods pricing. *The Bell Journal of Economics*, pages 112–128, 1981.

Ning Zhu. The local bias of individual investors. *Yale ICF Working Paper No. 02-30*, 2002.