Consumer Scores and Market Segmentation*

Alessandro Bonatti and Gonzalo Cisternas

June 1, 2018

Abstract

A long-lived consumer interacts with a sequence of firms in a stationary Gaussian setting. Firms rely on the consumer's score—an aggregate measure of past quantity signals with exponential discounting—to learn about her preferences and to set prices. In the unique stationary linear Markov equilibrium, strategic demand reduction by the consumer drives expected prices below their static benchmark. The precision of the firms' information is maximized by scores that induce more persistent beliefs than under public signals of past quantities. Hidden scores—observed by firms but not by consumers—reduce demand sensitivity, increase expected prices, and reduce expected quantities.

Keywords: Market Segmentation; Ratings; Signaling; Ratchet Effect; Persistence; Transparency.

JEL codes: C73, D82, D83.

*Bonatti: MIT Sloan School of Management, 100 Main St., Cambridge, MA 02142, bonatti@mit.edu. Cisternas: MIT Sloan School of Management, 100 Main St., Cambridge, MA 02142, gcistern@mit.edu. Bonatti gratefully acknowledges support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 670494), as well as the hospitality of the Toulouse School of Economics. We thank Rossella Argenziano, Renato Gomes, Johannes Hörner, Pino Lopomo, George Mailath, Juuso Toikka, Alexander White, and seminar participants at Bilkent, Boston University, Chicago, Duke, Harvard, Indiana, INFORMS, MIT, Miami, Napoli, Penn, Princeton, Texas, and Toulouse for helpful discussions. Tina Pavlovich provided excellent research assistance.

1 Introduction

The ability to tailor advertising, content, products, and prices to the consumers' preferences is a critical driver of firm profitability in online markets. To facilitate market segmentation, several data brokers (e.g., Acxiom, Equifax, Experian) collect consumer-level behavioral and demographic information from online and offline sources. Because of their potential adverse and discriminatory use, the collection and distribution of personal data have received significant attention by policymakers.[1] To correctly assess the impact of regulation, however, one must take into account how technological and market forces affect consumers' incentives. In particular, if the final use of information impacts the distribution of surplus, the mechanisms by which consumer data is aggregated and transmitted determine equilibrium behavior in the transactions where the data is collected. In turn, this may affect the equilibrium prices of those transactions and the informational content of the data itself.

A prominent way to operationalize individual-level data is to create *consumer scores*—metrics that "describe an individual and predict a consumer's behavior, habit, or predilection" (Dixon and Gellman, 2014). For example, a consumer score may combine information about individual customers' age, ethnicity, gender, household income, zip code, and purchase histories to quantify their profitability, health risk, job security, or credit worthiness.[2] From the consumer's perspective, the possible uses of a score range from beneficial (e.g., personalized content and advertising) to ambiguous (e.g., tailored pricing of goods and services) to harmful (racial discrimination or identity theft). At a broad level, any such score creates a link between a consumer's interaction with one firm and the terms of her future transactions with other firms and industries.[3]

In this paper, we explore the informational and welfare consequences of summarizing consumers' purchase histories into scores that are used for price discrimination. We focus on policy-relevant questions related to the composition and the transparency of these scores. For example, how heavily should the score weigh information about the consumer's most recent actions, and how quickly should it discount past information? Should consumers be aware of the score that data brokers have currently assigned to them? Our approach consists

[1]The European Union's General Data Protection Regulation focuses on improving consumers' control over the retention and the distribution of their personal data. In the United States, the policy debate has mainly focused on the transparency of data brokers' information (Federal Trade Commission, 2014).

[2]Specific examples of consumer scores include: the Acxiom Consumer Prominence Indicator Score that measures the size of a consumer's economic footprint; the Equifax Discretionary Spending Limit Index; the Experian Consumer View Profitability Score; and the AppNexus/ Jivox IQ Dynamic Audience Scoring tool. All these scores fall outside of the narrow definition of credit scores, and hence, they are not subject to the same regulation under the Fair Credit Reporting Act (Schmitz, 2014).

[3]For example, information about a consumers sporting goods purchases or eating habits can become a part of a predictive score for a health insurer. See "Very personal finance," *The Economist*, July 2nd 2012.

of developing a dynamic model where the consumer can manipulate her own score, and hence the prices she will face in the future, through the current quantity demanded.

Our key modeling tool is a continuous-time version of the ratchet-effect framework that is sufficiently tractable to analyze the role of information management, persistence, and transparency. In particular, we consider a model with linear signals, quadratic preferences, and Gaussian fundamentals. In particular, the consumer's private willingness to pay follows a mean-reverting diffusion process with known parameters. The consumer's actions generate noisy signals, i.e., the quantity purchased is recorded with error.[4] An (unmodeled) intermediary aggregates each consumer's past quantity signals into a one-dimensional score that discounts past signals at an exponential rate. The score is revealed to a sequence of short-lived firms that use it to set prices.

There are several reasons to focus on (third-degree) price discrimination based on purchase histories: (i) it is implicitly used in a number of markets, whether in the form of coupons, discounts, fees (e.g., Uber), return policies, order of search results, and even reserve prices in advertising auctions;[5] (ii) thanks to better data, market semgentation is becoming increasingly profitable and its use is likely to expand in the future (Dubé and Misra, 2017); and (iii) as a well-understood workhorse model, it allows us to highlight the key economic forces at play with other discriminatory uses of individual information. Because scores are used for pricing, we explicitly ignore any "horizontal" aspect of information transmission. Such aspects would be directly beneficial to the consumer, by facilitating the matching of content to her tastes. Finally, in our model, third-degree price discrimination is harmful to consumers in a one-period interaction. This allows us to focus on the equilibrium effect of information revelation in a dynamic model, which can benefit the consumer by affecting the dynamics of prices over time.

Overview of the Results For any exponential score (i.e., for any discount rate of past signals), there exists a unique equilibrium. Strategic consumers react to the possibility of firms ratcheting up prices by reducing their quantity demanded relative to the static optimal choices. Thus, the *ratchet effect* drives average equilibrium prices down for all levels of the score's persistence, but it allows firms to better tailor prices to the consumer's type. Not

[4]Consumer scores are far from perfect, and two types of attribution are more frequent than with credit scores. For example, a score may combine data from two different users with the same device ID, or fail to match a single users activity across multiple devices (Rafieian and Yoganarasimhan, 2018).

[5]Google's bidder-specific dynamic price floor is described in Ad Exchange auction model at https://support.google.com/adxbuyer/answer/6077702. For a discussion of the scope of price discrimination see Council of Economic Advisers (2015) and Kehoe, Larsen, and Pastorino (2018). For more details on personalized reserve prices in online auctions, see Ostrovsky and Schwarz (2016), Syrgkanis, Kempe, and Tardos (2017), and the references therein.

all consumers, however, have equal incentives to manipulate future prices. Higher consumer types, who have are likely to buy large quantities in the immediate future, have a stronger incentive to reduce the quantity they purchase to manipulate prices downward. This yields a lower sensitivity of the consumer's actions to the underlying type, which in turn reduces the amount of information revealed to the market in equilibrium through the consumer score.

A unique persistence level yields the same equilibrium outcome as revealing the full history of (noisy) signals, i.e., there exists a unique way to aggregates the data without concealing payoff-relevant information. The non-concealing score does not, however, maximize the amount of information conveyed to the firms in equilibrium. A more persistent score—one that places excessive weight on the past—narrows the gap between different consumer types' incentives to manipulate prices. Thus, a greater persistence level increases the responsiveness of the consumer's action to her type. The effect on the consumer's equilibrium behavior offsets the (second-order) loss in information from aggregating the signals in a suboptimal way and increases the amount of information conveyed by the score.

We then turn to the welfare implications of the score's persistence. We show that the firm's (ex ante) expected profits can be written in terms of the mean and variance of the equilibrium price. For very valuable market segments, firms would prefer to operate under no information, so to eliminate the ratchet effect and encourage the consumer to buy. Conversely, for less valuable segments, firms benefit from informative scores. In many cases, the firm-optimal score is more persistent than the maximally informative score. By further increasing the persistence of the score, firm are able to trade-off the precision of the information they receive in equilibrium for greater sensitivity of the consumer's demand to her true type, which enhances their ability to price discriminate.

In contrast, the consumer-optimal score is uninformative if the mean of the consumer's type is sufficiently low. If the mean of the consumer's type is sufficiently high, consumers prefer an informative score to not being tracked, because the presence of a score leads firms to reduce prices. Consumer surplus, however, displays an additional effect, whereby a quantity demanded that is responsive to the true type (i.e., buys more when the type is high) is per se beneficial. This third effect can lead the consumer-optimal score to be more or less persistent than the unique "public" score that yields the same equilibrium outcome as observing the full history of signals. Overall, consumer surplus can be maximized by a more or less persistent score, depending on the discount rate and the persistence of the underlying type.

Finally, when the level of the score is hidden from the consumer, the firms can signal their private information. Becaue prices are observed without noise, the contemporaneous price perfectly reveals the current score in a pure-strategy equilibrium. As such, the current price carries a signal of future prices that affects the consumer's incentives to manipulate the score.

In particular, a high current price signals that prices will be high in the future. Therefore, the consumer expects to buy few units in the near term, and the value of manipulating her score by reducing the current quantity demanded is diminished. Conversely, a low price signals a good opportunity to manipulate scores downward, which limits the consumer's response to a price cut. Therefore, with hidden scores, the consumer's demand is less price sensitive, average prices are higher, and average quantities are lower, relative to the case of public scores. Numerical results show the ranking of observable and hidden scores extends to welfare levels.

From a transparency standpoint, consumer scores stand today where credit scores stood in the 1950s: in the shadows (Dixon and Gellman, 2014). Even when consumers know they are being tracked, they have little recourse to learn, verify or correct the information in their scores. We therefore examine the case of scores that are hidden to the consumer.[6] Overall, our analysis suggests that transparent scores are beneficial to consumers, even if they are already aware of their existence.

Related Literature This paper builds on the literature on behavior-based price discrimination (Taylor, 2004; Acquisti and Varian, 2005), whose results are described in great detail in the surveys by Fudenberg and Villas-Boas (2006, 2015) and by Acquisti, Taylor, and Wagman (2016). Closest to our work is the two-period model with noiseless signals in Taylor (2004) that shows how demand reduction results in lower equilibrium prices when consumers are strategic. Thus, firms may want to commit to observing no information about their consumers. In our stationary model, the presence of noisy signals implies that the informational effects of the consumer's actions are endogenous. This allows us to study how the score's persistence and transparency affect the information available to firms in equilibrium.[7]

The economic force driving the dynamics of our model is the ratchet effect, which was initially studies in Freixas, Guesnerie, and Tirole (1985), Laffont and Tirole (1988), and Hart and Tirole (1988). Gerardi and Maestri (2016) develop an infinite-horizon version with two types, perfectly observable actions, and characterize a mixed-strategy equilibrium. The ratchet effect also underscores the analysis of privacy in a model with multiple principals. See, e.g., Calzolari and Pavan (2006) for the case of two principals and Dworczak (2017) for

[6]Data brokers make few attempts at improving transparency. Two exceptions are the Oracle/Bluekai Registry http://www.bluekai.com/registry/ that reveals to consumers which interest groups they belong to, and Acxioms About the Data initiative.

[7]Cummings, Ligett, Pai, and Roth (2016) and Shen and Miguel Villas-Boas (2017) study two-period models in which advertisement messages are targeted on the basis of the information about consumers' purchases. These papers highlight trade-offs similar to ours, where the value of targeted advertising impacts the equilibrium price of the first-period good and the amount of information revealed by the consumer. In contrast, McAdams (2011) assumes that the buyer chooses whether the seller can observe her purchase history, and hence, the buyer's decision to disclose information can be used to facilitate price discrimination.

the case of a single transaction followed by an aftermarket. In our paper, the noise in the signals and restriction to linear pricing limit the ratcheting forces and create conditions for firms to benefit from the transmission of information.[8]

Finally, our paper contributes to the literature on dynamic information design. Specifically, one can view our consumer as a privately informed sender whose preferences depend on the mean and variance of the receivers' beliefs. Our construction of consumer scores leverages the notion of a linear Gaussian *rating* pioneered in Hörner and Lambert (2017). Relative their paper, we maintain the assumptions of short-lived firms and additive signals. In other words, the agent does not control the precision of the information directly. We consider only linear scores with exponential weights, but we allow for a privately informed agent. This signaling component determines the equilibrium precision of signals. Moreover, unlike in the symmetric uncertainty model of Hörner and Lambert (2017), the consumer's nonlinear preferences yield optimal actions (quantities) that depend on the level of the firms' beliefs. As a result, the transparency question is important in our setting—prices and signaling incentives depend critically on whether the agent knows her own score.[9]

2 Model

We develop a continuous-time game analog of a repeated interaction between a long-run consumer and an infinite sequence of short-run firms characterized by two central features. First, the consumer faces a different monopolist in every period. Second, within any period, the consumer and the operating monopolist play a sequential-move stage game in which the monopolist moves first by posting a unit price for the good it produces; subsequently, having observed the price, the consumer chooses a quantity to purchase.

Players, types, and payoffs. Directly in continuous time, consider a long-lived consumer who interacts with a continuum of firms over an infinite horizon. The consumer discounts the future at rate $r > 0$ and, at any instant of time $t \geq 0$, consuming $Q_t = q$ units of the good at price $P_t = p$ delivers a flow utility

$$u(\theta, p, q) := (\theta - p)q - \frac{q^2}{2}, \tag{1}$$

[8] The ratchet effect appears, with a different interpretation or motivation, in relational contracts with and without private information (Halac, 2012; Fong and Li, 2016) and in dynamic games with symmetric uncertainty (Cisternas, 2017b).

[9] In related contributions, Heinsalu (2017) analyzes a dynamic game with noisy signaling and Gaussian information, and Di Pei (2016) analyzes coarse performance scores in a model with privately informed agents and Poisson learning. See also Kovbasyuk and Spagnolo (2016) for a related study of optimal memory and information design in markets.

where $\theta_t = \theta$ is the consumer's *type* at t, a measure of her willingness to pay at that instant. We assume that throughout the analysis that the type process is a stationary mean-reverting process with mean $\mu > 0$, speed of reversion $\kappa > 0$, and volatility $\sigma_\theta > 0$, i.e.,

$$d\theta_t = -\kappa(\theta_t - \mu)dt + \sigma_\theta dZ_t^\theta, \ t > 0, \tag{2}$$

where $(Z_t^\theta)_{t\geq 0}$ is a Brownian motion.[10] In particular, $(\theta_t)_{t\geq 0}$ is Gaussian and, by stationarity,

$$\mathbb{E}[\theta_t] = \mu \quad \text{and} \quad \text{Cov}[\theta_t, \theta_s] = \frac{\sigma_\theta^2}{2\kappa} e^{-\kappa|t-s|}, \quad \text{for all} \ \ t, s \geq 0. \tag{3}$$

Each firm interacts with the consumer for only one instant, and only one firm operates at any time t; we refer to such firm as *monopolist* t. Monopolist t's ex post profits are given by $P_t Q_t$ if the posted price is P_t and the quantity demanded is Q_t, $t \geq 0$.

In the population interpretation of the model, each consumer is identified with a different path of $(\theta_t)_{t\geq 0}$. Consumers with high willingness to pay are those who exhibit large upward deviations from μ, the average willingness to pay of the relevant market segment; as (2) suggest, however, most consumers fluctuate around μ. Finally, with different goods being offered, (3) states that the consumer encounters goods for which the associated measures of willingness to pay are positively correlated, but that such dependence weakens as the time between purchases increases (e.g., searching for specific good leads to offers of similar products, but subsequent offers becomes less related to the original one as shopping progresses).

Score process and information. At any $t \geq 0$, monopolist t observes only the current value Y_t of a *score process* $(Y_t)_{t\geq 0}$ that is provided by an (unmodeled) intermediary. In contrast, the consumer observes the entire history of scores $Y^t := (Y_s : 0 \leq s \leq t)$, in addition to past prices and quantities, and types realizations.[11]

Building a score process is a two-step procedure that involves data collection followed by data aggregation. We assume that the intermediary collects information about the consumer using a technology that records purchases with noise. Specifically, the intermediary observes

$$d\xi_t = Q_t dt + \sigma_\xi dZ_t^\xi, \ t > 0,$$

where $(Z^\xi)_{t\geq 0}$ is a Brownian motion independent of $(Z^\theta)_{t\geq 0}$, and where Q_t is the (realized) quantity demanded by the consumer at $t \geq 0$.

[10] Stationarity is equivalent to $\theta_0 \sim \mathcal{N}(\mu, \sigma_\theta^2/2\kappa)$, independent of $(Z_t^\theta)_{t\geq 0}$.

[11] In Section 5, we let the time-t monopolist to observe the entire history of scores Y^t. On the other hand, in Section 7 firm t observes Y_t only, but the score is hidden to the consumer.

The intermediary then operationalizes the data by aggregating every history of the form $\xi^t := (\xi_s : 0 \leq s < t)$ into a real number Y_t that corresponds to the consumer's time-t score, $t \geq 0$. Building on Hörner and Lambert (2017), we restrict attention to the family of *exponential scores*, i.e., to Ito processes of the form

$$Y_t = Y_0 e^{-\phi t} + \int_0^t e^{-\phi(t-s)} d\xi_s,\ t \geq 0, \tag{4}$$

where $\phi \in (0, \infty)$. Under this specification, the consumer's current score is a linear function of the contemporaneous history recorded purchases, and lower values of ϕ lead to scores processes that exhibit more *persistence*—they discount old information less quickly. In differential form, the score process satisfies

$$dY_t = -\phi Y_t dt + d\xi_t,\ t > 0.$$

Finally, the prior is that (θ_0, Y_0) is normally distributed; the exact distribution is determined in equilibrium so that to the joint process $(\theta_t, Y_t)_{t\geq 0}$ that is stationary Gaussian along the path of play.[12] In what follows, the expectation operator $\mathbb{E}[\cdot]$ is with respect to such prior, while $\mathbb{E}_0[\cdot]$ conditions on the realized value of (θ_0, Y_0)—the former operator is the relevant for welfare analysis, while the latter for equilibrium analysis. Conditional expectations for the consumer and monopolist t are therefore denoted by $\mathbb{E}_t[\cdot]$ and $\mathbb{E}[\cdot|Y_t]$, respectively.

Strategies and Equilibrium Concept. A strategy for the consumer specifies, for every $t \geq 0$, a quantity $Q_t \in \mathbb{R}$ to purchase as a function of the history of past prices, types, and score values, $(\theta_s, P_s, Y_s : 0 \leq s \leq t)$. Instead, monopolist t must choose a price $P_t \in \mathbb{R}$ that is measurable with respect to Y_t only, $t \geq 0$. We say that a strategy for the consumer is *linear Markov* if, at any $t \geq 0$, $Q_t = Q(p, \theta_t, Y_t)$ where $Q : \mathbb{R}^3 \to \mathbb{R}$ is linear and p denotes the current posted price (i.e., the consumer conditions her quantity demanded on the observation of the contemporaneous price). Similarly for firm t, $P_t = P(Y_t)$ with $P : \mathbb{R} \to \mathbb{R}$, $t \geq 0$.

We focus on Nash equilibria in linear Markov strategies with the property that $(\theta_t, Y_t)_{t\geq 0}$ is stationary Gaussian. From this perspective, given $P(\cdot)$ linear, an admissible strategy for the consumer is any process $(Q_t)_{t\geq 0}$ satisfying (i) progressive measurability with respect to the filtration generated by $(\theta_t, Y_t)_{t\geq 0}$, (ii) $\mathbb{E}_0\left[\int_0^T Q_s^2 ds\right] < \infty$, for all $T > 0$, and (iii) $\mathbb{E}_0\left[\int_0^\infty e^{-rt}(|\theta_t Q_t - Q_t^2/2| + |P_t(Y_t)Q_t|)dt\right] < \infty$: requirement (i) states that, at histories where firms have chosen prices as prescribed by any candidate equilibrium, the history $(\theta_s, Y_s : 0 \leq s \leq t)$ captures all the information that is relevant for future decision-making;

[12] See Proposition 1 in Section 4.1.

(ii) and (iii) are purely technical.[13]

Definition 1. *A pair (Q, P) of linear Markov strategies is a Nash equilibrium if*

(i) The process $(Q(P(Y_t^{Q,P}), \theta_t, Y_t^{Q,P}))_{t\geq 0}$, where $(Y_t^{Q,P})_{t\geq 0}$ denotes the solution to $dY_t = (-\phi Y_t + Q(P(Y_t), \theta_t, Y_t))dt + \sigma_\xi dZ_t^\xi$, $t \geq 0$, maximizes

$$\mathbb{E}_0 \left[\int_0^\infty e^{-rt} u(\theta_t, P(Y_t), Q_t) dt \right]$$

subject to $dY_t = (-\phi Y_t + Q_t)dt + \sigma_\xi dZ_t^\xi$, among all admissible strategies $(Q_t)_{t\geq 0}$;

(ii) At any time-t history such that $Y_t = y$, $p = P(y)$ maximizes $p\mathbb{E}[Q(p, \theta_t, y)|Y_t = y]$.

We say that a linear pair (P, Q) is a stationary *linear Markov equilibrium if, in addition, the joint process $(\theta_t, Y_t^{Q,P})_{t\geq 0}$ is stationary Gaussian.*

In a Linear Markov (Nash) equilibrium, the optimality of the consumer's strategy is verified only when firms set prices according to $P_t = P(Y_t)$ for all $t \geq 0$, i.e., on the path of play—we elaborate more on this in Section 4.2 where we refine our solution concept to provide an analog of Markov Perfect equilibrium in our continuous time setting. The stationarity notion in turn encompasses two ideas: the score must admit a proper long-run distribution, and such distribution prevails from time zero. These properties allow us to perform meaningful welfare analysis that, in addition, is time invariant.

3 Price Discrimination Benchmark

In order to address our policy questions around persistence and transparency of the consumer's score, it is useful to begin with a review the basic effects of (static) price discrimination and (dynamic) behavior-based price discrimination in our environment.

Consider first a consumer with preferences as in (1) who interacts with a firm only once. The firm's prior about the consumer's type has mean $\mu \in \mathbb{R}$ and variance $\text{Var}[\theta] > 0$. Before interacting with the consumer, a public signal Y about the consumer's type is realized. This signal can be general, i.e., it is not necessarily Gaussian.

[13]Requirement (iii) is a mild strengthening of the usual condition $\mathbb{E}_0 \left[\int_0^\infty e^{-rt} |u(\theta_t, P(Y_t), Q_t)| dt \right] < \infty$ usually imposed for verification theorems to hold (Sections 3.2 and 3.5 in Pham (2009)). In particular, strategies with the unappealing property of yielding high payoffs by making expenditures arbitrarily negative are ruled out, provided they exist. Finally, under (ii), the controlled score dynamic admits a strong solution given any initial conditions, so the best-response problem is well-defined (Section 3.2 in Pham (2009)).

Let $M := \mathbb{E}[\theta|Y]$. Given a posted price p, maximization of the consumer's flow payoff (1) yields a demand of unit slope $Q(p) = \theta - p$. The static equilibrium quantity and price are then given by $Q = \theta - M/2$ and $P = M/2$, yielding the following ex ante surplus levels:

$$\Pi_Y^{\text{static}} = \frac{\mu^2}{4} + \text{Var}[P] \quad \text{and} \quad CS_Y^{\text{static}} = \frac{1}{2}\text{Var}[\theta] + \frac{\mu^2}{8} - \frac{3}{2}\text{Var}[P]. \tag{5}$$

In this static version of our linear-demand model, the average price and quantity across all signal realizations are independent of the information structure (Schmalensee, 1981). Thus, by allowing the firm to better tailor its price to the demand θ, better information unequivocally raises profits: this is captured in the ex ante variance of the monopoly price, $\text{Var}[P]$, which measures the firm's ability to price discriminate. From the consumer's perspective, more precise information results in a higher degree of correlation between her type and the price, which reduces her expected surplus. In other words, we have deliberately chosen a setting where any benefit to consumers is derived from dynamic incentives.[14]

Now suppose the consumer interacts with two firms. Firm 1 sets a price P_1 using the common prior, while firm 2 observes a signal of the consumer's first-period quantity before choosing its price $P_2 = M/2$. Importantly, because the consumer realizes the impact of her first period quantity on the second period price, she attempts to manipulate firm 2's beliefs M downwards by adopting a lower demand function in the first period than in the static benchmark. In equilibrium, however, firm 1 anticipates the consumer's manipulation incentives: but facing a lower demand function, it lowers its price. In other words, firm 1 must compensate the consumer for the information she reveals through her actions, because firm 2 will use that information to price discriminate. Figure 1 illustrates this outcome.[15]

The outcome of a two-period interaction therefore features a consumer buying a lower quantity than in a static interaction without information, $Q' < \theta - \mu/2$, but who also pays a lower price, $P' < \mu/2$. It is intuitive that a small amount of demand reduction is unambiguously beneficial in the first period: the consumer gives up the marginal unit of consumption but receives an infra-marginal discount. Thus, the incentives to manipulate the firms' beliefs induce a tradeoff between lower prices today and tailored prices tomorrow.

The effects of strategic demand reduction on price levels are well-understood in the literature on behavior-based price discrimination (Taylor, 2004). In our setting where purchases are observed with noise, however, the consumer's value of manipulating the firm's beliefs is central to both the equilibrium price level and the *firms' information*. In particular, higher

[14] See Bergemann, Brooks, and Morris (2015) for a complete analysis of third-degree price discrimination in the static case.

[15] In Figure 1, the ratchet effect results in a parallel downward shift of the demand function. This is a property of equilibrium in our continuous-time setting, but not in a two-period model.

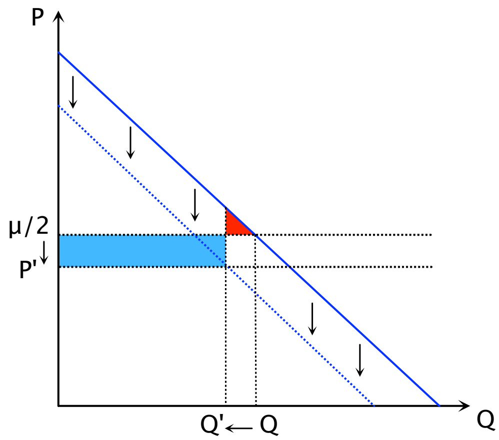

Figure 1: Strategic Demand Reduction

types θ (i.e., consumers who buy more units in period 2) stand to gain more by shading their period-1 demands. This effect reduces the responsiveness of actions to the consumer's unobserved willingness to pay, and hence, it lowers the informativeness of the first-period signal, relative to a setting without information spillovers. In other words, the sensitivity of the value of manipulation to the consumer's type θ is crucial to determine the amount of information available to the firms in equilibrium.

Having developed some intuition for the drivers of strategic demand reduction, we now turn to the role of score persistence and transparency in shaping the equilibrium outcomes. In particular, we want to distinguish between a small, but very persistent, increase in prices, and a large, but transient, price shock. This analysis requires a fully dynamic model where the information about the consumer's purchase history is long-lived, and hence, the value of strategic demand reduction is determined in equilibrium manipulation (i.e., it is not a function of second-period parameters only). The stationary setting that we analyze next is the simplest dynamic model that captures both features in a parsimonious way.

4 Equilibrium Analysis

We begin our equilibrium analysis by studying how each monopolist learns about the consumer from the observation of the contemporaneous level of the consumer's score (Section 4.1), and how it chooses a monopoly price (Section 4.2).

4.1 Learning under Linear Strategies

Along the path of play of a stationary linear Markov equilibrium, the consumer's quantity demanded at time t, Q_t, is a linear function of the contemporaneous pair (θ_t, Y_t), $t \geq 0$. As a result, the process $(\theta_t, Y_t)_{t\geq 0}$ evolves according to a linear stochastic differential equation (SDE) with a normally distributed initial condition, and so it is a Gaussian process.[16]

Importantly, since the time-t monopolist does not observe deviations by the consumer or previous monopolists, (θ_t, Y_t) is also Gaussian from her perspective. Thus, her posterior belief $\theta_t|Y_t$ is normally distributed with a mean and variance given by

$$\mathbb{E}[\theta_t|Y_t] = \mu + \frac{\text{Cov}[\theta_t, Y_t]}{\text{Var}[Y_t]}[Y_t - \mathbb{E}[Y_t]] \text{ and } \text{Var}[\theta_t|Y_t] = \text{Var}[\theta_t] - \frac{\text{Cov}^2[\theta_t, Y_t]}{\text{Var}[Y_t]}.$$

Furthermore, the stationarity requirement implies that the affine relationship between the posterior mean and the score, as well as the posterior variance, are time-independent. In what follows, we focus on the posterior mean to set up the consumer's best-response problem, and defer the discussion of the posterior variance to Section 5 in which we discuss learning.

Let $M_t := \mathbb{E}[\theta_t|Y_t]$, $t \geq 0$. By the previous discussion, we can characterize outcomes of stationary linear Markov equilibria using the firms' posterior belief process $(M_t)_{t\geq 0}$ instead of $(Y_t)_{t\geq 0}$. Specifically, we aim to find coefficients α, β and δ such that

$$Q_t = \alpha\theta_t + \beta M_t + \delta\mu, \; t \geq 0, \tag{6}$$

corresponds to the *quantity demanded along the path of play.* Such parameters will in turn depends on the degree of (inverse) persistence of the score $\phi > 0$.

The next result characterizes stationary Gaussian pairs $(\theta_t, Y_t)_{t\geq 0}$ when the quantity demanded is given by (6), with the consistency requirement that $M_t = \mathbb{E}[\theta_t|Y_t]$.

Proposition 1 (Stationarity and Beliefs)**.** *Suppose the process $(\theta_t, Y_t)_{t\geq 0}$ is stationary Gaussian when the quantity demanded is given by (6). Then:*

(i) $M_t = \mu + \lambda[Y_t - \bar{Y}]$, with $\bar{Y} = \mu(\alpha + \beta + \delta)/\phi$ and

$$\lambda = \frac{\alpha\sigma_\theta^2(\phi - \beta\lambda)}{\alpha^2\sigma_\theta^2 + \sigma_\xi^2\kappa(\phi - \beta\lambda + \kappa)}; \tag{7}$$

(ii) the score process (4) is mean reverting: $\phi - \beta\lambda > 0$;

[16] The technical details are in the proof of Proposition 1, which is stated below.

(iii) $(\theta_0, Y_0) \sim \mathcal{N}([\mu, \bar{Y}]^\top, \Gamma)$ *is independent of* $(Z_t^\theta, Z_t^\xi)_{t\geq 0}$*, where the long-run covariance matrix* Γ *is given in (A.4).*

Conversely, if $Q_t = \alpha\theta_t + \beta(\mu + \lambda[Y_t - \bar{Y}]) + \delta\mu$*, with* λ *and* $\bar{Y}$ *as in (i), and (ii) and (iii) hold, then* $(\theta_t, Y_t)_{t\geq 0}$ *is stationary Gaussian and* $M_t = \mu + \lambda[Y_t - \bar{Y}]$*,* $t \geq 0$*.*

When $(\theta_t, Y_t)_{t\geq 0}$ is stationary Gaussian under (6), $(\theta_t, Y_t) \sim \mathcal{N}([\mu, \bar{Y}]^\top, \Gamma)$ for all $t \geq 0$, for some scalar $\bar{Y} := \mathbb{E}[Y_t]$, $t \geq 0$, and some covariance matrix Γ, both depending on (α, β, δ). Part (i) in the proposition then states an equation for the *sensitivity of beliefs to changes in the score*, λ: as a regression coefficient, (7) is simply $\lambda = \text{Cov}[\theta_t, Y_t]/\text{Var}[Y_t] = \Gamma_{12}/\Gamma_{22}$ once we account for the fact that Γ depends on the weight λ that past firms have placed on the score to form their beliefs using the quantity demanded (6). Regarding (ii), observe that the effective rate of decay of $(Y_t)_{t\geq 0}$ under (6) is $\phi - \beta\lambda$ due to the contribution of Y_t to Q_t via $M_t = \mu + \lambda[Y_t - \bar{Y}]$; in this line, non-positive values of this rate would imply that the scores process' long-run distribution is improper, which violates stationarity.[17] Finally, equipped with a proper long-run distribution, (iii) says that if this one is to prevail from time zero, (θ_0, Y_0) must be drawn from the same distribution.

An implication of Proposition 1 is that the consumer can control the future firms' beliefs by affecting the evolution of the score. Specifically, since $M_t = \mu + \lambda[Y_t - \bar{Y}]$ holds path-by-path of $(Y_t)_{t\geq 0}$, given any admissible strategy $(Q_t)_{t\geq 0}$, the law of motion of $(M_t)_{t\geq 0}$ is

$$dM_t = \left[-\phi\left(M_t - \mu + \lambda\bar{Y}\right) + \lambda Q_t\right]dt + \lambda\sigma_\xi dZ_t^\xi,\ t \geq 0. \tag{8}$$

Thus, the consumer's choice of quantity Q_t affects firm t's beliefs linearly with slope λ, but this effect decays at rate ϕ. We can now formalize the key tradeoff between *persistence* and *sensitivity* of the firms' beliefs.

Lemma 1 (Persistence and Sensitivity)**.** *Suppose that* $\alpha > 0$*. Then the regression coefficient* λ *that solves* (7) *is strictly increasing in* ϕ*.*

Intuitively, the faster the score discounts old information, the easier it is to manipulate the firm's beliefs in the short run, and vice-versa. This tension underscores all the welfare and information properties of our equilibrium. We now connect the manipulation of beliefs to the manipulation of prices through the analysis of the time-t monopolist's problem.

4.2 Monopoly Pricing

In order to solve firm t's pricing problem, we must specify the sensitivity of the consumer's demand to intra-temporal price variation, or own-price sensitivity of demand: the weight

[17] In particular, (ii) and (7) imply that $\lambda > 0$ when $\alpha > 0$.

that a candidate equilibrium linear Markov strategy attaches to the contemporaneous price.

Crucially, because the score Y_t is publicly observed and firms adopt linear strategies $P_t = P(Y_t)$, the consumer is able to perfectly anticipate the equilibrium price. The sensitivity of the consumer's demand is then equivalent to the (optimal) change in her quantity demanded in response to a price deviation $p \neq P_t$. This poses a challenge in continuous time, because optimality of the consumer's strategy does not pin down the response to a deviation by a single (zero-measure) firm. Indeed, such a deviation does not affect the consumer's continuation payoff, and hence, every response is trivially optimal.[18]

To overcome this challenge, we refine our stationary linear Markov equilibrium concept: we impose the condition that prices are supported by the *limit sensitivity* of demand along a natural sequence of discrete-time analogs of our continuous-time model indexed by their period length. Along this family, as the period length shrinks to zero, the limit demand sensitivity is equal to -1.

Heuristically, consider a discrete-time version of our model in which the period length is given by $\Delta > 0$ small. Given any posted price p, we write the consumer's continuation value V_t recursively with M_t as a state as follows:

$$\begin{aligned} V_t &= \max_q \left[(\theta_t - p)q - \frac{q^2}{2}\right]\Delta + \mathbb{E}_t[V_{t+\Delta}] \\ &= \max_q \left[(\theta_t - p)q - \frac{q^2}{2}\right]\Delta + V_t + \frac{\partial V_t}{\partial M_t}\underbrace{\left[-\phi\left(M_t - \mu + \lambda\bar{Y}\right) + \lambda q\right]\Delta t}_{=\mathbb{E}[\Delta M_t] \text{ from } (8)} + \text{other terms.} \end{aligned}$$

When Δ is sufficiently small, the remaining terms that are affected by q on the right-hand side have only second-order effects on the consumer's payoff, and so the impact of quantities on the continuation value becomes asymptotically linear. Furthermore, because the noise in the score prevents price deviations from being observed by future firms, the continuation game is unaffected by the current price p. The consumer's best reply is then given by

$$Q_t = \theta_t - p + \lambda\frac{\partial V_t}{\partial M_t},$$

where $\partial V_t/\partial M_t$ is independent of the posted price. Since p was arbitrary, it follows that -1 is the own-price sensitivity of demand.[19] A key advantage of the continuous time is that it allows us to pin down this sensitivity independent of the rest of the equilibrium coefficients, which simplifies the analysis. The details of the formal argument can be found in the Online

[18]Similar challenges arise in other continuous-time games with sequential moves. See, for example, Kuvalekar and Lipnowski (2018).

[19]In particular, observe that the incentives to manipulate the firms' beliefs affect the intercept but not the slope of the demand function, as in Figure 1.

Appendix.[20]

We use our refinement to characterize the monopoly price process in the next result.

Lemma 2 (Monopoly Price)**.** *Consider a stationary linear Markov equilibrium with unit demand sensitivity in which the quantity demanded follows (6). Then, prices are given by*

$$P_t = (\alpha + \beta)M_t + \delta\mu,\ t \geq 0. \tag{9}$$

The intuition is simple: because demand has unit slope, the monopoly price along the path of play of such an equilibrium satisfies $P_t = \mathbb{E}[Q_t|Y_t]$, $t \geq 0$.

Equipped with this result, we can formulate the consumer's best-response problem as

$$\max_{(Q_t)_{t\geq 0}} \mathbb{E}\left[\int_0^\infty e^{-rt}\left[(\theta_t - P_t)Q_t - \frac{Q_t^2}{2}\right]dt\right]$$

subject to

$$\begin{aligned} d\theta_t &= -\kappa(\theta_t - \mu)dt + \sigma_\theta dZ_t^\theta \\ dM_t &= (-\phi[M_t - \mu + \lambda\bar{Y}] + \lambda Q_t)dt + \lambda\sigma_\xi dZ_t^\xi \\ P_t &= (\alpha + \beta)M_t + \delta\mu, \end{aligned}$$

where where λ satisfies (7). To tackle this problem—and, a fortiori, the task of finding stationary linear Markov equilibria—we use dynamic-programming tools.

4.3 Stationary Linear Markov Equilibria

Let $V(\theta, M)$ denote the value of the consumer's best-response problem when the current value of the state is $(\theta, M) \in \mathbb{R}^2$. The HJB equation is given by

$$\begin{aligned} rV(\theta, M) = \sup_{q\in\mathbb{R}}\Big\{ &\big(\theta - \underbrace{[(\alpha+\beta)M + \delta\mu]}_{=P_t}\big)q - \frac{q^2}{2} - \kappa(\theta - \mu)\frac{\partial V}{\partial \theta} \\ &+ (\lambda q - \phi[M - \mu + \lambda\bar{Y}])\frac{\partial V}{\partial M} + \frac{\lambda^2\sigma_\xi^2}{2}\frac{\partial^2 V}{\partial M^2} + \frac{\sigma_\theta^2}{2}\frac{\partial^2 V}{\partial \theta^2}\Big\}, \end{aligned} \tag{10}$$

If (α, β, δ) is such that the policy delivered by the HJB equation (10) (subject to standard transversality conditions) coincides with (6), and the tuple $(\alpha, \beta, \delta, \lambda)$ satisfies the station-

[20] The sequence of discrete-time games examined employ the traditional discrete-time analog of $(\theta_t, Y_t)_{t\geq 0}$ that involves noise scaled by $\sqrt{\Delta}$. Along this sequence, we can show that (i) linear best replies along the path of play are also optimal after having observed off-path prices, and (ii) the weight that linear best replies attach to the current price converges to -1 as the period length goes to zero.

arity condition $\phi - \beta\lambda > 0$ in (ii) of Proposition 1, then the coefficients (α, β, δ) fully characterize the outcome (i.e., prices and quantities) of a stationary linear Markov equilibrium.

The combination of (i) quadratic flow payoffs and (ii) Gaussian types and shocks makes the analysis particularly tractable. In particular, the consumer's best-response problem is a linear-quadratic optimization problem. We then look for a quadratic value function

$$V(\theta, M) = v_0 + v_1\theta + v_2 M + v_3 M^2 + v_4\theta^2 + v_5\theta M \tag{11}$$

that solves (10), and hence, for a linear best response

$$Q(\theta, M) = \theta - [(\alpha + \beta)M + \delta\mu] + \lambda \frac{\partial V(\theta, M)}{\partial M}. \tag{12}$$

Observe that the consumer's optimal quantity is the combination of the myopic purchase level, $\theta - P$, and the marginal value of manipulating the firms' beliefs via the score, $\lambda \cdot \partial V(\theta, M)/\partial M$. The latter term corresponds to the consumer's strategic demand reduction: thus, it captures the strength of the *ratchet effect*.

Imposing the condition that the firms correctly anticipate the consumer's behavior and substituting into the value function, we can obtain a sub-system of equations for the equilibrium coefficients (α, β, δ). This system is then coupled with equation (7) to pin down the equilibrium sensitivity of beliefs λ. This procedure allows us to establish existence and uniqueness of an equilibrium. Furthermore, the equilibrium can be computed in closed form, up to the solution of a single algebraic equation for the coefficient α.

We denote by $\Lambda(\phi, \alpha, \beta) \in \mathbb{R}_+$ the unique positive solution to (7) for λ when $(\phi, \alpha, \beta) \in (0, \infty) \times [0, 1] \times (-\infty, 0)$; the closed-form expression can be found in (A.12) in the Appendix.

Theorem 1 (Existence and uniqueness)**.** *For any $\phi > 0$, there exists a unique stationary linear Markov equilibrium with unit demand. In this equilibrium, $0 < \alpha < 1$ is characterized as the unique solution to the equation*

$$x \;=\; 1 + \frac{\Lambda(\phi, x, B(\phi, x))B(\phi, x)x}{r + \kappa + \phi},\; x \in [0, 1]. \tag{13}$$

Moreover, $\beta = B(\phi, \alpha) \in (-\alpha/2, 0)$ and $\delta = D(\phi, \alpha) \in \mathbb{R}$, where $B : (0, \infty) \times [0, 1] \to \mathbb{R}$ and $D : (0, \infty) \times [0, 1] \to \mathbb{R}$ are defined in (A.11) and (A.15). Finally, $\lambda = \Lambda(\phi, \alpha, B(\phi, \alpha)) > 0$.

Figure 2 illustrates the equilibrium coefficients $(\alpha, -\beta, \delta)$, the average equilibrium price

$$\mathbb{E}[P_t] = [\alpha + \beta + \delta]\mu, \tag{14}$$

as well as their static benchmark levels. We discuss their properties in the next subsection.

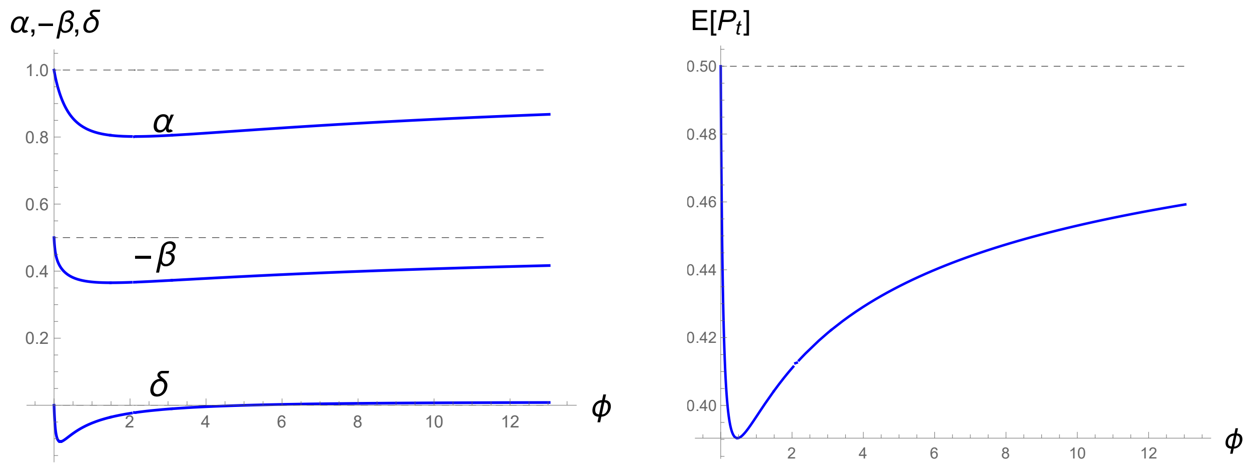

Figure 2: $(r, \sigma_\theta, \sigma_\xi, \kappa) = (1/10, 1, 1/3, 1)$.

4.4 Strategic Demand Reduction

The consumer's incentives to reduce her demand—which ultimately shape the equilibrium coefficients and prices—depend on both her current type θ and the firms' contemporaneous beliefs, M. In fact, having solved for the value function (11), we obtain the following linear expression for the *ratchet effect* term in the first-order condition (12):

$$\lambda \frac{\partial V(\theta, M)}{\partial M} = \frac{\lambda \alpha \beta}{r + \kappa + \phi} \theta - \frac{\lambda \beta (\alpha + \beta)}{r + 2\phi - \beta \lambda} M + 2\delta \mu. \tag{15}$$

To guide intuition, Proposition 2 introduces an alternate representation of the strength of the consumer's incentives to reduce demand.

Proposition 2 (Value of Future Savings)**.** *Equilibrium prices, $P_t = \delta\mu + (\alpha + \beta)M_t$, and quantities demanded, $Q_t = \delta\mu + \alpha\theta_t + \beta M_t$, $t \geq 0$, satisfy*

$$Q_t = \theta_t - P_t - \lambda \mathbb{E}_t \left[\int_t^\infty e^{-(r+\phi)(s-t)} (\alpha + \beta) Q_s ds \right], \ t \geq 0. \tag{16}$$

The last term on the right-hand side of (16) is the (equilibrium) *value of future savings*. By the Envelope Theorem, a marginal reduction in today's purchases along the "optimal trajectory" is equal to the net present value of future reduction in prices holding future quantities constant. In particular, lowering Q_t by one unit reduces the posterior belief M_{t+dt} by λ, and hence, the contemporaneous price P_{t+dt} by $\lambda(\alpha + \beta)$; the impact on subsequent prices then vanishes at the rate ϕ at which beliefs discount the initial change.[21]

In equilibrium, therefore, the marginal value of manipulating the firms' beliefs (15) corresponds to the value future savings: the larger the value of such savings, the stronger the incentive to engage in a downward deviation from the static Nash equilibrium; hence, the stronger the ratchet effect. Taking expectations under the prior distribution of $(\theta_t)_{t\geq 0}$ in (16), we obtain that the average strength of the ratchet effect is

$$\lambda \frac{\partial V(\mu,\mu)}{\partial M} = -\lambda \frac{(\alpha+\beta)(\alpha+\beta+\delta)\mu}{r+\phi}. \tag{17}$$

With non-negative average prices (i.e., $(\alpha+\beta+\delta)\mu > 0$—a feature of the equilibrium found), the right-hand side of (17) is strictly negative due to $\alpha+\beta > 0$, and hence the average quantity demanded contracts. Furthermore, because $\mathbb{E}[P_t] = \mathbb{E}[Q_t]$, the average equilibrium price lies below the static level $\mu/2$, as in Figure 2 (right panel).

Central to our analysis is how the incentives for demand reduction vary with the consumer's type, which is captured by the weight on θ in (15). Again using Proposition 2, the sensitivity of the value of future savings to the consumer's type can be written as

$$\frac{\partial}{\partial \theta}\left(-\lambda \mathbb{E}_t\left[\int_t^\infty e^{-(r+\phi)(s-t)}(\alpha+\beta)(\alpha\theta_s+\beta M_s+\delta\mu)ds \mid \theta_t=\theta\right]\right) = \frac{\lambda\alpha\beta}{r+\kappa+\phi}. \tag{18}$$

This ratio is in fact the last term of equation (13) that characterizes $\alpha(\phi)$ in Proposition 1, and hence, it affects the informativeness of the consumer's score. The denominator reflects that the impact of θ_t on future types $\theta_s,\ s>t$, depreciates at rate κ; thus, the impact of a marginal change in today's type on future savings decays at the augmented rate of $r+\kappa+\phi$. The numerator in turn reflects that an increase in the current type θ_t not only positively affects future types θ_s, but also future belief realizations M_s, $s>t$; in equilibrium, these two effects reduce to $\alpha\beta$ using expression (A.11).

We now use this intuition to derive properties of the equilibrium as a function of the persistence level, making explicit the dependence of (α,β,δ) on ϕ whenever required.

Proposition 3 (Equilibrium Properties)**.**

(i) $\lim_{\phi\to 0,+\infty}\alpha(\phi)=1$; $\lim_{\phi\to 0,+\infty}\beta(\phi)=-1/2$; $\lim_{\phi\to 0,+\infty}\delta(\phi)=0$; $\lim_{\phi\to 0,+\infty}\mathbb{E}[(P_t-\mu/2)^2]=0$.

[21] While we prove the result only in equilibrium, (16) is an optimality condition, and thus holds more generally. Specifically, if a best-response to a (potentially, non-Markov) price process exists, standard variational arguments show that (16) must hold.

(ii) for all $\phi > 0$,

$$1/2 < \frac{r+\kappa+\phi}{r+\kappa+2\phi} < \alpha(\phi) < 1; \;\; -\alpha(\phi)/2 < \beta(\phi) < 0; \;\; \textit{and } \mathbb{E}[P_t] \in (\mu/3, \mu/2) \;\; \textit{if } \mu \neq 0.$$

(iii) $\alpha(\cdot)$ *is strictly quasiconvex.*

(iv) $\alpha(\phi)$ *and* $\mathbb{E}[P_t]$ *are increasing in* σ_ξ/σ_θ *for all* $\phi > 0$.

Part (i) shows that the value of manipulation (15) vanishes when scores become uninformative: the equilibrium coefficients and price all converge to the static benchmark as $\phi \to 0$ and $+\infty$. In fact, if the score has no memory ($\phi \to \infty$), new information is forgotten instantaneously: in this case, the consumer's action has no impact on future prices almost surely, and hence, behaving myopically is optimal. On the other hand, if the score is fully persistent ($\phi \to 0$), it places an arbitrarily large weight on arbitrarily old information that is uncorrelated with the current type: in this case, the firms' beliefs are simply not sensitive to new information, yielding the same conclusion.[22]

Part (ii) formalizes the intuition that the benefits of a downward belief manipulation are larger for higher types: because types are persistent, a high θ_t type is more likely to buy larger quantities in the future, and hence, is more willing to invest in strategic demand reduction. Formally, because $\beta < 0$, the last term in (18) is strictly negative, and hence, $\alpha < 1$. Furthermore, the equilibrium coefficient α and the expected price level are also bounded from below: if firms conjectured $\alpha = 0$ then beliefs and prices would not respond to scores, and consumers would not attempt to manipulate them.

Part (iii) shows that the sensitivity of the value of manipulation to the type θ is strongest for an intermediate level of persistence ϕ. Consider (18) and recall the persistence vs. sensitivity tradeoff in Lemma 1: strategic demand reduction has a long-lasting effect for a low ϕ, i.e., if the score is persistent; but a low ϕ is associated with a low value of λ, and hence, a lower sensitivity of the value of manipulation to the consumer's type, everything else equal. Conversely, because $\phi \mapsto \lambda(\phi)$ is bounded due to the noise present in the score, the net present value of reducing prices vanishes for all types θ as ϕ grows without bound.

Finally, part (iv) considers the effects of noise on the value of demand reduction to contrast it with the effect of persistence. As the exogenous noise in the purchase signals increases, beliefs become less responsive to changes in the score, so the consumer's incentives to manipulate decrease and the equilibrium price rises in expectation. Furthermore, the incentives to manipulate decrease more quickly for higher types, which reduces the equilibrium

[22] As part of the proof of the proposition, we also show that $\lim_{\phi\to\infty} \text{Var}[Y_t] = 0$ and $\lim_{\phi\to 0} \lambda(\phi) = 0$.

α. In contrast, both the expected price level and the coefficient α attain their minimum at some interior values of ϕ.

5 Information Revelation

To quantify the extent of the firms' learning, we consider the variance of the firms' posterior beliefs

$$\mathrm{Var}[\theta_t|Y_t] = \mathrm{Var}[\theta_t]\left(1 - \frac{\mathrm{Cov}^2[\theta_t, Y_t]}{\mathrm{Var}[Y_t]\mathrm{Var}[\theta_t]}\right).$$

Thus, the amount of the firms' learning is captured by strength of the correlation between the consumer's type and the score. Given a persistence level ϕ and coefficients $(\alpha, \beta) \in \mathbb{R}_+ \times \mathbb{R}_-$ (i.e., not necessarily the equilibrium ones), if $(\theta_t, Y_t)_{t\geq 0}$ is stationary (Proposition 1),

$$\frac{\mathrm{Cov}^2[\theta_t, Y_t]}{\mathrm{Var}[Y_t]\mathrm{Var}[\theta_t]} = \frac{\alpha\Lambda(\phi,\alpha,\beta)}{\phi+\kappa-\beta\Lambda(\phi,\alpha,\beta)} := G(\phi,\alpha,\beta) \in [0,1]. \tag{19}$$

We refer to G as the *gain* function: with Gaussian signals, the coefficient of determination (R^2) measures how much information about the consumer's current type is gained by observing the contemporaneous score relative to the prior distribution.

A natural benchmark is the persistence level $\nu(\alpha, \beta)$ that maximizes the amount of learning for a given strategy of the consumer (6). Specifically, let

$$\nu(\alpha,\beta) := \arg\max_{\phi} G(\phi,\alpha,\beta) = \kappa + \frac{\gamma(\alpha)\alpha(\alpha+\beta)}{\sigma_\xi^2}, \tag{20}$$

where $\gamma(\alpha) > 0$ is the steady state variance of beliefs for an econometrician who has access to the histories $\xi^t := (\xi_s : 0 \leq s \leq t)$, $t \geq 0$.[23] Accordingly, we derive an equivalence result between observing a score of persistence $\nu(\alpha, \beta)$ and observing the full history of signals.

Proposition 4 (Learning under Public Histories)**.** *Consider a quantity process $(Q_t)_{t\geq 0}$ as in (6) with $\alpha > 0$ and $\beta < 0$, and such that $\nu(\alpha, \beta) > 0$.*

1. *If firms observe the histories of $(\xi_t)_{t\geq 0}$ and their beliefs are stationary, the posterior mean process is affine in a (stationary Gaussian) score (4) with $\phi = \nu(\alpha, \beta)$.*

2. *Conversely, if firms only observe the current value of a stationary Gaussian score with $\phi = \nu(\alpha, \beta)$, their beliefs coincide with an observer with access to the history of $(\xi_t)_{t\geq 0}$.*

[23] $\gamma(\alpha)$ is the unique positive root of the quadratic $x \mapsto \alpha^2 x^2/\sigma_\xi^2 + 2\kappa x - \sigma_\theta^2 = 0$. See Lemma 6 in the Appendix for a derivation of the optimal persistence $\nu(\alpha, \beta)$.

By Bayes' rule, if firm t could observe the entire history ξ^t, $t \geq 0$, and her belief were stationary, the posterior mean would aggregate past signals linearly with an exponential weight $\nu(\alpha, \beta)$. Thus, the stationary belief process of an outside observer with access to the full history of signals is an exponential score with persistence $\phi = \nu(\alpha, \beta)$; and conversely, those beliefs can be induced by an exponential score with the persistence $\phi = \nu(\alpha, \beta)$.

We now define scores that induce the same outcome as the observation of the full history of signals in *equilibrium*, i.e., when the consumer's behavior to is allowed to respond.

Definition 2 (Concealing Information)**.** *A score with persistence $\phi > 0$ does not conceal information about the consumer's* equilibrium *behavior if and only if it satisfies*

$$\phi = \nu(\alpha(\phi), \beta(\phi)). \tag{21}$$

In other words, if a score with persistence ϕ generates equilibrium coefficients (α, β) such that $\nu(\alpha(\phi), \beta(\phi)) = \phi$, then the aggregation of signals into a score does not conceal any further information about the consumer's history. We establish the existence and uniqueness of a fixed point in the next Proposition.[24]

Proposition 5 (Non-Concealing Score)**.**

(i) There exists a unique $\phi^* \in \mathbb{R}_+$ solving $\phi = \nu(\alpha(\phi), \beta(\phi))$.

(ii) The fixed point ϕ^* satisfies $\kappa < \phi^* < \sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2}$.

(iii) The coefficient $\alpha(\cdot)$ is strictly decreasing at the fixed point $\phi = \phi^*$.

Combined with Proposition 4, part (i) indirectly establishes the uniqueness of stationary linear Markov equilibrium for the case of observable histories ξ^t, $t \geq 0$. To build intuition, consider the consumer's best reply to the firms' conjecture about her strategy when such histories are public: if the firms expect low sensitivity of quantities to the underlying types, they also view the signals as uninformative, but the consumer then has no reason to manipulate. The opposite holds if the firms assign a large weight to the purchase signals. Thus, the firms' conjecture and the sensitivity of the consumer's actual behavior are strategic substitutes.

Part (ii) is intuitive: beliefs discount past signals more heavily than the type process discounts past shocks to taste, reflecting the identification problem faced by the firms. The upper bound follows from α—and hence, the informativeness of $(\xi_t)_{t\geq 0}$—being also bounded.

[24]In the career concerns model of Hörner and Lambert (2017), the agent has no private information on the equilibrium path. The existence of a non-concealing score is then immediate because the persistence of the firms' beliefs does not depend on the agent's strategy.

Part (iii) suggests that the degree of persistence ϕ^* that conceals no information about the consumer's behavior does not, however, maximize learning in equilibrium. This is due to the effect of the score's persistence on the consumer's behavior. In particular, a lower persistence level $\phi < \phi^*$ induces the consumer to be more responsive to her current type (i.e., a larger α), and hence, increases the informativeness of the purchase signals. As a result, the precision of the firms' beliefs may increase, even if the score conceals some of the information contained in the signals.

We formalize this intuition in the next result where, without fear of confusion, we let $G(\phi) := G(\phi, \alpha(\phi), \beta(\phi))$, $\phi \in [0, \infty)$, denote the *equilibrium gain* function.

Proposition 6 (Maximizing Learning). *$\phi \in [0, \infty) \mapsto G(\phi) \in [0, 1]$ is maximized in $(0, \phi^*)$.*

By definition of ϕ^*, we have $G_\phi(\phi^*, \alpha(\phi^*), \beta(\phi^*)) = 0$, and hence, changing the persistence of the score has only a second-order effect on learning, holding α and β constant. Marginally increasing β at $(\phi^*, \alpha(\phi^*), \beta(\phi^*))$ has no first-order effect on the amount of information transmitted either: this is because β corresponds to the coefficient on the firm's beliefs in the consumer's strategy, and at ϕ^*, the score perfectly accounts for the contribution of beliefs to contemporaneous recorded purchases. Increasing α, however, has a first-order positive effect on learning, as the score is now more sensitive to the consumer's type.

Figure 3 plots the equilibrium gain G as a function of the difference $\phi - \nu(\phi)$: its maximum is located to the left of the vertical axis that corresponds to the fixed point of the $\nu(\cdot)$ map, i.e., to the fully revealing score ϕ^*.

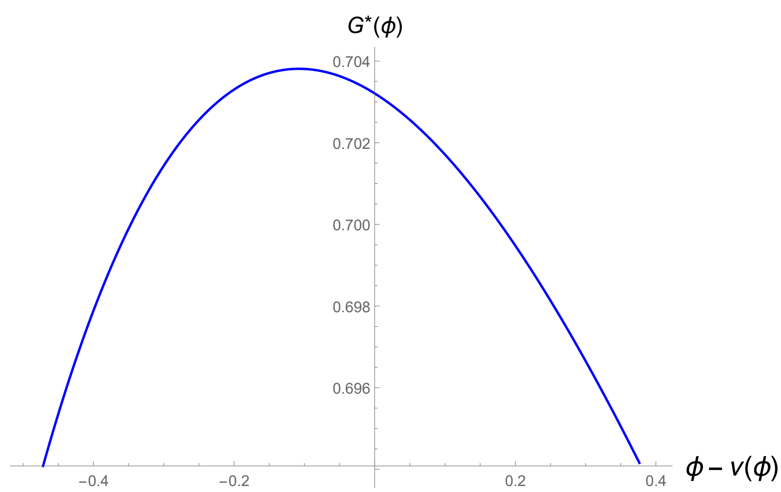

Figure 3: $(r, \sigma_\theta, \sigma_\xi, \kappa) = (3, 1, 2, 1)$.

To understand why a reduction in ϕ from ϕ^* increases the equilibrium coefficient α, recall that an increase in the score's persistence has two effects. First, because beliefs are an affine function of the score, beliefs also acquire more persistence. Thus, any change in the score resulting from a change in demand has a more prolonged effect on prices, which makes the consumer more eager to manipulate her own score (i.e., to scale back more). Second, because now the score attaches an excessive weight to past signals, it correlates less with the consumer's current type: in other words, Bayesian updating punishes the additional persistence by reducing the sensitivity of beliefs (and prices) to changes in the score. This in turn makes the consumer less concerned about purchasing large quantities.[25]

One would expect the consumer's discount rate to non-trivially affect the relative strength of these effects on average. However, (iii) in Proposition 5 holds *for all discount rates*, meaning that the "score-sensitivity effect" is relatively stronger for all $r > 0$. In particular, even a very patient consumer finds it optimal to attach a higher weight on her type, despite the consequences that more persistent scores can have on long-term prices.

To understand the reason behind the result, recall that the coefficient α does not capture the average value of future savings. Instead, α reflects the *relative* value of future savings for a marginally higher type θ_t, as derived in (18). From Section 4.4, the sensitivity of the value of future savings to the consumer's type satisfies

$$\frac{\partial^2 V(\theta, M)}{\partial\theta\partial M} \propto \int_t^{\infty} \lambda e^{-(s-t)(r+\phi+\kappa)} ds = \frac{\lambda}{r+\phi+\kappa}.$$

As explained earlier, this expression is the total value of the impulse response triggered by a shock to θ_t: a combination of the direct impact on future types (which decays at rate κ) and the indirect impact on future prices (which depreciates at rate $r + \phi$) via the change in the quantity demanded.

With linear dynamics, however, covariances—and hence, firms' learning—are also determined by impulse responses. In fact, recalling that $\lambda = \text{Cov}[\theta_t, Y_t]/\text{Var}[Y_t]$ and $G(\phi) \propto \lambda \cdot \text{Cov}[\theta_t, Y_t]$, the linearity of the covariance operator yields (up to constant terms)

$$\text{Cov}[\theta_t, Y_t] = \int_{-\infty}^{t} e^{(s-t)(\phi-\beta\lambda+\kappa)} ds = \frac{\lambda}{\phi-\beta\lambda+\kappa} \Rightarrow \lambda \cdot \text{Cov}[\theta_t, Y_t] = \frac{\lambda}{\phi-\beta\lambda+\kappa},$$

where $e^{-(s-t)(\phi-\beta\lambda)}$ is the weight of past scores Y_s on the current score Y_t and $e^{-(s-t)\kappa}$ is the covariance of (θ_s, θ_t) up to a multiplicative constant.[26] Intuitively, a shock to a past type θ_s, $s < t$, has an impact proportional to λ on the past score Y_s, which depreciates at rate $\phi - \beta\lambda$. But the type shock itself depreciates at rate κ, and hence, enters the covariance of

[25] The trade-off between persistence and sensitivity also arises in signal-jamming models with symmetric uncertainty. See, for instance, Cisternas (2017a) in the context of career concerns.

[26] This expression is written "as if" the game had an infinite past. An identical expression can be obtained by integrating on $[0, t]$ and using the covariance of the stationary distribution of (θ_0, Y_0).

θ_t and Y_t with weight $e^{-(s-t)(\phi-\beta\lambda+\kappa)}$.

Consider now a marginal change in ϕ holding $(\alpha, \beta) \in [0,1] \times \mathbb{R}_-$ fixed. In this case, the only change in the previous impulse responses is that λ becomes $\Lambda(\phi,\alpha,\beta)$, the unique positive root to (7) when $\alpha > 0$ and $\beta < 0$. Thus,

$$\nu(\alpha,\beta) = \arg\max \frac{\Lambda(\phi,\alpha,\beta)}{\kappa+\phi-\beta\Lambda(\phi,\alpha,\beta)} = \arg\max \frac{\Lambda(\phi,\alpha,\beta)}{\kappa+\phi} < \arg\max \frac{\Lambda(\phi,\alpha,\beta)}{r+\kappa+\phi},$$

which means that $\partial^2 V/\partial\theta\partial M$ falls whenever $G(\phi,\alpha,\beta)$ peaks (i.e., at $\phi = \nu(\alpha,\beta)$).

In other words, with a linear relationship between $(Y_t)_{t\geq 0}$ and $(\theta_t)_{t\geq 0}$, the covariance between Y_t and θ_t is affine in the sensitivity of the former to the latter. However, because both processes have persistence, such sensitivity is in the form of an impulse response. With Gaussian learning, moreover, such impulse response shapes the problem of maximizing learning. Critically, since the price, via the belief, is itself a linear function of the score, the problems of maximizing learning and of maximizing the value of future savings to shocks to θ_t (for a fixed consumer behavior) differ only in that discounting gives the immediate future more relevance in the latter problem. As a result, the sensitivity-persistence tradeoff is tilted in the favor of the sensitivity effect, thus leading to lower α when when $\phi > \phi^*$, i.e., when the impact of quantities on future prices is *frontloaded* relative to the Bayesian benchmark.

6 Optimal Persistence

In this section, we examine the welfare implications of a score persistence ϕ. In particular, since $(\alpha(\phi), \beta(\phi), \delta(\phi), G(\phi))$ converges to $(1, -1/2, 0, 0)$ as $\phi \to 0$, the equilibrium found converges to the repetition of static Nash when firms price using the prior distribution only. Thus, we can include $\phi = 0$ as part of the policy space, understood as the case in which firms possess no information about the consumer.

We begin with the firms' profits. Omitting the dependence of P_t and M_t on ϕ, and using that $\mathbb{E}[Q_t|Y_t] = P_t$, firm t's ex ante profits are given by

$$\Pi(\phi) := \mathbb{E}[P_t Q_t] = \mathbb{E}[P_t^2] = \mathbb{E}[P_t]^2 + \text{Var}[P_t],\ t \geq 0. \tag{22}$$

In other words, a firm's profits depend on the expected price level, on the amount of information available in equilibrium, and on its ability to tailor prices to that information.

By the projection theorem for Gaussian random variables, the ex ante variance of the firms' price is proportional to the variance of the fundamental, scaled by the equilibrium gain factor. We then have

$$\mathbb{E}[P_t] = (\alpha(\phi) + \beta(\phi) + \delta(\phi))\mu \quad \text{and}$$
$$\text{Var}[P_t] = (\alpha(\phi) + \beta(\phi))^2 \text{Var}[M_t] = (\alpha(\phi) + \beta(\phi))^2 \text{Var}[\theta_t] G(\phi),$$

because larger values of $G(\phi)$ indicate that a greater fraction of more of the variability of the consumer's type is inherited by the posterior mean from an ex ante perspective.

The variance of the equilibrium price is not only a driver of firm profits; it also captures the equilibrium value of information. In particular, it can be written as

$$\text{Var}[P_t] = \mathbb{E}[P_t Q_t] - \mathbb{E}[P_t] \cdot \mathbb{E}[Q_t].$$

In other words, the value of information is the difference between equilibrium profits and the profits from pricing under knowledge of the prior distribution and the equilibrium coefficients, when the consumer follows her equilibrium strategy. We now establish properties of the persistence level that maximizes the value of information.

Proposition 7 (Value of Information)**.**

(i) $\text{Var}[P_t]$ *is maximized to the left of* ϕ^*.

(ii) If $r \geq \kappa$, *then* $\text{Var}[P_t]$ *is maximized to the left of* $\phi^\dagger := \arg\max_{\phi \geq 0} G(\phi) < \phi^*$.

The variance of the equilibrium price attains its global maximum for $\phi < \phi^*$, i.e., for scores that are more persistent than the non-concealing one. Furthermore, if $r \geq \kappa$, the price variance-maximizing score is even more persistent than the learning-maximizing one, because further increasing persistence induces prices that are more responsive to beliefs.[27]

Consider now the case of the consumer. The (normalized) ex ante consumer surplus is defined by

$$CS(\phi) := r \cdot \mathbb{E}\left[\int_0^\infty e^{-rt}\{(\theta_t - P_t)Q_t - Q_t^2/2\}dt\right].$$

where $P_t = (\alpha + \beta)M_t + \delta\mu$ and $Q_t = (\alpha\theta + \beta M_t + \delta\mu$ for $t \geq 0$.[28] Moreover, as we show in the Online Appendix,

$$CS(\phi) = \mathbb{E}[P_t]\left(\mu - \frac{3}{2}\mathbb{E}[P_t]\right) + L(\phi)\text{Var}[\theta_t]G(\phi) + \alpha(\phi)\left(1 - \frac{\alpha(\phi)}{2}\right)\text{Var}[\theta_t], \tag{23}$$

[27] The condition $r \geq \kappa$ does not appear necessary from numerical examples.

[28] By stationarity, $CS(\phi) = \mathbb{E}\left[\mathbb{E}_t\left[\int_t^\infty e^{-r(s-t)}\{(\theta_s - P_s)Q_s - Q_s^2/2\}dt\right]\right]$ for all $t \geq 0$.

where,

$$L(\phi) := \frac{\alpha(\phi)^2}{2} + \beta(\phi) - \frac{3}{2}(\alpha(\phi) + \beta(\phi))^2 < 0, \quad \text{for all } \phi > 0.$$

Thus, the welfare consequences of persistence depend critically on three closely related *equilibrium* objects: the expected price level $\mathbb{E}[P_t]$; the sensitivity $\alpha(\phi)$ of the consumer's purchases to her type; and the firms' ability to learn the consumer's type in equilibrium, as measured by $G(\phi)$.

We know from the bounds in Proposition 3 that the first term in $CS(\phi)$ is positive and decreasing in the expected price. Thus, the ratchet-type forces identified here benefit the consumer through a lower price. Moreover, such benefit is higher when her average willingness to pay μ is high and discounts are applied to a larger number of units. Thus, the first term in (23) is proportional to μ^2.

Opposing this dynamic benefit are two forces. First, by shading down her demand ($\alpha(\phi) < 1$), the consumer moves away from her static optimum. This reduces consumer surplus, as reflected in the third term, which is increasing in α by Proposition 3. Second, the consumer transmits information about her willingness to pay to future firms. This makes the price positively covary with the consumer's type and reduces her surplus proportionally to the firms' information gain $G(\phi)$.

Unlike the expected price level, however, the costs of price discrimination for the consumer are independent of the average willingness to pay μ. It is then natural that consumers identified with high-μ segments benefit more from the availability of information than those coming low-μ counterparts, and that firms benefit from the score mechanism only if μ is low enough. The next result formalizes this intuition; to this end, let $\phi^f := \arg\max_{\phi \geq 0} \Pi(\phi)$ and $\phi^c := \arg\max_{\phi \geq 0} CS(\phi)$.

Proposition 8 (Optimal Persistence)**.**

(i) Informative optima: there exist $\underline{\mu}_f > 0$ *such that* ϕ^f *is interior if* $\mu < \underline{\mu}_f$*; and there exists* $\bar{\mu}_c > 0$ *such that* ϕ^c *is interior if* $\mu > \bar{\mu}_c$*.*

(ii) Uninformative optima: there exist $\underline{\mu}_c > 0$ *such that* $\phi^c \in \{0, \infty\}$ *if* $\mu < \underline{\mu}_c$*; and there exists* $\bar{\mu}_f > 0$ *such that* $\phi^f \in \{0, \infty\}$ *if* $\mu > \bar{\mu}_f$*.*

In the limiting case $\mu = 0$, the firm's optimal score is strictly positive and, under the conditions of Proposition 7, it is given by $\phi^f < \phi^\dagger$: for consumers with low average willingness to pay, the firms' ideal score is more persistent than the score that maximizes learning. Conversely, for consumers with high average willingness to pay, firms would prefer to commit to not observing any information, so that equilibrium prices rise to the static benchmark.

For the same reason, when μ is sufficiently large, consumers prefer an informative score to anonymous purchases. Figure 4 illustrates consumer surplus as a function of ϕ for different levels of μ, and Figure 6 compares profit levels for the same parameter values.

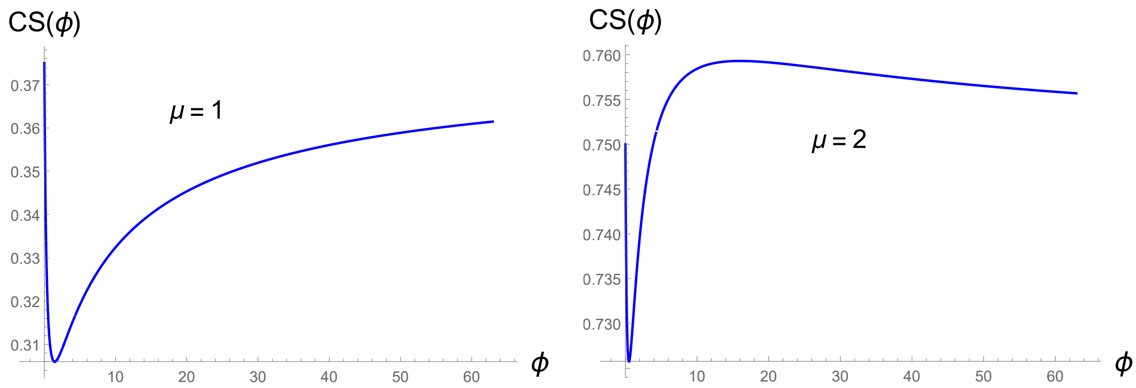

Figure 4: Consumer Surplus, $(\sigma_\theta, \sigma_\xi, \kappa, r) = (1, 2, 1, 1)$

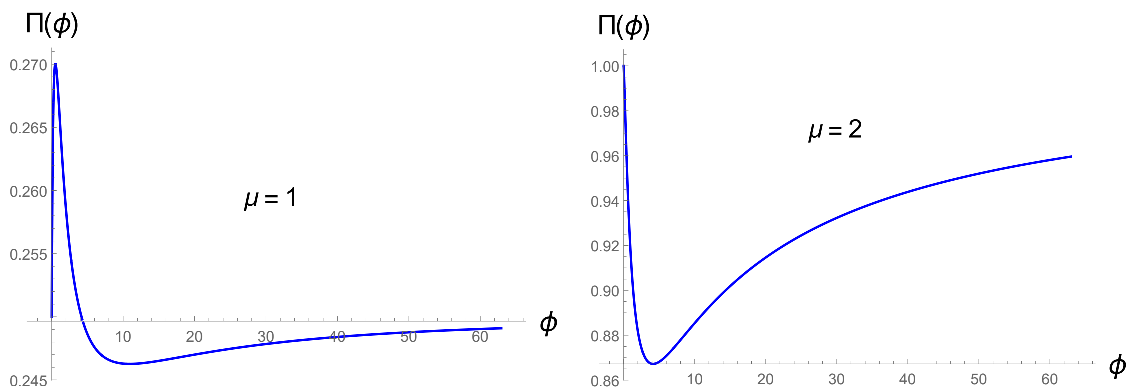

Figure 5: Producer Surplus, $(\sigma_\theta, \sigma_\xi, \kappa, r) = (1, 2, 1, 1)$

When interior, the consumer-optimal persistence level can be higher or lower than the fully revealing ϕ^*. In the case $\mu \to \infty$, the comparative statics of consumer surplus are driven by the expected price level, which in turn reflects the consumer's *average* value of future savings (16). Consistent with intuition, very patient consumers value having a long-term impact on the price, and hence, a very persistent score increases the benefit of reducing the quantity *today* to obtain a lower price. Conversely, impatient consumers prefer having an immediate effect on the price. Therefore, a higher score that is very sensitive to new information (but forgets it quickly) maximizes their value of manipulation and minimizes the price. Proposition (9) provides support for these claims.

Proposition 9 (Minimum of the Expected Price)**.** *Let $\alpha(\phi; r)$ denote the equilibrium coefficient on the type as a function of $r \geq 0$. Then,* *(i)* $\arg\min \mathbb{E}[P] \in (0, \arg\min \alpha(\phi; r))$*;*

(ii) There exists $0 < \underline{r} < \kappa$ *such that for all* $r < \underline{r}$*,* $\arg\min \mathbb{E}[P] < \phi^*(r) < \arg\min \alpha(\phi; r)$*;*

(iii) There exists $\bar{r} > \kappa$ *such that for all* $r > \bar{r}$ $\mathbb{E}[P_t]$ *is strictly decreasing at* $\phi^*(r)$*.*

Furthermore, numerical simulations show that the optimal persistence level is increasing in the discount rate r, with the consumer-optimal persistence level falling even below the learning-maximizing $\phi^\dagger$, which is lower than ϕ^*:

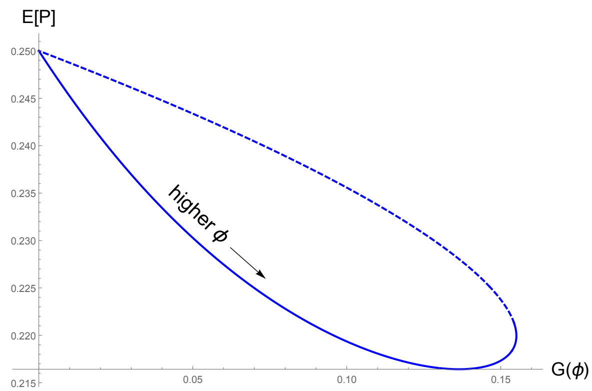

Figure 6: Average price minimizer vs. learning maximizer: $(\sigma_\theta, \sigma_\xi, \kappa) = (1, 2, 1)$; $r = 1/2$ (left); $r = 3$ (right). For a more patient consumer (left), the average price is minimized to the left of $\phi^\dagger$.

7 Hidden Scores

To study the role of transparency, we analyze the case in which the score Y_t observed by firm t is hidden to the consumer for all $t \geq 0$. The critical difference with the baseline model is that the firms' beliefs are private, and hence, prices carry information relevant to the consumer. In particular, by potentially signaling the current level of the firms' beliefs, today's price can provide information about future prices.[29]

Because in this context the consumer does not directly observe her score, we define a strategy for the consumer to be *linear Markov* if it is a linear function of (θ_t, p) only, where p is the contemporaneous price. In contrast, the notion of a linear strategy for the firms remains as in the baseline model. Our objects of interest are

$$\begin{aligned} Q(\theta, p) &= \delta^h \mu + \alpha^h \theta + \zeta^h p \text{ and} \\ P(Y) &= \pi_0^h + \pi_1^h Y, \end{aligned}$$

[29]The issue of transparency is moot in the career concerns model of Hörner and Lambert (2017) where, as in Holmström (1999), the *level* of the firms' beliefs is irrelevant for the worker's incentives.

where the superscript h stands for hidden. The corresponding concepts of admissible strategies, equilibrium, and stationarity are all straightforward modifications of the ones introduced in Section 2.[30] The focus is stationary linear Markov equilibrium.

Before turning to equilibrium analysis, we elaborate on two strategic implications of hidden scores. First, observe that ζ^h denotes the own-price sensitivity of demand. Critically, unlike in the observable-score case where price deviations must be examined in order to select a value for the sensitivity of the demand, we now determine ζ^h along the equilibrium path. The intuition comes from discrete time: if a score process is hidden and has full-support noise, (i) the consumer is not able to predict next period's price using today's observation, and (ii) any price realization is possible. Thus, the (discrete time) price process induced by a linear strategy exhibits the required intra-temporal variation along the path of play to enable the consumer's best-response problem to pin down the sensitivity of demand.[31]

Second, when scores are hidden, both the firms and the consumer can signal their private information. In particular, if (as we will show) $\alpha^h > 0$, we will have that $\pi_1 > 0$. The price is then a fully separating signal, i.e., the consumer perfectly learns her score along the path of play. Thus, on the equilibrium path, the consumer has the same information as in the observable case. Nonetheless, the signaling effect of prices deeply affects the consumer's incentives, as we illustrate next.[32]

Turning to the equilibrium analysis, notice that any equilibrium must entail $\zeta^h \neq 0$, and hence, firm t sets a price $P(Y_t) = -[\delta^h \mu + \alpha^h M_t(Y_t)]/[2\zeta^h]$. We therefore seek to characterize an equilibrium in which the on-path purchases process is of the form

$$Q_t = \delta^h \mu + \alpha^h \theta_t + \zeta^h \underbrace{\left[-\frac{\delta^h \mu + \alpha^h M_t}{2\zeta^h} \right]}_{P_t=} = \frac{\delta^h}{2}\mu + \alpha^h \theta_t + \beta^h M_t, \tag{24}$$

with $\beta^h := -\alpha^h/2$, and where $M_t = \mu + \lambda^h [Y_t - \bar{Y}^h]$, some λ^h and $\bar{Y}^h$, $t \geq 0$. In particular, realized prices and quantities satisfy $P_t = -\mathbb{E}[Q_t | Y_t]/\zeta^h$, $t \geq 0$, along the path of play.

[30] Concretely, there are only two changes: the suppression of Y_t in the notion of a linear Markov strategy for the consumer; and the conditioning on $(\theta_t, P_t)_{t\geq 0}$, rather than on $(\theta_t, Y_t)_{t\geq 0}$, by any admissible strategy. The latter change is innocuous in the consumer's best-response problem to a linear pricing strategy—we work with $(\theta_t, P_t)_{t\geq 0}$ as states for consistency only.

[31] As in the observable case, we view the program solved in this section as the limiting case of a sequence of discrete-time games in which the period length shrinks to zero under appropriately scaled noise. Directly in continuous time, however, the price process that results from a linear Markov pricing strategy will have continuous paths, so a deviation by a single firm can be detected. Because with full-support noise this issue arises only in continuous time, we are refining our equilibrium in the continuous-time game by assuming that the firms conjecture that the consumer responds to the deviation with a sensitivity that coincides with the sensitivity of the quantity demanded along the path of play of a candidate Nash equilibrium. Thus, as it occurs in discrete time, the same candidate policy $Q(\theta, p)$ is used by the firms in their pricing problem.

[32] For a model of hidden signals see Xu and Dukes (2017).

Since the quantity demanded (24) has the same structure as (6), the characterization of stationary beliefs in Proposition 1 applies to the hidden case with the additional restriction that $\beta = -\alpha^h/2$. Moreover, recall that when the score process is observed by the consumer, Proposition 1 reduces the quest for stationary linear Markov equilibria to a single equation (13) for α. With hidden scores, existence and uniqueness reduce to an almost identical equation for the coefficient on the consumer's type.[33]

Proposition 10 (Existence and Uniqueness)**.** *There exists a unique stationary linear Markov equilibrium. In this equilibrium, $\alpha^h \in (0,1)$ is the unique solution to*

$$x \;=\; 1 + \frac{\Lambda(\phi, x, -x/2)x[x/2]}{r+\kappa+\phi},\ x \in [0,1]. \tag{25}$$

Moreover, $\lambda^h = \Lambda(\phi, \alpha^h, -\alpha^h/2) > 0$, where Λ is defined in (A.12), and the own-price sensitivity of demand is given by

$$\zeta^h = -\frac{2(r+2\phi)}{2(r+2\phi)+\alpha^h\lambda^h} \in \left(-1, -\frac{r+2\phi}{r+3\phi}\right). \tag{26}$$

The first part of the proposition states that there is always a unique stationary linear Markov equilibrium, and the quantity demanded responds positively to changes in the consumer's type. Moreover, the coefficient α^h is determined by the same equation as α in the observable case replacing $B(\phi,\alpha) \in (-\alpha/2, 0)$ by $-\alpha/2$, $\alpha \in (0,1)$.

The second part of the Proposition 10 states a key result: demand is less sensitive to price in the hidden case than in the baseline model. The reason for this result lies in the informational content of prices when scores are hidden: by informing the consumer that her score is high, a high price today is a signal of high prices in the future, and hence a signal of a low value of reducing today's quantity demanded. This effect is due to the convexity property of the consumer's value as a function of the current price, which is analogous to the convexity of the indirect utility function in a static model: the advantage of reducing prices is greater when prices are low and the consumer is likely to buy more units in the near future. Everything else equal, therefore, the price-signaling effect reduces the ratchet effect (at any given price) relative to the observable-scores case.

The next result summarizes the economic implications hiding scores from consumers, and

[33] All formal proofs are in the Online Appendix. Appendix B, however, outlines some of the arguments, and lists further equilibrium properties that are direct corollaries from the baseline model (e.g., $\alpha^h(\phi)$ is quasiconvex, uniqueness of a non-concealing rating $\phi^{*,h}$, gain factor G maximized to the left of $\phi^{*,h}$, etc.).

hence, of lower demand sensitivity. To unify notation, we let

$$\begin{aligned} Q_t &= \delta^o\mu + \alpha^o\theta_t + \zeta^o P_t \\ P_t &= \pi_0^o + \pi_1^o Y_t \end{aligned} \tag{27}$$

denote the realized demand and prices along the path of play of a stationary linear Markov equilibrium (henceforth, an equilibrium) when the score is observed by the consumer.[34] Also, let $(\mathrm{Q}^o, \mathrm{P}^o)$ and $(\mathrm{Q}^h, \mathrm{P}^h)$ the average price-quantity pairs in the observable and hidden cases, respectively. Recall that all the coefficients are functions of ϕ.

Proposition 11 (Role of Transparency)**.** *In equilibrium, for all $\phi > 0$:*

(i) Sensitivity of demand to price: $0 > \zeta^h(\phi) > \zeta^o(\phi) > -1$.

(ii) Sensitivity of price to score: $\pi_1^h(\phi) > \pi_1^o(\phi) > 0$.

(iii) Sensitivity of demand to type: $1 > \alpha^o(\phi) > \alpha^h(\phi) > 0$

(iv) Average prices and quantities: $\mu/2 > \mathrm{P}^h(\phi) > \mathrm{P}^o(\phi) = \mathrm{Q}^o(\phi) > \mathrm{Q}^h(\phi) > \mu/4$.

To understand these results, begin with $0 > \zeta^h > -1$, i.e., demand is more inelastic when scores are hidden. Facing a demand with less own-price sensitivity, each firm charges a higher price that perfectly offsets the consumer's tendency to react less strongly to price increases (as in all linear-demand monopoly problems, the slope ζ^h cancels out in (24)). This implies that future prices become more sensitive to changes in the score relative to the observable case where demand has a unit slope. This also the case in equilibrium, reflecting that the demand sensitivity effect is always strong enough to offset the reduction in signaling.

With prices that are more sensitive to the score, the equilibrium amount of strategic demand reduction must increase ($\mathrm{Q}^o > \mathrm{Q}^h$ in part (iv)). By the envelope theorem, an analog of the representation result (16) in Proposition 2 holds for the hidden case:

$$Q_t = \theta_t - P_t - \pi_1^h \mathbb{E}_t\left[\int_t^\infty e^{-(r+\phi)(s-t)} Q_s ds\right],\ t \geq 0.$$

Because $\pi_1^h > \pi_1^o = \lambda(\alpha+\beta)$, the value of future savings increases (everything else equal), thereby inducing a lower average quantity demanded. Furthermore, since any change in the consumer's type translates into a stronger ratcheting of future prices in this case, the sensitivity of value of future savings to the consumer's type falls, which explains (iii). Conversely,

[34] It is easy to see that type-price space $\alpha^o = \alpha$, where α solves (13). The coefficient on the price $\zeta^o \neq -1$ combines the slope of demand (-1) and the change in the value of future savings associated with change in the current price, as described above for the hidden case.

since in the observable-score case prices are ratcheted up less strongly, the consumer has a relatively weak incentive to deviate from her myopic demand (which has unit sensitivity). As a result, the sensitivity of the value of future savings is relatively small in the observable case, which results in a corresponding sensitivity ζ^o of the *quantity demanded* that falls in between the two own-price sensitivities of demand, -1 and ζ^h (part (i)).

Finally, while average prices remain below the no-information (i.e., $\phi = 0$) case, they are higher than when scores are observed by the consumer ($\mu/2 > \mathrm{P}^h > \mathrm{P}^o$ in (iv)). This implies that the compensatory effect that emerges as a response to the consumer transmitting information to the firms, diminishes on average.

The previous properties on average prices and quantities, coupled with the lower own-price sensitivity of demand when scores are hidden, strongly suggest that firms are better off without transparency, while consumers worse off. While establishing a comparison of ex ante surpluses across cases is challenging due to the different nonlinear components affecting each one, numerical results nevertheless show that the ranking of prices and quantities in fact holds point-wise for all values of ϕ. We illustrate this in the next two figures.

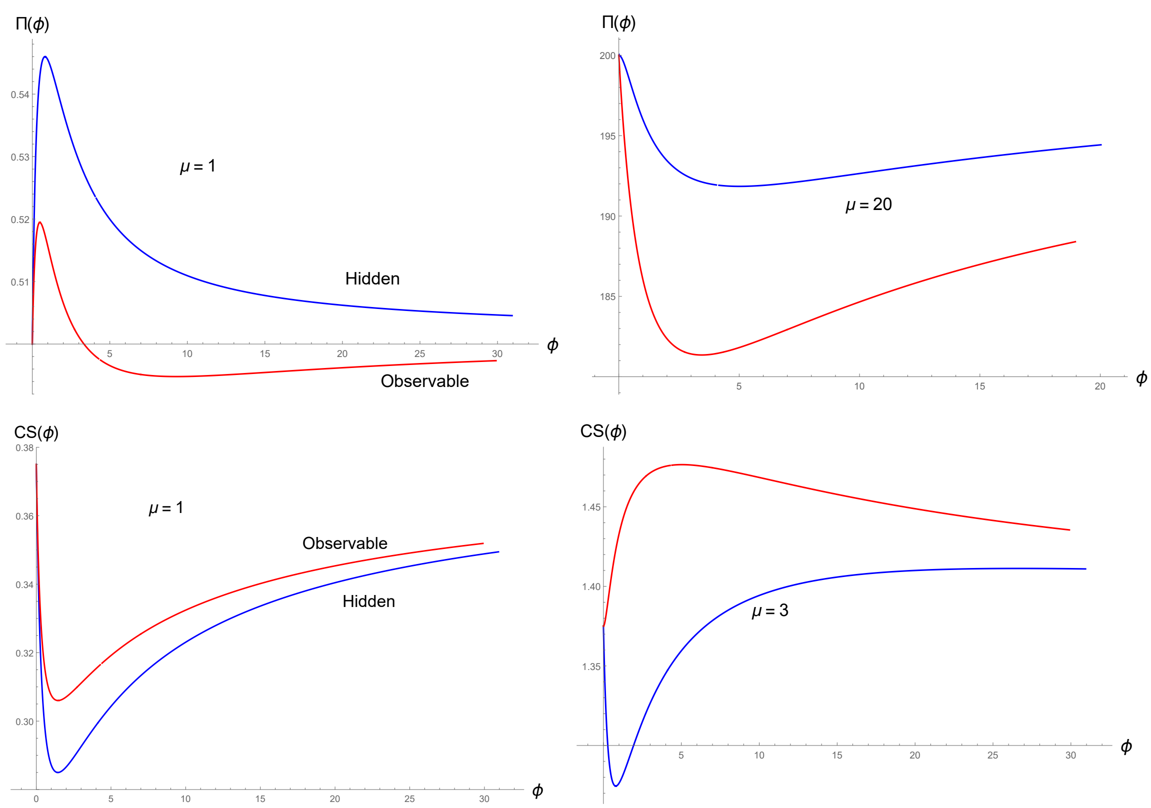

Figure 7: Producer and consumer surplus. Parameter values $\sigma_\xi = 1/3; \sigma_\theta = 1; r = 3; \kappa = 1$

8 Conclusion

We have explored the allocation and welfare consequences of score consumers based on their purchase histories and using the information so-gained to price discriminate. More specifically, we have focused on the effects of a score's persistence and transparency on the level and terms of trade when a consumer purchases a good from a different monopolist in every period.

Discussion of the Assumptions. First, instead of studying the reputation dynamics of a long-lived firm that practices behavior-based price discrimination, we focus on a consumer who is in the market for different products over time. In other words, the consumer interacts with a different (imperfectly competitive) industry in each period, and only returns to the same industry after a long time. In this sense, our simplifying assumptions are that the consumer is never in the market for the same product twice (i.e., firms are short lived), and that each firm is a local monopolist. Finally, we observe that the assumption of short-lived firms is useful to examine the role of scores but that, under public (noisy) purchase histories, our Markovian equilibrium remains such in a model with a single firm.[35]

Second, we paid attention to one-dimensional, continuous scores. The score Y_t, as opposed to a coarse categorization of the consumer's tastes, is an approximation of the products offered by big-data brokers. Brokers collect data about consumers' behavior from a wide variety of sources. Data brokers' clients (e.g., retail firms) can access information about several attributes for each individual prospect. However, even a multidimensional map of a consumer's preferences does not often contain detailed time-level information. Thus, the score does not correspond to a full history of consumers' actions. On the other hand, information on each attribute is collected in detailed categories (e.g., point estimates of annual income vs. coarse income brackets), which we approximate with a continuous score.

Third, the exponential score (4) is a tractable aggregator of past signals that retains some attractive properties: it has a natural counterpart in Bayesian updating when all signals are public, and it is analogous to stochastic forgetting. The latter property is perhaps a more realistic representation of the lifespan of Internet cookies than a deterministic duration.[36]

[35]In a long-lived firm's best-response problem to a linear Markov strategy of the consumer, the only state variable is her belief about the consumer's type. However, a choice of price today does not affect future beliefs, as the additivity of the price and the type in the consumer's strategy eliminates any experimentation effects. This in turn renders myopic behavior optimal.

[36]For example, different Internet browsers retain information for varying amounts of time, and different types of cookies have different automatic expiration dates.

Appendix A

Proofs for Section 4

Proof of Proposition 1. Suppose that $(\theta_t, Y_t)_{t\geq 0}$ is stationary Gaussian. By stationarity, $\mathbb{E}[Y_t]$ and $\text{Cov}[\theta_t, Y_t]/\text{Var}[Y_t]$ are independent of time; let $\bar{Y}$ and λ denote their respective values (to be determined). Moreover, by normality,

$$M_t := \mathbb{E}[\theta_t|Y_t] = \mu + \lambda[Y_t - \bar{Y}],\ t \geq 0.$$

Let $\hat{\delta} := \delta\mu + \beta(\mu - \lambda\bar{Y})$ and $\hat{\beta} = \beta\lambda$. We can then write the quantity demanded (6) as

$$Q_t = \delta\mu + \alpha\theta_t + \beta M_t = \hat{\delta} + \alpha\theta_t + \hat{\beta}Y_t,\ t \geq 0. \tag{A.1}$$

Using that $d\xi_t = Q_t dt + \sigma_\xi dZ_t^\xi$, we can conclude that $(\theta_t, Y_t)_{t\geq 0}$ evolves according to

$$\begin{aligned} d\theta_t &= -\kappa(\theta - \mu)dt + \sigma_\theta dZ_t^\theta, \\ dY_t &= [-(\phi - \hat{\beta})Y_t + \hat{\delta} + \alpha\theta_t]dt + \sigma_\xi dZ_t^\xi\ t > 0. \end{aligned} \tag{A.2}$$

The previous system is linear, and thus admits an analytic solution. Specifically, letting

$$X := \begin{bmatrix} \theta \\ Y \end{bmatrix},\ A_0 := \begin{bmatrix} \kappa\mu \\ \hat{\delta} \end{bmatrix},\ A_1 := \begin{bmatrix} \kappa & 0 \\ -\alpha & \phi - \hat{\beta} \end{bmatrix},\ B := \begin{bmatrix} \sigma_\theta & 0 \\ 0 & \sigma_\xi \end{bmatrix} \text{ and } Z := \begin{bmatrix} Z_t^\theta \\ Z_t^\xi \end{bmatrix},$$

we can write $dX_t = [A_0 - A_1 X_t]dt + BdZ_t$, $t > 0$, which has as unique (strong) solution

$$X_t = e^{-A_1 t}X_0 + \int_0^t e^{-A_1(t-s)}A_0 dt + \int_0^t e^{-A_1(t-s)}BdZ_s,\ t \geq 0, \tag{A.3}$$

where $e^{A_1 t}$ denotes the matrix exponential (Section 1.7 in Platen and Bruti-Liberati (2010)).

From the additive structure of (A.3), X_t is Gaussian for all $t \geq 0$ if and only if X_0 is Gaussian. But this implies that X_0 must be independent of $Z := (Z_t)_{t\geq 0}$ for Z to be a Brownian motion under the (null-sets augmented) filtration generated by Z and X_0.[37] Letting $\mathcal{N}(\vec{\mu}, \Gamma)$ denote the stationary distribution of X_t, $t \geq 0$, it follows that $\vec{\mu} \in \mathbb{R}^2$ and

[37] Denote such filtration by $(\mathcal{G}_t)_{t\geq 0}$. In the absence of independence, there must $t \geq 0$ such that Z_t is not independent of $\mathcal{G}_0$; but this violates the independent increments requirement of a Brownian motion.

the 2×2 covariance matrix Γ must satisfy the equations

$$\mathbb{E}[X_t] = \hat{\mu} \quad \Leftrightarrow \quad e^{-A_1 t}\hat{\mu} + [A_1^{-1} - e^{-A_1 t}A_1^{-1}]A_0 = \hat{\mu} \ \text{ and}$$
$$\text{Var}[X_t] = \Gamma \quad \Leftrightarrow \quad e^{-A_t t}\Gamma e^{-A_1^T t} + e^{-A_1 t}\text{Var}\left[\int_0^t e^{A_1 s} B dZ_s\right] e^{-A_1^T t} = \Gamma,$$

where $\text{Var}[\cdot]$ denotes the covariance matrix operator and T the transpose operator.

Observe that the first condition leads to $\vec{\mu} = A_1^{-1}A_0$ provided A_1 is invertible, which happens when $\phi - \beta\lambda \neq 0$—we assume this in what follows. Regarding the second condition, differentiating it and using that $\text{Var}\left[\int_0^t e^{A_1 s} B dZ_s\right] = \int_0^t e^{A_1 s} B^2 e^{A_1^T s} ds$ yields

$$-A_1\Gamma - \Gamma A_1^T + B^2 = 0.$$

Using that $\vec{\mu} = (\mathbb{E}[\theta_t], \mathbb{E}[Y_t])^T = (\mu, \bar{Y})^T$ and that $\Gamma_{11} = \text{Var}[\theta_t]$, $\Gamma_{12} = \Gamma_{21} = \text{Cov}[\theta_t, Y_t]$ and $\Gamma_{22} = \text{Var}[Y_t]$, it is then easy to verify that the previous system has as solution

$$\vec{\mu} = \begin{bmatrix} \mu \\ \frac{\hat{\delta}+\alpha\mu}{\phi-\hat{\beta}} \end{bmatrix} \quad \text{and} \quad \Gamma = \begin{bmatrix} \frac{\sigma_\theta^2}{2\kappa} & \frac{\alpha\sigma_\theta^2}{2\kappa(\phi-\beta\lambda+\kappa)} \\ \frac{\alpha\sigma_\theta^2}{2\kappa(\phi-\beta\lambda+\kappa)} & \frac{\alpha^2\sigma_\theta^2+\kappa\sigma_\xi^2(\phi-\beta\lambda+\kappa)}{2\kappa(\phi-\beta\lambda)(\phi-\beta\lambda+\kappa)} \end{bmatrix} \tag{A.4}$$

The last step required to confirm that the previous expressions indeed correspond to the first two moments of stationary Gaussian process is to verify that the covariance matrix Γ is positive definite. Since $\sigma_\theta^2/2\kappa > 0$, this condition reduces to

$$\det(\Gamma) > 0 \Leftrightarrow \frac{\sigma_\xi^2\kappa(\phi - \hat{\beta} + \kappa)^2 + \alpha^2\sigma_\theta^2\kappa}{(\phi - \hat{\beta} + \kappa)^2(\phi - \hat{\beta})} > 0 \Leftrightarrow \phi - \underbrace{\hat{\beta}}_{=\beta\lambda} > 0.$$

This proves (ii) and (iii). To finish the proof, we need to determine the scalars λ and $\bar{Y}$ that are consistent with Bayes' rule given a score process that is driven by the quantity demanded (6). Using (A.4), and after some simplification,

$$\lambda = \frac{\text{Cov}[\theta_t, Y_t]}{\text{Var}[Y_t]} = \frac{\alpha\sigma_\theta^2(\phi - \beta\lambda)}{\alpha^2\sigma_\theta^2 + \kappa\sigma_\xi^2(\phi + \kappa - \beta\lambda)} \quad \text{and} \tag{A.5}$$

$$\bar{Y} = \frac{\mu[\alpha + \beta + \delta]}{\phi} \tag{A.6}$$

where the last equality follows from $\bar{Y} = [\hat{\delta} + \alpha\mu]/(\phi - \beta\lambda)$ and $\hat{\delta} = \delta\mu + \beta(\mu - \lambda\bar{Y})$; this proves (i). The converse part of the Proposition is true by the previous constructive argument. This concludes the proof. □

Proof of Lemma 1. The result follows from partially differentiating (7) with respect to ϕ.□

Proof of Lemma 2. Consider a linear Markov strategy $Q(p, \theta, Y)$ for the consumer with weight equal to -1 on the contemporaneous price. Because the time-t monopolist assumes that past purchases followed (6), we have that $M_t = \mu + \lambda[Y_t - \bar{Y}]$, $t \geq 0$, where $\bar{Y}$ and λ are given in (i) in Proposition 1. Thus, we can write $Q(p, \theta_t, M_t) = q_0 + \alpha\theta_t + q_2 M_t - p$ for some coefficients q_0, α and q_2. Importantly, the weight that the strategy attaches to the contemporaneous price does not change under this linear transformation.

The monopolist operating at time t therefore solves

$$\max_p \; p\mathbb{E}[q_0 + \alpha\theta_t + q_2 M_t - p | Y_t] \Leftrightarrow P(M_t) = \frac{q_0}{2} + \frac{\alpha + q_2}{2} M_t,$$

which leads to a realized purchase

$$Q_t \;\; = \;\; q_0 + \alpha\theta_t + q_2 M_t - P(M_t) = \frac{q_0}{2} + \alpha\theta_t + \frac{q_2 - \alpha}{2} M_t, \; t \geq 0.$$

We conclude that when demand is linear with unit sensitivity, if realized purchases are given by $Q_t = \delta\mu + \alpha\theta_t + \beta M_t$, contemporaneous prices satisfy $P_t = \delta\mu + (\alpha + \beta)M_t$, $t \geq 0$. Importantly, once the coefficients (α, β, δ) are determined, simple algebra shows that prices are supported by the linear Markov strategy

$$Q(p, \theta_t, Y_t) = 2\delta\mu + [\mu - \lambda\bar{Y}][\alpha + 2\beta] + \alpha\theta_t + \lambda[\alpha + 2\beta]Y_t - p,$$

where λ satisfies (7). This conclude the proof. □.

Proof of Theorem 1. Under the set of admissible strategies defined in Section 2, Verification Theorem 3.5.3 in Pham (2009) applies. Specifically, we look for a quadratic solution $V = v_0 + v_1\theta + v_2 M + v_3 M^2 + v_4\theta^2 + v_5\theta M$ to the HJB equation (10)

$$\begin{aligned} rV(\theta, M) = \sup_{q \in \mathbb{R}} & \Big\{ (\theta - [(\alpha + \beta)M + \delta\mu])q - q^2/2 - \kappa(\theta - \mu)V_\theta \\ & [\lambda q - \phi\left(M - \mu + \lambda\bar{Y}\right)] \frac{\partial V}{\partial M}(\theta, M) + \frac{\lambda^2\sigma_\xi^2}{2}\frac{\partial^2 V}{\partial M^2} + \frac{\sigma_\theta^2}{2}\frac{\partial^2 V}{\partial \theta^2} \Big\} \end{aligned}$$

subject to standard transversality conditions, with the additional property that the optimal policy is of the form $\delta\mu + \alpha\theta + \beta M$.

The first-order condition reads

$$\begin{aligned} q &= \theta - [\delta\mu + (\alpha+\beta)M] + \lambda[v_2 + 2v_3 M + v_5\theta] \\ &= -\delta\mu + \lambda v_2 + [1+\lambda v_5]\theta + [2\lambda v_3 - (\alpha+\beta)]M \end{aligned}$$

which leads to the following system matching-coefficient conditions:

$$\delta\mu = -\delta\mu + \lambda v_2, \;\; \alpha = 1 + \lambda v_5, \;\; \text{and } \beta = 2\lambda v_3 - (\alpha+\beta). \tag{A.7}$$

By the Envelope Theorem, moreover,

$$\begin{aligned} (r+\phi)[v_2 + 2v_3 M + v_5\theta] &= -(\alpha+\beta)[\delta\mu + \alpha\theta + \beta M] - \kappa(\theta-\mu)v_5 \\ &\quad + [\lambda(\delta\mu + \alpha\theta + \beta M) - \phi(M - \mu + \lambda\bar{Y})]2v_3, \end{aligned} \tag{A.8}$$

which yields the following system of equations

$$\begin{cases} (r+\phi)v_2 = -(\alpha+\beta)\delta\mu + \kappa\mu v_5 + [\lambda\delta\mu + \phi(\mu - \lambda\bar{Y})]2v_3 \\ (r+2\phi)2v_3 = -(\alpha+\beta)\beta + 2v_3\lambda\beta \\ (r+\kappa+\phi)v_5 = -(\alpha+\beta)\alpha + 2v_3\lambda\alpha. \end{cases} \tag{A.9}$$

Using that v_2, v_3 and v_5 can be written as a function of α, β and $\delta\mu$, this system becomes

$$\begin{cases} (r+\phi)\frac{2\delta\mu}{\lambda} = -(\alpha+\beta)\delta\mu + \kappa\mu\frac{\alpha-1}{\lambda} + [\lambda\delta\mu + \phi(\mu - \lambda\bar{Y})]\frac{\alpha+2\beta}{\lambda} \\ (r+2\phi)\frac{\alpha+2\beta}{\lambda} = \underbrace{-(\alpha+\beta)\beta + \beta(\alpha+2\beta)}_{=(\beta)^2} \\ (r+\kappa+\phi)\frac{\alpha-1}{\lambda} = \underbrace{-(\alpha+\beta)\alpha + \alpha(\alpha+2\beta)}_{=\alpha\beta}. \end{cases} \tag{A.10}$$

where we have assumed that $\lambda \neq 0$. In fact, since $\phi - \beta\lambda > 0$ in any stationary linear Markov equilibrium, the equation for λ (i.e., (7)) implies that $\lambda \neq 0$ as long as $\alpha \neq 0$; but the latter is a corollary of the following lemma.

Lemma 3. *Any stationary linear Markov equilibrium must satisfy* $\alpha \in (0,1)$.

Proof. Consider a stationary linear Markov equilibrium with coefficients (α, β, δ). Straightforward integration shows that the consumer's equilibrium payoff is quadratic, and thus the system of equations (A.10) holds.

Suppose that $\alpha = 0$. From (7), $\lambda = 0$, and so $M_t = \mu$ for all $t \geq 0$; but this implies that prices do not respond to changes in the score, and hence, it is optimal for the consumer to behave myopically by choosing $Q_t = \theta_t - \mu$, a contradiction. If instead $\alpha < 0$, the last equation in (A.10) yields

$$\phi - \beta\lambda = (r + \kappa)\left(\frac{1}{\alpha} - 1\right) + \frac{\phi}{\alpha} < 0,$$

which is a contradiction with the equilibrium being stationary ((ii) in Proposition 1).

The case $\alpha = 1$ can be easily ruled out too: since $\lambda > 0$ in this case, the last equation in the system (A.10) yields that $\beta = 0$, but the second equation then implies that $\alpha = 0$, a contradiction. As a corollary, $\beta \neq 0$ in a stationary linear Markov equilibrium.

Suppose now that $\alpha > 1$. The last two equations of (A.10) can be used to solve for β and thus to find an expression for λ as a function of ϕ and α and the parameters r and κ. In addition, from the last equation in (A.10),

$$L := \phi - \beta\lambda = \frac{\phi - \alpha(\kappa + r) + \kappa + r}{\alpha},$$

and hence, we can solve for $\phi = \phi(\alpha, L)$. We conclude that in the equation for λ, (7), ϕ can be replaced by expressions that depend on L and α. Specifically, the resulting equation is

$$\frac{\alpha L \sigma_\theta^2}{\kappa(\kappa + L)\sigma_\xi^2 + \alpha^2\sigma_\theta^2} + \frac{(\alpha - 1)(\kappa + L + r)(3\alpha(\kappa + L + r) - 3\kappa + L - r)}{\alpha(2\alpha(\kappa + L + r) - 2\kappa - r)} = 0.$$

By stationarity, $L > 0$. But this coupled with $\alpha > 1$ implies that the left-hand side of this expression is strictly positive, which is a contradiction. Thus, $\alpha \in (0, 1)$. □

We continue with the proof of the proposition. From the proof of the previous lemma, $\beta \neq 0$. In the system (A.10), we can multiply the second equation by $\alpha \neq 0$ and the third by $\beta \neq 0$ to obtain $(r + 2\phi)\alpha(\alpha + 2\beta) = (r + \kappa + \phi)\beta(\alpha - 1)$. From here, $\beta = B(\phi, \alpha)$ where

$$B(\phi, x) := \frac{-x^2(r + 2\phi)}{2(r + 2\phi)x - (r + \kappa + \phi)(x - 1)} \in (-\frac{x}{2}, 0) \text{ when } x \in (0, 1). \tag{A.11}$$

Moreover, since $\alpha \in (0, 1)$ and $\phi - \beta\lambda > 0$, it follows from (7) that $\lambda > 0$. However, when $\alpha > 0$ and $\beta < 0$, the unique strictly positive root of (7) is given by

$$\Lambda(\phi, \alpha, \beta) := \frac{\sigma_\theta^2\alpha(\alpha + \beta) + \kappa\sigma_\xi^2(\kappa + \phi) - \sqrt{[\sigma_\theta^2\alpha(\alpha + \beta) + \kappa\sigma_\xi^2(\kappa + \phi)]^2 - 4\kappa(\sigma_\theta\sigma_\xi)^2\alpha\beta\phi}}{2\beta\kappa\sigma_\xi^2}. \tag{A.12}$$

In particular, since $\alpha^2 + \alpha B(\phi, \alpha) = \alpha[\alpha + B(\phi, \alpha)] \geq \alpha^2/2 > 0$ when $\alpha \in [0, 1]$, $\sigma_\theta^2\alpha(\alpha +$

$B(\phi,\alpha)) + \kappa\sigma_\xi^2(\kappa+\phi) > 0$ over the same range.

We conclude that $\lambda = \Lambda(\phi, \alpha, B(\phi,\alpha))$ in equilibrium, and so, using the last equation of (A.10), we arrive to equation (13): namely, $\alpha \in (0,1)$ must satisfy $A(\phi,\alpha) = 0$, where

$$A(\phi, x) := (r+\kappa+\phi)(x-1) - \lambda(\phi, x, B(\phi,x))xB(\phi,x), \; x \in [0,1]. \tag{A.13}$$

Lemma 4. *For every $\phi > 0$, there exists a unique $\alpha \in (0,1)$ satisfying the previous equation. Moreover, the resulting function $\alpha : (0,\infty) \to (0,1)$ is of class of class C^2.*

Proof: Fix $\phi > 0$. Observe that as $x \to 1$, $B(\phi,x) \to -1/2$ and $\lim_{x\to 1} \lambda(\phi, x, B(\phi,x)) > 0$. Hence, $\lim_{x\to 1} A(\phi,x) > 0$. Similarly, as $x \to 0$, $B(\phi,x) \to 0$ and $B(\phi,x)\lambda(\phi,x,B(\phi,x)) \to 0$. Hence, $\lim_{x\to 0} A(\phi,x) < 0$. The existence of $\alpha \in (0,1)$ satisfying $A(\phi,\alpha) = 0$ follows from the continuity of $x \in [0,1] \mapsto g(\phi,x)$ and the Intermediate Value Theorem.

To show uniqueness, we prove that $x \mapsto -\Lambda(\phi, x, B(\phi,x))xB(\phi,x)$ is strictly increasing in $[0,1]$. To this end, notice first that since

$$H(\phi,x) := -\Lambda(\phi, x, B(\phi,x))B(\phi,x) > 0, \; x \text{ in } (0,1),$$

it suffices to show that $x \mapsto H(\phi,x)$ is strictly increasing in the same region.

From the previous limits, $\lim_{x\to 0} H(\phi,x) = 0$ and $\lim_{x\to 1} H(\phi,x) > 0$; thus, there must exist a point at which $H_x(\phi,x) > 0$. Towards a contradiction, suppose that there is $\hat{x} \in (0,1)$ s.t. $H_x(\phi,\hat{x}) = 0$, where H_x denotes the partial derivative of H with respect to x, and let $\ell(\phi,x) := \sigma_\theta^2 x(x + B(\phi,x)) + \kappa\sigma_\xi^2(\kappa+\phi)$. At any such $\hat{x}$,

$$\ell_x(\phi,\hat{x}) \underbrace{\left[\ell(\phi,\hat{x}) - \left(\ell^2(\phi,\hat{x}) - 4\kappa\sigma_\xi^2\sigma_\theta^2 B(\phi,\hat{x})\hat{x}\phi\right)^{1/2}\right]}_{<0, \text{ as } B<0} = 2\kappa(\sigma_\theta\sigma_\xi)^2[B_x(\phi,\hat{x})\hat{x} + B(\phi,\hat{x})]\phi. \tag{A.14}$$

Moreover, straightforward algebra shows that

$$B_x(\phi,x)x = \underbrace{B(\phi,x)}_{<0} - \underbrace{\frac{x^2(r+2\phi)(r+\kappa+\phi)}{[2x(r+2\phi) - (r+\kappa+\phi)(x-1)]^2}}_{>0} < 0 \;\text{ for }\; x \in [0,1],$$

so $B_x(\phi,x)x + B(\phi,x) < 0$ for all $x \in [0,1]$. It then follows that $\ell_x(\phi,\hat{x}) = \sigma_\theta^2[2x + B_x(\phi,\hat{x})\hat{x} + B(\phi,\hat{x})] > 0$, otherwise the left-hand side of the previous condition is positive, while the right-hand side is negative.

Isolating the square root and squaring both sides in the first-order condition leads to the cancellation of $\ell^2\ell_x^2$ in (A.14). Dividing the resulting expression by $4\kappa(\sigma_\theta\sigma_\xi)^2\phi$ then yields

$$
\begin{aligned}
0 &= \ell_x(\phi,\hat{x})\underbrace{\{\ell(\phi,\hat{x})[-B_x(\phi,\hat{x})\hat{x}-B(\phi,\hat{x})]+\ell_x(\phi,\hat{x})B(\phi,\hat{x})\hat{x}\}}_{K:=} \\
&\quad+\underbrace{\kappa(\sigma_\theta\sigma_\xi)^2[-B_x(\phi,\hat{x})\hat{x}-B(\phi,\hat{x})]^2\phi}_{>0}.
\end{aligned}
$$

But since $\ell_x(\phi,\hat{x})>0$, we must have that $K<0$. In particular, using that $\ell(\phi,x)=\sigma_\theta^2 x[x+B(\phi,x)]+\kappa\sigma_\xi^2(\phi+\kappa)$ and $\kappa\sigma_\xi^2(\phi+\kappa)[-B_x(\phi,\hat{x})\hat{x}-B(\phi,\hat{x})]>0$, it must be that

$$
\begin{aligned}
\sigma_\theta^2\{[\hat{x}^2+\hat{x}B(\phi,\hat{x})][-B_x(\phi,\hat{x})\hat{x}-B(\phi,\hat{x})]+[2\hat{x}+\hat{x}B_x(\phi,\hat{x})+B(\phi,\hat{x})]\hat{x}B(\phi,\hat{x})\} &< 0 \\
\Leftrightarrow \hat{x}^2[-\hat{x}B_x(\phi,\hat{x})+B(\phi,\hat{x})] &< 0.
\end{aligned}
$$

But from the expression for $B_x(\phi,x)x$, we have that $-xB_x(\phi,x)+B(\phi,x)=x^2(r+2\phi)(r+\kappa+\phi)/[2x(r+2\phi)-(r+\kappa+\phi)(x-1)]^2\geq 0$, reaching a contradiction. The continuity of H_x implies that $x\mapsto H(\phi,x)$ is strictly increasing.

To conclude, since $(0,1)\times(0,\infty)\mapsto A(x,\phi)$ is of class C^1 and $\partial A/\partial x>0$, the Implicit Function Theorem guarantees that $\phi\in(0,1)\mapsto\alpha(\phi)\in(0,1)$, where $\alpha(\phi)$ satisfies $A(\phi,\alpha(\phi))=0$, $\phi>0$, is of class C^1. In fact, since such $\alpha(\cdot)$ is unique, the local property of continuous differentiability on the implicit function delivered by the theorem trivially extends to the whole domain. Moreover, using $H(\phi,x)$ defined above, it follows that

$$
\alpha'(\phi)=\frac{1-\alpha(\phi)-\alpha(\phi)H_\phi(\phi,\alpha(\phi))}{r+\kappa+\phi+H(\phi,\alpha(\phi))+\alpha(\phi)H_\alpha(\phi,\alpha(\phi))}.
$$

It is straightforward to verify that the right-hand side of the previous equality is of class C^1 as a function of $\phi\in(0,\infty)$. Thus, $\alpha(\cdot)$ is of class C^2. □

To characterize δ, recall that the first equation in (A.10) reads

$$
(r+\phi)\frac{2\delta\mu}{\lambda}=-(\alpha+\beta)\delta\mu+\kappa\mu\frac{\alpha-1}{\lambda}+[\lambda\delta\mu+\phi(\mu-\lambda\bar{Y})]\frac{\alpha+2\beta}{\lambda},
$$

where $\bar{Y}=\mu[\delta+\alpha+\beta]/\phi$. Plugging this expression in the previous equation yields

$$
\left[\frac{2(r+\phi)}{\lambda}+\alpha+\beta\right]\delta\mu \;=\; \mu\left[\frac{\kappa(\alpha-1)}{\lambda}+\frac{\alpha+2\beta}{\lambda}[\phi-(\alpha+\beta)\lambda]\right].
$$

Observe that since $\alpha+\beta>0$, the bracket on the left-hand side is strictly positive. If $\mu=0$ this equation is trivially satisfied, i.e., the price and quantity demanded along the path of play have no deterministic intercept (and $v_2=0$, leaving the rest of the system unaffected).

If $\mu \neq 0$, we have that $\delta = D(\phi, \alpha)$ where

$$D(\phi, x) := \frac{\kappa(\alpha - 1) + [\alpha + 2B(\phi, \alpha)][\phi - (\alpha + B(\phi, \alpha))\Lambda(\phi, \alpha, B(\phi, \alpha))]}{2(r + \phi) + (\alpha + B(\phi, \alpha))\Lambda(\phi, \alpha, B(\phi, \alpha))}, \qquad \text{(A.15)}$$

for $(\phi, x) \in (0, \infty) \times (0, 1)$, which is well-defined for all values $\phi > 0$.

To conclude the proof of the Proposition, two steps are remaining:

1. Determination of the remaining coefficients. From the three matching coefficient conditions (A.7), v_2, v_3 and v_5 are determined using δ, α and β as follows:

$$v_2 = \frac{2\delta\mu}{\lambda}, \; v_3 = \frac{\alpha + 2\beta}{2\lambda} > 0, \quad \text{and} \quad v_5 = \frac{\alpha - 1}{\lambda} < 0.$$

As for v_1 and v_4 (corresponding to θ and θ^2 in the value function) these can be obtained by differentiating the HJB equation with respect to θ. Specifically,

$$(r + \kappa)[v_1 + 2v_4\theta + v_5 M] \;=\; (\delta\mu + \alpha\theta + \beta M)\,[1 + v_5\lambda] - v_5\phi\left[M - \mu + \lambda\bar{Y}\right] - 2v_4\kappa(\theta - \mu)$$

leads to the additional equations

$$\begin{aligned} 2(r + \kappa)v_4 &= \alpha \underbrace{[1 + \lambda v_5]}_{=\alpha;\text{ system } (A.7)} -2v_4\kappa \Rightarrow v_4 = \frac{\alpha^2}{2(r + 2\kappa)}, \quad \text{and,} \\ (r + \kappa)v_1 &= \delta\mu\alpha + v_5\phi(\mu - \lambda\bar{Y}) \Rightarrow v_1 = \frac{\delta\mu\alpha}{r + \kappa} + \frac{\phi(\mu - \lambda\bar{Y})\alpha\beta}{(r + \kappa + \phi)(r + \kappa)}. \end{aligned}$$

The coefficient v_0 can be found by equating the constant terms in the HJB equation - since the value function is quadratic, there is no constraint on this coefficient.

2. Transversality conditions and admissibility of the candidate equilibrium strategy (6). Refer to the Online Appendix.

This concludes the proof. □

Proofs for Section 4.4

Proof of Proposition 2. Consider the partial differential equation (PDE)

$$\begin{aligned} -(\alpha + \beta)[\delta\mu + \alpha\theta + \beta M] + \mathcal{L}F(\theta, M) - (r + \phi)F(\theta, M) &= 0 \\ \lim_{r \to \infty} e^{-rt}\mathbb{E}_0[F(\theta_t^{\theta^o}, M_t^{m^o})] &= 0, \end{aligned}$$

where $\mathcal{L}F := -\kappa(\theta-\mu)F_\theta + [-\phi(M-\mu+\lambda\bar{Y}) + \lambda(\delta\mu+\alpha\theta+\beta M)]F_M + \frac{\sigma_\theta^2}{2}F_{\theta\theta} + \frac{(\lambda\sigma_\xi)^2}{2}F_{MM}$ and $(\theta_t^{\theta^o}, M_t^{m^o}))_{t\geq 0}$ is the type-belief process starting from $(\theta_0, M_0) = (\theta^o, m^o) \in \mathbb{R}^2$.

From the proof of Proposition 1, the previous equation admits as solution the function

$$V_M(\theta, M) = v_2 + 2v_3 M + v_5\theta$$

where v_2, v_3 and v_5 are the coefficients of the consumer's value function on M, M^2, and $M\theta$, respectively. In fact, display (A.8) shows that the previous function satisfies the PDE, while the transversality condition follows from $(\theta_t^{\theta^o}, M_t^{m^o}))_{t\geq 0}$ being stationary Gaussian and V_M being linear.

Importantly, $V_M(\cdot,\cdot)$ (i) is of class C^2 and (ii) exhibits quadratic growth. Thus, the Feynman-Kac Representation Theorem (Remark 3.5.6 in Pham 2009) applies: namely,

$$V_M(\theta_t, M_t) = -\mathbb{E}_t\left[\int_t^\infty e^{-(r+\phi)(s-t)}(\alpha+\beta)Q_s ds\right],\ \forall t\geq 0,$$

where we used that $Q_t = \delta\mu + \alpha\theta_t + \beta M_t$ in equilibrium. This concludes the proof. □

Proof of Proposition 3. (i) Limit values. Let $\ell(\phi,\alpha) := \alpha\sigma_\theta^2[\alpha + B(\phi,\alpha)] + \kappa\sigma_\xi^2(\phi+\kappa)$ and

$$J(\phi) := \sqrt{[\ell(\phi,\alpha(\phi))]^2 - 4\kappa(\sigma_\xi\sigma_\theta)^2 B(\phi,\alpha(\phi))\alpha(\phi)\phi} - \ell(\phi,\alpha(\phi)).$$

With this in hand, observe that (13) (or, equivalently, $A(\phi,\alpha(\phi)) = 0$, where $A(\phi,x)$ is defined in (A.13)), becomes $(r+\kappa+\phi)(\alpha(\phi)-1) + \alpha(\phi)J(\phi)/[2\kappa\sigma_\xi^2] = 0$.

Since $\alpha(\phi)\in(0,1)$ for all $\phi>0$, and $0<|B(\phi,\alpha)|<1/2$ for all $\alpha\in(0,1)$ and $\phi>0$, we have that $0 < -4\kappa(\sigma_\xi\sigma_\theta)^2\beta(\phi)\alpha(\phi)\phi \to 0$ as $\phi\to 0$. In addition, because $\alpha(\phi)+\beta(\phi)>0$, it follows that $\ell(\phi,\alpha) > \kappa^2\sigma_\xi^2$. Using that $\beta(\phi) = B(\phi,\alpha(\phi))$ this yields,

$$0 < J(\phi) = \frac{-4\kappa(\sigma_\xi\sigma_\theta)^2\beta(\phi)\alpha(\phi)\phi}{\sqrt{[\ell(\phi,\alpha(\phi))]^2 - 4\kappa(\sigma_\xi\sigma_\theta)^2\beta(\phi)\alpha(\phi)\phi} + \ell(\phi,\alpha(\phi))} < \frac{-4\kappa(\sigma_\xi\sigma_\theta)^2\beta(\phi)\alpha(\phi)\phi}{2\kappa^2\sigma_\xi^2}.$$

We conclude that $\lim_{\phi\to 0}\alpha(\phi)$ exists and takes value 1.

As for the limit to $+\infty$, notice that since $\ell(\phi,\alpha(\phi)) \geq \kappa\sigma_\xi^2\phi$ and $\alpha(\phi)B(\phi,\alpha(\phi)) < 0$,

$$0 < J(\phi) = \frac{-4\kappa(\sigma_\xi\sigma_\theta)^2 B(\phi,\alpha(\phi))\alpha(\phi)}{\sqrt{\left[\frac{\ell(\phi,\alpha(\phi))}{\phi}\right]^2 - \frac{4\kappa(\sigma_\xi\sigma_\theta)^2 B(\phi,\alpha(\phi))\alpha(\phi)}{\phi}} + \frac{\ell(\phi,\alpha(\phi))}{\phi}} \leq -\frac{4\kappa(\sigma_\xi\sigma_\theta)^2 B(\phi,\alpha(\phi))\alpha(\phi)}{2\sigma_\xi^2\kappa}.$$

But since $\alpha(\cdot)$ and $B(\cdot, \alpha(\cdot))$ are bounded, $L(\cdot)$ is bounded. Thus, from $g(\phi, \alpha(\phi)) = 0$,

$$1 - \alpha(\phi) = \underbrace{\frac{\alpha(\phi)J(\phi)}{2\kappa\sigma_\xi^2}}_{\text{bounded}} \underbrace{\frac{1}{(r+\kappa+\phi)}}_{\to 0 \text{ as } \phi\to\infty} \to 0 \text{ as } \phi \to \infty.$$

Regarding the limit values for $\beta(\phi) = B(\phi, \alpha(\phi))$, these follow from the limit behavior of $\alpha(\phi)$ and the definition of B given by (A.11). As for $\delta(\phi)$, recall that from (A.15),

$$\delta(\phi) = \frac{\kappa(\alpha(\phi) - 1) + [\alpha(\phi) + 2\beta(\phi)][\phi - (\alpha(\phi) + \beta(\phi))\lambda(\phi)]}{2(r+\phi) + (\alpha(\phi) + \beta(\phi))\lambda(\phi)}.$$

But from Lemma 5 below, $\lambda(\phi) \to 0$ as $\phi \to 0$. Using that $\alpha(\phi) \to 1$ and $\alpha(\phi) + 2\beta(\phi) \to 0$ as $\phi \to 0$, and that $\alpha(\phi) + \beta(\phi) > 0$, it follows that $\delta(\phi) \to 0$ as $\phi \to 0$. The same lemma shows that $\lambda(\phi) \to \sigma_\theta^2/\kappa\sigma_\xi^2$ as $\phi \to \infty$. Thus, $[\phi - (\alpha(\phi) + \beta(\phi))\lambda(\phi)]/[2(r+\phi) + (\alpha(\phi) + \beta(\phi))\lambda(\phi)] \to 1/2$ as $\phi \to \infty$. The limit $\delta(\phi) \to 0$ as $\phi \to \infty$ then follows from $\alpha(\phi) \to 1$ and $\alpha(\phi) + 2\beta(\phi) \to 0$ as $\phi \to \infty$.

It remains to show the limit result on prices. To this end, we start with a preliminary

Lemma 5. *$\lim_{\phi\to 0} \lambda(\phi) = 0$, $\lim_{\phi\to\infty} \lambda(\phi) = \sigma_\theta^2/\kappa\sigma_\xi^2$ and $\lim_{\phi\to 0} \lambda(\phi)/\phi = 2\sigma_\theta^2/[\sigma_\theta^2 + 2\sigma_\xi^2\kappa^2]$.*

Proof. $\lim_{\phi\to 0} \lambda(\phi) = 0$ follows directly from the first bound in (A.18) that we establish in the proof of the bounds section (ii) shortly. Also, letting $\ell(\phi, \alpha) := \sigma_\theta^2[\alpha + B(\phi, \alpha)] + \kappa\sigma_\xi^2[\phi + \kappa]$, it is straightforward to verify that

$$\lambda(\phi) = \frac{4\kappa(\sigma_\xi\sigma_\theta)^2\alpha(\phi)}{2\kappa\sigma_\xi^2\left(\sqrt{\left[\frac{\ell(\phi,\alpha(\phi))}{\phi}\right]^2 - \frac{4\kappa(\sigma_\xi\sigma_\theta)^2 B(\phi,\alpha(\phi))\alpha(\phi)}{\phi}} + \frac{\ell(\phi,\alpha(\phi))}{\phi}\right)} \to \frac{4\kappa(\sigma_\xi\sigma_\theta)^2}{2\kappa\sigma_\xi^2[\kappa\sigma_\xi^2 + \kappa\sigma_\xi^2]} = \frac{\sigma_\theta^2}{\kappa\sigma_\xi^2}$$

as $\phi \to \infty$, and thus the second limit holds. The third limit follows directly from the first equality in the previous display. This ends the proof of the lemma. □

Using the lemma, we first show that $\lim_{\phi\to\infty} \text{Var}[\lambda(\phi)Y_t] = \lim_{\phi\to 0} \text{Var}[\lambda(\phi)Y_t] = 0$. Recall that $(\delta(\phi), \alpha(\phi), \beta(\phi)) \to (0, 1, -1/2)$ as $\phi \to 0, \infty$. Also, from (A.4),

$$\text{Var}[Y_t] = \frac{1}{2(\phi - \beta(\phi)\lambda(\phi))}\left[\sigma_\xi^2 + \frac{\alpha(\phi)\sigma_\theta^2}{\kappa[\phi - \beta(\phi)\lambda(\phi) + \kappa]}\right]. \tag{A.16}$$

By the previous lemma, therefore, $\lim_{\phi\to\infty} \text{Var}[Y_t] = 0$, and so $\lim_{\phi\to\infty} \text{Var}[\lambda(\phi)Y_t] = 0$ as $\phi \to \infty$.

As for the other limit, this one follows from writing (A.16) as

$$\mathrm{Var}[\lambda(\phi)Y_t] \quad = \quad \underbrace{\frac{1}{2(\frac{\phi}{\lambda(\phi)} - \beta(\phi))}}_{\to \text{ constant}} \underbrace{\lambda(\phi)}_{\to 0} \underbrace{\left[\sigma_\xi^2 + \frac{\alpha(\phi)\sigma_\theta^2}{\kappa[\phi - \beta(\phi)\lambda(\phi) + \kappa]}\right]}_{\to \sigma_\xi^2 + \sigma_\theta^2/\kappa^2} \to 0 \text{ as } \ \phi \to 0.$$

The L^2-limits then follow directly from the following results: $(\delta(\phi), \alpha(\phi), \beta(\phi)) \to (0, 1, -1/2)$ as $\phi \to 0, \infty$; $P_t = \delta\mu + (\alpha+\beta)M_t$ and $M_t = \mu + \lambda[Y_t - \bar{Y}]$; $\mathbb{E}[P_t] = \mu[\alpha(\phi) + \beta(\phi) + \delta(\phi)] \to \mu/2$ as $\phi \to 0$ and $+\infty$; and the triangular inequality.

(ii) Bounds. Observe that the bounds for $\beta(\phi)$ were already determined from (A.11) and $\alpha(\phi) \in (0, 1)$. As for the lower bound for α, we will show the stronger result

$$\max\left\{\frac{r+\kappa+\phi}{r+\kappa+2\phi}, \frac{r+\kappa+\phi}{r+\kappa+\phi+\sigma_\theta^2/2\kappa\sigma_\xi^2}\right\} \le \alpha(\phi).$$

The bound is tight in the sense that it converges to 1 when $\phi \to 0, \infty$.

To obtain the bound, observe that from (A.12), $\lambda(\phi)$ satisfies

$$\lambda(\phi) = \frac{2\sigma_\theta^2\alpha(\phi)\phi}{\sqrt{\ell^2(\phi, \alpha(\phi)) - 4\kappa(\sigma_\theta\sigma_\xi)^2\alpha(\phi)B(\phi, \alpha(\phi))\phi} + \ell(\phi, \alpha(\phi))} < \frac{\sigma_\theta^2\alpha(\phi)\phi}{\ell(\phi, \alpha(\phi))}, \quad \text{(A.17)}$$

where $\ell(\phi, \alpha(\phi)) = \sigma_\theta^2\alpha(\phi)[\alpha(\phi) + B(\phi, \alpha(\phi))] + \kappa\sigma_\xi^2[\phi + \kappa]$. From here, we get two bounds:

$$\lambda < \frac{2\phi}{\alpha(\phi)} \ \text{ and } \ \lambda < \frac{\sigma_\theta^2}{\kappa\sigma_\xi^2}. \quad \text{(A.18)}$$

In fact, since $\ell(\phi, \alpha(\phi)) = \sigma_\theta^2\alpha(\phi)[\alpha(\phi) + B(\phi, \alpha(\phi))] + \kappa\sigma_\xi^2[\phi + \kappa]$, $\ell(\phi, \alpha(\phi)) > \sigma_\theta^2\alpha(\phi)^2/2$ due to $B(\phi, \alpha(\phi)) \ge -\alpha/2$; but using this in (A.17) leads to the first inequality in (A.18). Similarly, the second upper bound follows from (A.17) using that $\alpha(\phi) < 1$ and that $\ell(\phi, \alpha(\phi)) > \kappa\sigma_\xi^2\phi$ due to $\alpha + B(\phi, \alpha) > 0$. In particular, $\lambda(\phi)$ is bounded over $\mathbb{R}_+$.

Using the first bound in (A.18) in $A(\phi, \alpha(\phi)) = 0$ yields

$$0 = (r+\kappa+\phi)(\alpha(\phi) - 1) + \underbrace{\lambda(\phi)}_{\le 2\phi/\alpha(\phi)} \alpha(\phi) \underbrace{[-B(\phi, \alpha(\phi))]}_{\in(0, \alpha(\phi)/2)} < (r+\kappa+\phi)(\alpha(\phi) - 1) + \phi\alpha(\phi)$$

$$\Rightarrow \alpha(\phi) > \frac{r+\kappa+\phi}{r+\kappa+2\phi} > \frac{1}{2}, \ \text{ for all } \ \phi > 0.$$

Similarly, using the second bound, $0 < (r+\kappa+\phi)(\alpha(\phi) - 1) + [\sigma_\theta\alpha(\phi)]^2/[2\kappa\sigma_\xi^2]$. Using that $\alpha^2 < \alpha$ in the previous inequality delivers the second lower bound.

To establish the bounds for the expected price, we omit the dependence of $(\alpha, \beta, \delta, \lambda)$

on ϕ in what follows. Observe that $\mathbb{E}[P_t] = \delta\mu + (\alpha+\beta)\mathbb{E}[M_t] = [\delta+\alpha+\beta]\mu$. Now, adding the second and third equation in the system (A.10) yields $(\alpha+2\beta)(\alpha+\beta)\lambda = (r+2\phi)(\alpha+2\beta) + (r+\kappa+\phi)(\alpha-1) + (\alpha+\beta)^2\lambda$. Thus, in equilibrium,

$$\begin{aligned}\delta &= \frac{\kappa(\alpha-1)+[\alpha+2\beta][\phi-(\alpha+\beta)\lambda]}{2(r+\phi)+(\alpha+\beta)\lambda}\\ &= \frac{\kappa(\alpha-1)+(\alpha+2\beta)\phi-(r+2\phi)(\alpha+2\beta)-(r+\kappa+\phi)(\alpha-1)-(\alpha+\beta)^2\lambda}{2(r+\phi)+(\alpha+\beta)\lambda}\\ &= \frac{-(r+\phi)[2(\alpha+\beta)-1]-(\alpha+\beta)^2\lambda}{2(r+\phi)+(\alpha+\beta)\lambda},\end{aligned}$$

from where it is easy to conclude that

$$\mathbb{E}[P_t] \quad = \quad \mu[\alpha+\beta+\delta] = \mu\frac{r+\phi}{2(r+\phi)+(\alpha+\beta)\lambda}. \tag{A.19}$$

In particular, $\mathbb{E}[P_t] < \mu/2$ when $\mu \neq 0$ follows directly from $\lambda(\alpha+\beta) > 0$.

On the other hand, from (A.17) and $\ell(\phi,\alpha) > \sigma_\theta^2\alpha[\alpha+B(\phi,\alpha)] = \sigma_\theta^2\alpha[\alpha+\beta]$,

$$(\alpha+\beta)\lambda < (\alpha+\beta)\frac{\sigma_\theta^2\alpha\phi}{\ell(\phi,\alpha)} < (\alpha+\beta)\frac{\sigma_\theta^2\alpha(\phi)\phi}{\sigma_\theta^2\alpha[\alpha+\beta]} = \phi.$$

Using this bound in (A.19) and the fact that $r > 0$ leads to $\mathbb{E}[P_t] > \mu/3$ whenever $\mu \neq 0$,.

(iii) Quasiconvexity of α. To prove this property, it is more useful to solve the last two equations in the system (A.10) for λ and β, namely,

$$\lambda(\phi,\alpha) = -\frac{(\alpha-1)(\kappa+r+\phi)(\kappa+\alpha(-\kappa+r+3\phi)+r+\phi)}{\alpha^3(r+2\phi)}, \tag{A.20}$$

$$\beta(\phi,\alpha) = -\frac{\alpha^2(r+2\phi)}{\kappa+\alpha(-\kappa+r+3\phi)+r+\phi}. \tag{A.21}$$

Substituting both expressions into (7) that defines λ, and recalling $s := \sigma_\xi^2/\sigma_\theta^2$, we obtain an alternate locus $(\phi, \alpha(\phi))$ that satisfies

$$\tilde{A}(\phi,\alpha) := \frac{(\alpha-1)(\kappa+r+\phi)(\kappa+\alpha(-\kappa+r+3\phi)+r+\phi)}{\alpha^3(r+2\phi)} + \frac{\alpha(\kappa-\alpha(\kappa+r)+r+\phi)}{\alpha^3+\kappa s(\kappa-\alpha r+r+\phi)} = 0. \tag{A.22}$$

Observe that $\tilde{A}$ is increasing in α whenever $\tilde{g}(\phi,\alpha) = 0$. In fact, since Proposition 1 establishes the uniqueness of an equilibrium, there is a unique $\alpha(\phi) \in [0,1]$ solving $\tilde{A}(\phi,\alpha) = 0$. In addition, $\tilde{A}(\phi,1) = \phi/[\kappa s(\kappa+\phi)+1] > 0$. Thus $\tilde{A}(\phi,\cdot)$ must cross zero from below.

Now, the second partial derivative

$$\frac{\partial^2 \tilde{A}(\phi,\alpha)}{(\partial\phi)^2} = -\frac{2(\alpha-1)^2(2\kappa+r)^2}{(r+2\phi)^3} - \frac{2\alpha^5\kappa s\left(\alpha^2+\kappa^2 s\right)}{\left(\alpha^3+\kappa s(\kappa-\alpha r+r+\phi)\right)^3}$$

is strictly negative because, by inspection, the first term is nonpositive and the second term is strictly negative. Furthermore, we know from Proposition 1 that $\alpha(\phi)$ is twice continuously differentiable. Combined with the fact that $\tilde{A}$ is increasing in its second argument whenever $\tilde{A}=0$, the Implicit Function Theorem implies that $\alpha''(\phi)>0$ at any critical point $\alpha'(\phi)=0$.

(iv) Effect of noise terms σ_ξ/σ_θ. To show that $\phi\mapsto\alpha(\phi)$ is increasing in σ_ξ/σ_θ pointwise, consider the locus $\tilde{A}(\phi,\alpha)=0$ in (A.22) and differentiate with respect to $s:=\sigma_\xi^2/\sigma_\theta^2$. We obtain

$$\frac{\partial\tilde{A}}{\partial s} = -\frac{\alpha^4\kappa((1-\alpha)(\kappa+r)+\phi)(\kappa+(1-\alpha)r+\phi)}{\left(\alpha^3+\kappa s(\kappa+(1-\alpha)r+\phi)\right)^2}<0,$$

and because $\tilde{A}$ is increasing in α at $(\phi,\alpha(\phi))$, we conclude that α is increasing in s.

Finally, using the three equations (A.20)-(A.22), the derivative of the expected price $[\alpha+\beta+\delta]\mu$ with respect to α can be written as

$$\mu\frac{\alpha(r+\phi)(r+2\phi)(\kappa+r+\phi)(2(\kappa+r+\phi)-\alpha(2\kappa+r))}{\left[\alpha^2\left(\kappa^2+r(\kappa+2r)+5r\phi+3\phi^2\right)-\alpha(2\kappa+r)(\kappa+r+\phi)+(\kappa+r+\phi)^2\right]^2}>0.$$

Furthermore, when using (A.20)-(A.22), the expected price does not depend on s directly. Using that $\phi\mapsto\alpha(\phi)$ is increasing in σ_ξ/σ_θ, therefore, the expected price is also increasing in s. This concludes the proof. □

Proofs for Section 5

Proof of Proposition 4. Refer to the Online Appendix. □

The following results are used in the subsequent analysis, and their proofs can be found in the Online Appendix.

Lemma 6. *$\phi\mapsto G(\phi,\alpha,\beta)$ has a unique maximizer located at $\phi=\nu(\alpha,\beta)$, where $\nu(\alpha,\beta)$ is defined in (20). Moreover,*

(i) $\Lambda_\phi(\nu(\alpha,\beta),\alpha,\beta)=\Lambda(\nu(\alpha,\beta),\alpha,\beta)/(\nu+\kappa)$, where $\Lambda_\phi(\phi,\alpha,\beta)$ denotes the partial of $\phi\mapsto\Lambda(\phi,\alpha,\beta)$ with respect to ϕ, and,

(ii) $\Lambda(\nu(\alpha,\beta),\alpha,\beta)=\alpha\gamma(\alpha)/\sigma_\xi^2$*, where* $\gamma(\alpha)=\sigma_\xi^2[(\kappa^2+\alpha^2\sigma_\theta^2/\sigma_\xi^2)^{1/2}-\kappa]/\alpha^2$ *is the posterior belief's stationary variance when the histories* ξ^t*,* $t\geq 0$*, are public.*

Lemma 7. $[\alpha+\beta]'(\phi)<0$ *for* $\phi\in[\kappa,\arg\min\alpha]$*; hence, at any non-concealing point (i.e., satisfying (21)). Moreover, if* $r>\kappa$*,* $[\alpha+\beta]'(\phi)<0$ *for all* $\phi\in[0,\arg\min\alpha]$.

Proof of Proposition 5. We first show that $\alpha'(\phi)<0$ at any ϕ satisfying (21), i.e., $\phi=\nu(\alpha(\phi),\beta(\phi))$. To this end, recall that $\alpha(\phi)$ is the only value in $(0,1)$ satisfying $(r+\kappa+\phi)(\alpha(\phi)-1)+\alpha(\phi)H(\phi,\alpha(\phi))=0$, where

$$H(\phi,\alpha):=-\Lambda(\phi,\alpha,B(\phi,\alpha))B(\phi,\alpha)=\frac{\sqrt{\ell^2(\phi,\alpha)-4\kappa(\sigma_\xi\sigma_\theta)^2\alpha B(\phi,\alpha)\phi}-\ell(\phi,\alpha)}{2\kappa\sigma_\xi^2}$$

and $\ell(\phi,\alpha)=\sigma_\theta^2\alpha[\alpha+B(\phi,\alpha)]+\kappa\sigma_\xi^2[\phi+\kappa]$. Also, recall from the proof of Lemma 4 in the proof of proposition 1 that $\alpha\mapsto H(\phi,\alpha)$ is strictly increasing over $[0,1]$.

Thus, denoting the partial derivatives with subindices,

$$\alpha'(\phi)\left[r+\kappa+\phi+H(\phi,\alpha(\phi))+\alpha(\phi)H_\alpha(\phi,\alpha(\phi))\right]=1-\alpha(\phi)-\alpha(\phi)H_\phi(\phi,\alpha(\phi))$$

Consequently, because $H>0$, we conclude that the sign of α' is always determined by the sign of the right-hand side of the previous expression. We now show that the latter side is negative at any point ϕ s.t. $\phi=\nu(\alpha(\phi),\beta(\phi))=\kappa+\alpha(\phi)\gamma(\alpha(\phi))[\alpha(\phi)+\beta(\phi)]/\sigma_\xi^2$.

To simplify notation, let $\Delta(\phi,\alpha):=\sqrt{\ell^2(\phi,\alpha)-4\kappa(\sigma_\xi\sigma_\theta)^2\alpha B(\phi,\alpha)\phi}$. Omitting the dependence $(\phi,\alpha(\phi))$ of H, Δ, ℓ, B, and their respective partial derivatives,

$$H_\phi=\frac{1}{2\kappa\sigma_\xi^2}\left[\frac{\ell\ell_\phi-2\kappa(\sigma_\xi\sigma_\theta)^2\alpha[\phi B_\phi+B]}{\Delta}-\ell_\phi\right].$$

Using that $\ell_\phi=\sigma_\theta^2\alpha B_\phi+\kappa\sigma_\xi^2$ we can write

$$H_\phi\;=\;\frac{\kappa\sigma_\xi^2[\ell-\Delta]-2\kappa(\sigma_\xi\sigma_\theta)^2\alpha B}{2\kappa\sigma_\xi^2\Delta}+\frac{\sigma_\theta^2\alpha B_\phi[\ell-\Delta]-2\kappa(\sigma_\xi\sigma_\theta)^2\alpha\phi B_\phi}{2\kappa\sigma_\xi^2\Delta}$$

But observe that

$$\frac{\kappa\sigma_\xi^2[\ell-\Delta]-2\kappa(\sigma_\xi\sigma_\theta)^2\alpha B}{2\kappa\sigma_\xi^2\Delta}=-B\frac{\partial\Lambda}{\partial\phi}(\phi,\alpha,B).$$

From Lemma 6, however, $\Lambda_\phi(\nu(\alpha,\beta),\alpha,\beta)=\Lambda(\nu(\alpha,\beta),\alpha,\beta)/[\nu(\alpha,\beta)+\kappa]$ so the previous equality holds at any ϕ such that $(\phi,\alpha,\beta)=(\nu(\alpha(\phi),\beta(\phi)),\alpha(\phi),\beta(\phi))$.

On the other hand, the second term can be written as

$$\frac{\sigma_\theta^2 B_\phi}{\Delta}\left[\alpha\frac{\ell-\Delta}{2\kappa\sigma_\xi^2}-\phi\alpha\right]=\frac{\sigma_\theta^2 B_\phi}{\Delta}\left[(r+\kappa+\phi)(\alpha-1)-\phi\alpha\right]$$

where we used that $\alpha H = \alpha(\Delta-\ell)/2\kappa\sigma_\xi^2$. We deduce that, at the point of interest,

$$1-\alpha-\alpha H_\phi \;\;=\;\; \underbrace{1-\alpha+\frac{\lambda\alpha\beta}{\phi+\kappa}}_{K_1:=}-\underbrace{\frac{\sigma_\theta^2\alpha B_\phi}{\Delta}\left[(r+\kappa+\phi)(\alpha-1)-\phi\alpha\right]}_{K_2:=}$$

Straightforward differentiation shows that

$$B_\phi=\frac{\partial}{\partial\phi}\left(\frac{-\alpha^2(r+2\phi)}{2(r+2\phi)\alpha-(r+\kappa+\phi)(\alpha-1)}\right)=\frac{\alpha^2(\alpha-1)(r+2\kappa)}{[2(r+2\phi)\alpha-(r+\kappa+\phi)(\alpha-1)]^2}<0,$$

so $K_2>0$. As for the other term, $(r+\kappa+\phi)(\alpha-1)-\lambda\alpha\beta=0$ yields

$$K_1=\frac{(\phi+\kappa)(1-\alpha)+\lambda\alpha\beta}{\phi+\kappa}=\frac{r(\alpha-1)}{\phi+\kappa}<0.$$

We conclude that $\alpha'(\phi)<0$ at any point satisfying (21), provided such a point exists.

For existence, let $\eta(\phi):=\phi-\nu(\alpha(\phi),\beta(\phi))$, where $\nu(\alpha,\beta)=\kappa+\alpha\gamma(\alpha)[\alpha+\beta]/\sigma_\xi^2$, and

$$\gamma(\alpha):=\frac{\sigma_\xi^2}{\alpha^2}\left[\sqrt{\kappa^2+\alpha^2\frac{\sigma_\theta^2}{\sigma_\xi^2}}-\kappa\right],$$

(i.e., the unique positive solution of $0=\sigma_\theta^2-2\kappa\gamma-(\alpha\gamma/\sigma_\xi)^2$). Since $\alpha\in(1/2,1)$, γ is bounded, and so $\eta(\phi)>0$ for ϕ large. Also $L(0)$. The existence of ϕ s.t. $\eta(\phi)=0$ follows from the continuity of $L(\cdot)$.

Now, observe that since $\alpha>0$ and $\alpha+\beta>\alpha/2>0$, $\nu(\alpha(\phi),\beta(\phi))>\kappa$. On the other hand, since $\beta<0$ and $\alpha<1$,

$$\nu(\alpha(\phi),\beta(\phi))<\kappa+\frac{\alpha(\phi)^2\gamma(\alpha(\phi))}{\sigma_\xi^2}=\sqrt{\kappa^2+\alpha(\phi)^2\frac{\sigma_\theta^2}{\sigma_\xi^2}}<\sqrt{\kappa^2+\frac{\sigma_\theta^2}{\sigma_\xi^2}},$$

thus establishing (ii).

We conclude by proving the uniqueness of a non-concealing point. Notice that, by the previous lower bound and α's quasiconvexity, any point satisfying $\eta(\phi)=0$ must lie $[\kappa,\arg\min\alpha]$. We now show that $[\nu(\alpha(\phi),\beta(\phi))]'<0$ over that interval—however, this implies uniqueness. In fact, because the identity function is increasing, the existence of two such points would imply the existence of an intermediate third point at which $L(\cdot)$ vanishes and $[\nu(\alpha(\phi), \beta(\phi))]' > 0$, yielding a contradiction.

Observe that

$$[\nu(\alpha(\phi), \beta(\phi))]' = \frac{d}{d\phi}\left(\frac{\alpha(\phi)\gamma(\alpha(\phi))}{\sigma_\xi^2}\right)(\alpha(\phi) + \beta(\phi)) + \left(\frac{\alpha(\phi)\gamma(\alpha(\phi))}{\sigma_\xi^2}\right)\frac{d(\alpha(\phi) + \beta(\phi))}{d\phi}.$$

From Lemma 7, $\alpha(\phi) + \beta(\phi)$ is strictly decreasing over $[\kappa, \arg\min \alpha]$. Since $\alpha + \beta > 0$ and $\alpha\gamma(\alpha(\phi)) > 0$ is suffices to show that $[\alpha(\phi)\gamma(\alpha(\phi))]' < 0$ over the same region. However,

$$\frac{\alpha\gamma(\alpha)}{\sigma_\xi^2} = \frac{1}{\alpha}\left[\sqrt{\kappa^2 + \alpha^2\frac{\sigma_\theta^2}{\sigma_\xi^2}} - \kappa\right] = \frac{\sigma_\theta^2}{\sigma_\xi^2}\left(\sqrt{\frac{\kappa^2}{\alpha} + \frac{\sigma_\theta^2}{\sigma_\xi^2}} + \frac{\kappa}{\alpha}\right)^{-1}$$

which is strictly increasing in α. We conclude by using that $\alpha' < 0$ over $[\kappa, \arg\min \alpha]$. □

Proof of Proposition 6. Recall that $G(\phi) = \alpha(\phi)\lambda(\phi)/[\phi + \kappa - \beta(\phi)\lambda(\phi)] \geq 0$ for all $\phi \geq 0$. Since $\lambda(\phi)$ is bounded (second bound in A.18), $\lim_{\phi\to\infty} G(\phi) = 0$. Also $G(0) = 0$. By continuity, therefore, G has a global optimum that is interior.

From the definition of $\nu(\alpha, \beta)$,

$$G(\phi) := G(\phi, \alpha(\phi), \beta(\phi)) \leq G(\nu(\alpha(\phi), \beta(\phi)), \alpha(\phi), \beta(\phi)),$$

with equality only at ϕ^* defined in Proposition 5. Also, from Lemma 6, $\Lambda(\nu(\alpha, \beta), \alpha, \beta) = \alpha\gamma(\alpha)/\sigma_\xi^2$. Thus, letting $\nu(\phi) := \nu(\alpha(\phi), \beta(\phi))$,

$$G(\nu(\phi), \alpha(\phi), \beta(\phi)) = \frac{\alpha(\phi)\Lambda(\nu(\phi), \alpha(\phi), \beta(\phi))}{\nu(\phi) + \kappa - \beta(\phi)\Lambda(\nu(\phi), \alpha(\phi), \beta(\phi))} = \frac{\alpha^2(\phi)\gamma(\alpha(\phi))}{\alpha^2(\phi)\gamma(\alpha(\phi)) + 2\kappa\sigma_\xi^2}, \quad \text{(A.23)}$$

where we used that $\nu(\phi) = \kappa + \alpha(\phi)\gamma(\alpha(\phi))[\alpha(\phi) + \beta(\phi)]/\sigma_\xi^2$.

However, from Lemma 6, $\alpha^2\gamma(\alpha) = \sigma_\xi^2[(\kappa^2 + \alpha^2(\phi)\sigma_\theta^2/\sigma_\xi^2)^{1/2} - \kappa]$, and so (A.23) implies that $G(\nu(\phi), \alpha(\phi), \beta(\phi))$ is decreasing when $\alpha(\phi)$ is decreasing. Since $G(\phi)$ is bounded from above by a decreasing function of ϕ on $[\phi^*, \arg\min \alpha]$, $G(\phi^*) > G(\phi)$ over the same domain.

To conclude the proof, we show that $G(\phi) :=$ is decreasing when $\alpha(\phi)$ is increasing, i.e., over $(\arg\min \alpha(\phi), \infty)$. Using that $G(\phi, \alpha, \beta) := \alpha\Lambda(\phi, \alpha, \beta)/[\phi + \kappa - \beta\Lambda(\phi, \alpha, \beta)]$, and equations (A.20) and (A.21) to substitute for λ and β, we obtain the following expression for G in terms of α and ϕ only:

$$\tilde{G}(\alpha, \phi) := (1 - \alpha)\frac{(\kappa + r + \phi)}{(\kappa + (1 - \alpha)r + \phi)}\frac{(\kappa + \alpha(-\kappa + r + 3\phi) + r + \phi)}{\alpha(r + 2\phi)}. \quad \text{(A.24)}$$

Using that $\alpha \in (0,1)$, it is easy to verify that the two fractions on the right-hand side are strictly positive and strictly decreasing in ϕ. Thus, $\frac{\partial \tilde{G}}{\partial \phi} < 0$. Also, up to a positive term,

$$\frac{\partial \tilde{G}}{\partial \alpha} = \alpha^2 \left(-\left(-\kappa^2 + r^2 + 2\phi(\kappa + r) + 2\kappa r + 3\phi^2\right)\right) + 2\alpha r(\kappa + r + \phi) - (\kappa + r + \phi)^2.$$

Observe that the right-hand side of the previous expression is a concave quadratic in α when $\phi > \kappa$. Moreover, over the same region, straightforward maximization yields that such quadratic is negative for all parameter values. Also, by Proposition 5 and quasiconvexity of α, $\kappa < \arg\min \alpha$. Thus, $\frac{d\tilde{G}}{d\phi} = \frac{\partial \tilde{G}}{\partial \alpha}\alpha'(\phi) + \frac{\partial \tilde{G}}{\partial \phi} < 0$, $\phi \in (\arg\min \alpha, \infty)$. □

Proofs for Section 6

Proof of Proposition 7. Observe that $\text{Var}[P_t] = [\alpha(\phi) + \beta(\phi)]^2 \text{Var}[\theta_t] G(\phi)$. From the proof of Proposition 6, $\lim_{\phi \to 0, \infty} G(\phi) = 0$. Also, $\alpha + \beta$ is bounded. By continuity, we conclude that $\text{Var}[P_t]$ has a global optimum that is interior.

From Proposition 6, however, $G(\phi)$ is maximized to the left of ϕ^*. Also, from Lemma 7, $\alpha'(\phi) + \beta'(\phi) < 0$ over $[\kappa, \arg\min \alpha]$. Since $\kappa < \phi^* < \arg\min \alpha$, $\text{Var}[P_t]$ cannot attain a maximum in $[\phi^*, \arg\min \alpha]$. One can then verify that the total derivative of $\text{Var}[P_t]$ with respect to ϕ is negative over $[\arg\min \alpha, +\infty)$ for all parameter values $(r, \kappa, \sigma_\theta, \sigma_\xi)$.

As for part (ii), this one follows from the second part of Lemma 7, i.e., $\alpha + \beta$ being strictly decreasing in $[0, \arg\min \alpha]$ when $r \geq \kappa$, and the same previous argument over $[\arg\min \alpha, +\infty)$. □

In order to prove Proposition 8, we make use of the following auxiliary lemma which is proved in the Online Appendix.

Lemma 8. *$\lim_{\phi \to 0} \alpha'(\phi)$ exists and is negative; $\lim_{\phi \to 0} \beta'(\phi)$ exists; $\lim_{\phi \to 0} [\alpha + \beta + \delta]'(\phi)$ exists and is negative; and $\lim_{\phi \to 0} G'(\phi)$ exists and is positive.*

Proof of Proposition 8. We start with (i) for profits. From (A.19), $\alpha + \beta + \delta > 1/3$. Thus, omitting the dependence of the equilibrium coefficients on ϕ,

$$\Pi(\phi) := \mu^2 [\alpha + \beta + \delta]^2 + \frac{\sigma_\theta^2}{2\kappa}(\alpha + \beta)^2 G(\phi) \geq \frac{\mu^2}{9} + \frac{\sigma_\theta^2 (\alpha + \beta)^2}{2\kappa} G(\phi), \quad \text{for all} \quad \phi > 0.$$

On the other hand, from the proof of Proposition 7, $\lim_{\phi \to 0, \infty} (\alpha + \beta)^2 G(\phi) = 0$, and so

$\lim_{\phi \to 0,\infty} \Pi(\phi) = \mu^2/4$. Thus, if

$$\frac{\mu^2}{9} + \frac{\sigma_\theta^2(\alpha+\beta)^2}{2\kappa} G(\phi) \geq \frac{\mu^2}{4} \Leftrightarrow \mu^2 \leq \frac{18\sigma_\theta^2}{5\kappa}(\alpha+\beta)^2 G(\phi),$$

it follows that $\Pi(\phi) > 1/4$. Since $\phi \mapsto [\alpha(\phi) + \beta(\phi)]^2 G(\phi)$ is continuous, strictly positive, and converges to 0 as $\phi \to 0, \infty$, it has a global maximum; denote it by $\phi^\dagger$. Letting

$$\underline{\mu}_f = \left[\frac{18\sigma_\theta^2}{5\kappa}(\alpha(\phi^\dagger) + \beta(\phi^\dagger))^2 G(\phi^\dagger)\right]^{1/2} > 0$$

the result follows.

As for the consumer, let $CS_\mu(\phi)$ denote consumer surplus given $\mu \geq 0$ and observe that

$$CS_\mu(\phi) = CS_0(\phi) + \mu^2 R(\phi).$$

where $R(\phi) := [\alpha(\phi) + \beta(\phi) + \delta(\phi)] \left(1 - \frac{3}{2}[\alpha(\phi) + \beta(\phi) + \delta(\phi)]\right)$ and

$$CS_0(\phi) = \frac{\sigma_\theta^2}{2\kappa} G(\phi) L(\phi) + \frac{\sigma_\theta^2}{2\kappa}\left[\alpha(\phi) - \frac{(\alpha(\phi))^2}{2}\right].$$

Importantly since $\alpha(\phi) + \beta(\phi) + \delta(\phi) \to 1/2$ as $\phi \to 0$ and $+\infty$, we have that $\lim_{\phi \to 0} R(\phi) = \lim_{\phi \to +\infty} R(\phi) = 1/8$. In addition we know that $1/3 < \alpha(\phi) + \beta(\phi) + \delta(\phi) < 1/2$ for all $\phi > 0$. Because the function $x \mapsto x - 3x^2/2$ is strictly decreasing in $[1/3, 1/2]$, we have that $R(\phi) > 1/8$, for all $\phi > 0$.

Fix any $\hat{\phi} > 0$. Then, using that $CS_\mu(0) = \mu^2/8$,

$$CS_\mu(\hat{\phi}) - CS_\mu(0) = \mu^2\left[R(\hat{\phi}) - \frac{1}{8}\right] + \frac{\sigma_\theta^2}{2\kappa} \underbrace{\left[G(\hat{\phi})L(\hat{\phi}) + \alpha(\hat{\phi}) - \frac{(\alpha(\hat{\phi}))^2}{2} - \frac{1}{2}\right]}_{=:K(\hat{\phi})}.$$

Observe that $K(\cdot)$ and $R(\cdot)$ are independent of μ, we can choose μ arbitrarily large such that the right-hand side is strictly positive. Since $CS_\mu(0) = \lim_{\phi \to \infty} CS_\mu(0)$, the consumer's global maximum ϕ^c must be interior.

We now turn to (ii), starting with the firms' case. Towards a contradiction, suppose that there are sequences $\mu_n \nearrow \infty$ and $\phi_n > 0$ for all $n \in \mathbb{N}$ such that $\Pi_{\mu_n}(\phi_n) \geq \Pi_{\mu_n}(0) = \Pi_{\mu_n}(+\infty)$. Then,

$$\Pi_{\mu_n}(\phi_n) \geq \Pi_{\mu_n}(0) \Leftrightarrow \frac{\sigma_\theta^2}{2\kappa}[\alpha(\phi_n) + \beta(\phi_n)]^2 G(\phi_n) \geq \mu_n^2 \left[\frac{1}{4} - [\alpha(\phi_n) + \beta(\phi_n) + \delta(\phi_n)]^2\right]$$

Observe first that $(\phi_n)_{n\in\mathbb{N}}$ cannot have a cluster point different from zero. Otherwise, along such subsequence, say $(\phi_{n_k})_{k\in\mathbb{N}}$, both $[\alpha(\phi_{n_k})+\beta(\phi_{n_k})]^2G(\phi_{n_k})$ and $1/4-[\alpha(\phi_{n_k})+\beta(\phi_{n_k})+\delta(\phi_{n_k})]^2$ converge to strictly positive numbers; the inequality is then violated for large k.

Suppose now that there is a subsequence $(\phi_{n_k})_{k\in\mathbb{N}}$ that diverges. Using that $\alpha+\beta+\delta=(r+\phi)/[2(r+\phi)+\lambda(\alpha+\beta)]$ and that $G=\alpha\lambda/[\phi+\kappa-\beta\lambda]$ we have that

$$\Pi_{\mu_{n_k}}(\phi_{n_k})\geq\Pi_{\mu_n}(0)\Leftrightarrow\frac{\sigma_\theta^2}{2\kappa}\frac{4(\alpha+\beta)^2\alpha\lambda[2(r+\phi)+\lambda(\alpha+\beta)]^2}{[4(r+\phi)(\alpha+\beta)\lambda+(\alpha+\beta)^2\lambda^2](\phi+\kappa-\beta\lambda)}\Big|_{\phi=\phi_{n_k}}\geq\mu_{n_k}.$$

But since α,β and λ are all bounded and $(\alpha,\beta,\lambda)\to(1,-1/2,\sigma_\theta^2/[\kappa\sigma_\xi^2])$ as $\phi\to+\infty$, both the numerator and denominator are $O(\phi^2)$ for large ϕ, so the limit of the left-hand side of the second inequality exists. The inequality is then violated for large k, a contradiction.

From the previous argument, the only remaining possibility is that $(\phi_n)_{n\in\mathbb{N}}$ converges to zero. However, from Proposition 3, $\lim\limits_{\phi\to0}(\alpha(\phi),\beta(\phi),\lambda(\phi))=(1,-1/2,0)$, and so

$$\frac{4(\alpha+\beta)^2\alpha\lambda[2(r+\phi)+\lambda(\alpha+\beta)]^2}{[4(r+\phi)(\alpha+\beta)\lambda+(\alpha+\beta)^2\lambda^2](\phi+\kappa-\beta\lambda)}=\frac{4(\alpha+\beta)^2\alpha[2(r+\phi)+\lambda(\alpha+\beta)]^2}{[4(r+\phi)(\alpha+\beta)+(\alpha+\beta)^2\lambda](\phi+\kappa-\beta\lambda)}\to\frac{2r}{\kappa},\quad\text{as}\quad\phi\to\infty,$$

and so the same inequality is again violated, a contradiction.

The case for the consumer is proved in an analogous fashion. Namely, towards a contradiction, assume that there are $(\mu_n)_{n\in\mathbb{N}}$ decreasing towards zero and $(\phi_n)_{n\in\mathbb{N}}$ strictly positive such that $CS_{\mu_n}(\phi_n)\geq CS_{\mu_n}(0)$. Straightforward algebraic manipulation shows that

$$CS_{\mu_n}(\phi_n)\geq CS_{\mu_n}(0)\Leftrightarrow\frac{1}{\frac{\text{Var}[\theta_t]}{2}\left(\frac{(\alpha(\phi_n)-1)^2}{R(\phi_n)-1/8}-2L(\phi_n)\frac{G(\phi_n)}{R(\phi_n)-1/8}\right)}\geq\frac{1}{\mu_n},$$

with $R(\phi)$ defined in part (i) of the proof. As in the firms' case, there can't be a subsequence of $(\phi_n)_{n\in\mathbb{N}}$ converging to a value different from zero; otherwise, the left-hand side of the inequality on the right converges, but the right-hand side converges. In the Online Appendix we show that

$$\lim_{\phi\to0,+\infty}\frac{[\alpha(\phi)-1]^2}{R(\phi)}=0,\quad\text{and}\quad\lim_{\phi\to0,+\infty}\frac{G(\phi)}{R(\phi)-1/8}>0.$$

But since $\lim\limits_{\phi\to0,+\infty}L(\phi)=-1/8$, the left-hand side of the same inequality is again violated, a contradiction. Thus, there must exist $\underline{\mu}^c>0$ such that for all $\mu<\underline{\mu}^c$, $CS_\mu(\phi)\geq CS_\mu(\phi)$ for all $\phi\in(0,\infty)$. This concludes the proof. □

Proof of Proposition 9. Let $\phi_{\underline{\alpha}}(r):=\arg\min\{\alpha(\phi;r):\phi\geq0\}$, where the dependence of α on $r>0$ is explicit. Similarly, we write $\phi^*(r)$ whenever needed, and notice that, by

Proposition 5, $\phi^*(r) \in [\kappa, \sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2}]$ for all $r > 0$.

Let $H(\phi) := \lambda(\phi, \alpha(\phi), \beta(\phi))[\alpha(\phi) + \beta(\phi)]$, and recall that from (A.19), $\mathrm{P}(\phi) := \mathbb{E}[P^\phi] = \mu(r+\phi)/[2(r+\phi) + H(\phi)]$. Thus, we need to study $H(\phi)/(r+\phi)$. However, using that, in equilibrium, $(r+\kappa+\phi)(\alpha-1) - \lambda\beta\alpha = 0$ and $\beta = -\alpha^2(r+2\phi)/[2(r+2\phi)\alpha - (r+\kappa+\phi)(\alpha-1)]$,

$$
\begin{aligned}
\lambda(\alpha+\beta) &= \underbrace{\frac{(r+\kappa+\phi)(\alpha-1)}{\alpha\beta}}_{=\lambda}(\alpha+\beta) = \frac{(1-\alpha)(r+\kappa+\phi)[\alpha(r+2\phi) - (r+\kappa+\phi)(\alpha-1)]}{\alpha^2(r+2\phi)} \\
\Rightarrow \frac{H(\phi)}{r+\phi} &= \frac{(1-\alpha)(r+\kappa+\phi)}{\alpha(r+\phi)} \times \frac{\alpha(r+2\phi) - (r+\kappa+\phi)(\alpha-1)}{\alpha(r+2\phi)} =: h_1(\phi)h_2(\phi).
\end{aligned}
$$

To prove (i), observe that since $h_i > 0$, it suffices to show that $h_i' < 0$ at $\phi > \phi_{\underline{\alpha}}(r)$ for all $r > 0$, $i = 1, 2$. To this end, it is easy to see that

$$
\begin{aligned}
h_1' < 0 \;&\Leftrightarrow\; [-\alpha'[r+\kappa+\phi] + 1 - \alpha]\alpha[r+\phi] - (1-\alpha)[r+\kappa+\phi][\alpha'[r+\phi] + \alpha] < 0 \\
&\Leftrightarrow\; \underbrace{-\alpha'[r+\kappa+\phi]\alpha[r+\phi] - (1-\alpha)[r+\kappa+\phi]\alpha'[r+\phi]}_{-[r+\kappa+\phi]\alpha'[r+\phi]} - \kappa\alpha(1-\alpha) < 0, \quad \text{(A.25)}
\end{aligned}
$$

which clearly holds when $\alpha' > 0$. On the other hand,

$$
\begin{aligned}
h_2'(\phi) < 0 \;&\Leftrightarrow\; [\{\alpha[r+2\phi]\}' + 1 - \alpha - \alpha'[r+\kappa+\phi]]\alpha[r+2\phi] \\
&\quad -\{\alpha[r+2\phi]\}'[\alpha(r+2\phi) - (r+\kappa+\phi)(\alpha-1)] < 0 \\
&\Leftrightarrow\; [1-\alpha-\alpha'[r+\kappa+\phi]]\alpha[r+2\phi] - (r+\kappa+\phi)(1-\alpha)[\alpha'[r+2\phi] + 2\alpha] < 0 \\
&\Leftrightarrow\; \underbrace{(1-\alpha)\alpha[r+2\phi] - (r+\kappa+\phi)(1-\alpha)2\alpha}_{=(1-\alpha)\alpha[-r-2\kappa]<0} - (r+\kappa+\phi)\alpha'[r+2\phi] < 0, \quad \text{(A.26)}
\end{aligned}
$$

which is also true when $\alpha' > 0$.

In order to prove (ii), we show that for small enough discount rates, the function

$$
R(\phi) = \frac{\nu(\phi) - \kappa}{\nu(\phi) + r},
$$

with $\nu(\phi) := \kappa + \alpha(\phi)\gamma(\alpha(\phi))[\alpha(\phi) + \beta(\phi)]/\sigma_\xi^2$, is a decreasing upper bound to $H(\phi)/(r+\phi)$ over $[\phi^*(r), \phi_{\underline{\alpha}}(r)]$ that coincides with it at $\phi^*(r)$; as a byproduct, $1/[2 + R(\phi)]$ is increasing lower bound for the expected price that guarantees the latter cannot attain its global minimum in that region. The proof relies on the following technical lemma, which is proved in the Online Appendix

Lemma 9. $(\alpha(\phi), \beta(\phi)) \in I := [1/2, 1] \times [-1/2, -1/8]$ *for all* $\phi > \kappa$. *Moreover, there exists*

$0 < \underline{r} < \kappa$ such that for all $r < \underline{r}$

$$\phi \mapsto \frac{\lambda(\phi, \alpha, \beta)}{r + \phi}$$

is decreasing over $[\kappa, \infty)$ for all $(\alpha, \beta) \in I$.

Consider now $r < \underline{r}$ as in the lemma, and observe that $\underline{r}$ is independent of the specific values of $(\alpha, \beta) \in I$ (but it depends on the latter set, of course). For notational simplicity, we omit the dependence of all variables on r. Proposition 5 shows that $\nu(\phi)$ crosses the identity (only once) from above. Thus, we have that $\phi \geq \nu(\phi)$ for all $\phi \in [\phi^*, \phi_{\underline{\alpha}}]$ with equality only at ϕ^*. Since $\phi^* > \kappa$, the lemma yields

$$\frac{H(\phi)}{r+\phi} := \frac{\lambda(\phi, \alpha(\phi), \beta(\phi))[\alpha(\phi) + \beta(\phi)]}{r + \phi} \leq \frac{\lambda(\nu(\phi), \alpha(\phi), \beta(\phi))[\alpha(\phi) + \beta(\phi)]}{r + \nu(\phi)}, \tag{A.27}$$

for all $\phi \in [\phi^*, \phi_{\underline{\alpha}}]$, with equality only at ϕ^*. However, from Lemma 6, $\lambda(\nu(\alpha, \beta), \alpha, \beta) = \alpha\gamma(\alpha)/\sigma_\xi^2$, where $\gamma(\alpha) := \sigma_\xi^2[(\kappa^2 + \sigma_\theta^2/\sigma_\xi^2)^{1/2} - \kappa]/\alpha^2$. Thus,

$$\frac{H(\phi)}{r+\phi} \leq \frac{1}{\nu(\phi) + r}\left[\frac{\alpha(\phi)\gamma(\alpha(\phi))[\alpha(\phi) + \beta(\phi)]}{\sigma_\xi^2}\right] = \frac{\nu(\phi) - \kappa}{\nu(\phi) + r} = R(\phi)$$

with equality only at ϕ^*. However, $R'(\phi) = \nu'(\phi)(r + \kappa)/[\nu(\phi) + r]^2$, and $\nu'(\phi) < 0$ over $[\kappa, \phi_{\underline{\alpha}}]$ was already established as part of the uniqueness step in the proof of Proposition 5.

Finally, to prove (iii), we show that $H(\phi)/(r + \phi)$ is increasing at $\phi^*(r) < \phi_{\underline{\alpha}}(r)$ for r large enough. From the last expressions in (A.25)–(A.26) in part (i) of the proof, it suffices to show that

$$\begin{aligned} \Xi_1(\phi) &:= -[r + \kappa + \phi]\alpha'[r + \phi] - \kappa\alpha(1 - \alpha) > 0 \text{ and} \\ \Xi_2(\phi) &:= (1 - \alpha)\alpha[-r - 2\kappa] - (r + \kappa + \phi)\alpha'[r + 2\phi] > 0 \end{aligned}$$

at $\phi^*(r)$ for large r. We start with Ξ_1.

From the equation that defines α, i.e., (13), $\alpha(\phi; r) \nearrow 1$ as $r \to \infty$ for all $\phi > 0$. Since $\phi^* = \phi^*(r) \in K := [\kappa, \sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2}]$ and $\{\alpha(\cdot; r) : K \to [0, 1] \mid r > 0\}$ is a family of continuous functions, it follows that $\{\alpha(\cdot; r) : r > 0\}$ converges uniformly to 1 over K (Dini's Theorem). Thus, $\alpha(\phi^*(r); r) \to 1$ as $r \to \infty$ as well. Omitting the dependence on r, therefore, it is direct to verify that:

$$\begin{aligned}
\lim_{r\to\infty} \beta(\phi^*) &= -1/2 \\
\lim_{r\to\infty} \phi^* &= \kappa + \frac{\gamma(1)}{2\sigma_\xi^2} > \kappa, \\
\lim_{r\to\infty} \ell(\phi^*, \alpha(\phi^*)) &= \frac{\sigma_\theta^2}{2} + \kappa\sigma_\xi^2[\lim_{r\to\infty} \phi^* + \kappa] > 0 \\
\lim_{r\to\infty} -\lambda(\phi^*)\beta(\phi^*) &= \lim_{r\to\infty} \frac{\sqrt{\ell^2(\phi,\alpha) - 4\kappa(\sigma_\xi\sigma_\theta)^2\alpha B(\phi,\alpha)\phi} - \ell(\phi,\alpha)}{2\kappa\sigma_\xi^2}\Big|_{(\phi,\alpha)=(\phi^*,\alpha(\phi^*))} > 0 \\
\lim_{r\to\infty} r(1-\alpha(\phi^*)) &= \lim_{r\to\infty} -\lambda(\phi^*)\beta(\phi^*)\alpha(\phi^*) > 0,
\end{aligned}$$

where the last equality follows from the equation that defines α. Moreover, from the proof of Proposition 5, letting $H = H(\phi,\alpha) := -\Lambda(\phi,\alpha,B(\phi,\alpha)B(\phi,\alpha)$,

$$\alpha'(\phi^*) = \left[\frac{1}{r+\kappa+\phi+H+\alpha H_\alpha} \times \left(\frac{r(\alpha-1)}{\phi^*+\kappa} - \underbrace{\frac{\sigma_\theta^2\alpha B_\phi}{\Delta(\phi,\alpha)}[(r+\kappa+\phi)(\alpha-1)-\phi\alpha]}_{K(\phi):=}\right)\right]_{\phi=\phi^*}$$

where $\Delta(\phi,\alpha) := \sqrt{\ell^2(\phi,\alpha) - 4\kappa(\sigma_\xi\sigma_\theta)^2\alpha B(\phi,\alpha)\phi}$ and

$$B_\phi = \frac{\partial}{\partial\phi}\left(\frac{-\alpha^2(r+2\phi)}{2(r+2\phi)\alpha - (r+\kappa+\phi)(\alpha-1)}\right) = \frac{\alpha^2(\alpha-1)(r+2\kappa)}{[2(r+2\phi)\alpha-(r+\kappa+\phi)(\alpha-1)]^2}.$$

It is easy to see that $\lim_{r\to\infty} \Delta(\phi^*,\alpha(\phi^*)) > 0$ and $\lim_{r\to\infty} B_\phi(\phi^*,\alpha(\phi^*)) = 0$. Thus, $\lim_{r\to\infty} K(\phi^*) = 0$. Also, $H_\alpha > 0$ (proof of Proposition 1) and $\lim_{r\to\infty} H_\alpha(\phi^*,\alpha(\phi^*))$ exists and is finite. This yields

$$\lim_{r\to\infty}(r+\kappa+\phi^*)\alpha'(\phi^*) = \lim_{r\to\infty}(r+\phi^*)\alpha'(\phi^*) = \lim_{r\to\infty}\frac{r(\alpha(\phi^*)-1)}{\phi^*+\kappa} < 0,$$

from where $\Xi_1(\phi^*) > 0$ for large $r > 0$ (the first term goes to $+\infty$, while the second vanishes).

Regarding $\Xi_2 := (1-\alpha)\alpha[-r-2\kappa] - (r+\kappa+\phi)\alpha'[r+2\phi]$ observe that when $\phi > \kappa$,

$$\Xi_2 > (1-\alpha)\alpha[-r-\kappa-\phi] - (r+\kappa+\phi)\alpha'[r+2\phi] = -[r+\kappa+\phi][\alpha(1-\alpha)+\alpha'(r+2\phi)].$$

We conclude that $\Xi_2(\phi^*(r)) > 0$ for large r, as

$$\lim_{r\to\infty}(r+2\phi^*)\alpha'(\phi^*) = \lim_{r\to\infty}\frac{r(\alpha(\phi^*)-1)}{\phi^*+\kappa} < 0 \quad\text{and}\quad \lim_{r\to\infty} -\alpha(\phi^*)[1-\alpha(\phi^*)] = 0.$$

This ends the proof of the proposition. □

Appendix B: Hidden Scores

We briefly elaborate on some key arguments used in the proofs of the observable-scores model that have a direct analog in the hidden case. Some of them are not used in the economic analysis of Section 7, but we provide them for completeness.

Existence and uniqueness of stationary linear Markov equilibria. An key step in the proof of Theorem 1 is Lemma 4, which establishes the existence and uniqueness of a solution to (13). As we show in the Online Appendix, replacing $B(\phi, \alpha) \in (-\alpha/2, 0)$ as defined in (A.11) by $-\alpha/2$ in (13) still leads to $\alpha \mapsto H^h(\phi, \alpha) := -\Lambda(\phi, \alpha, -\alpha/2)[-\alpha/2]$ being strictly increasing in $(0, 1)$, which was the key property used to show existence and uniqueness in the observable counterpart.

α^h is decreasing at any non-concealing point Since $\alpha \mapsto H^h(\phi, \alpha)$ as defined above is strictly increasing in $(0, 1)$, and $\alpha^h(\phi)$ solves $(r + \kappa + \phi)(\alpha^h - 1) + \alpha H^h(\phi, \alpha^h) = 0$, we have that at any point $\phi^{*,h}$ solves $\phi = \nu(\alpha(\phi), -\alpha(\phi)/2)$,

$$\begin{aligned} & \text{sign}([\alpha^h]'(\phi^{*,h})) \\ = \; & \text{sign}([1 - \alpha^h(\phi) - \alpha^h(\phi) H_\phi(\phi, \alpha^h(\phi))]|_{\phi=\phi^{*,h}}) = \text{sign}\left(\left[1 - \alpha^h + \frac{\lambda \alpha^h[-\alpha^h/2]}{\phi + \kappa} \right] \Big|_{\phi=\phi^{*,h}} \right) < 0. \end{aligned}$$

Existence and uniqueness of a non-concealing point. Since $\alpha^h + \beta^h = \alpha^h/2$ in the hidden case, $\nu(\phi) = \kappa + \frac{[\alpha^h(\phi)]^2 \gamma(\alpha^h(\phi))}{2\sigma_\xi^2}$. Thus, the existence of $\phi^{*,h}$ and the corresponding bounds follow directly. Moreover,

$$\nu(\phi) = \kappa + \frac{1}{2} \left[\sqrt{\kappa^2 + \alpha^2 \frac{\sigma_\theta^2}{\sigma_\xi^2}} - \kappa \right],$$

so $\text{sign}(\nu'(\phi^{*,h})) = \text{sign}(\alpha'(\phi^{*,h}))$. Since $[\alpha^h]'(\phi^{*,h}) < 0$ at any such $\phi^{*,h} > 0$, we conclude that there is only one such point.

Quasiconvexity of α^h and $G^h(\phi) := G(\phi, \alpha^h(\phi), -\alpha^h(\phi)/2)$ maximized to the left of $\phi^{*,h}$. They follow identical arguments as the ones used in the observable case.

References

Acquisti, A., C. Taylor, and L. Wagman (2016): "The Economics of Privacy," *Journal of Economic Literature*, 54(2), 442–92.

Acquisti, A., and H. R. Varian (2005): "Conditioning prices on purchase history," *Marketing Science*, 24(3), 367–381.

Bergemann, D., B. Brooks, and S. Morris (2015): "The Limits of Price Discrimination," *American Economic Review*, 105(3), 921–957.

Calzolari, G., and A. Pavan (2006): "On the optimality of privacy in sequential contracting," *Journal of Economic Theory*, 130(1), 168–204.

Cisternas, G. (2017a): "Career Concerns and the Nature of Skills," *American Economic Journal: Microeconomics*, forthcoming.

——— (2017b): "Two-Sided Learning and the Ratchet Principle," *Review of Economic Studies*, forthcoming.

Council of Economic Advisers (2015): *Big Data and Differential Pricing*.

Cummings, R., K. Ligett, M. Pai, and A. Roth (2016): "The Strange Case of Privacy in Equilibrium Models," in *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 659–659. ACM.

Di Pei, H. (2016): "Reputation with Strategic Information Disclosure," Discussion paper, MIT.

Dixon, P., and R. Gellman (2014): "The Scoring of America: How Secret Consumer Scores Threaten Your Privacy and Your Future," Discussion paper, World Privacy Forum.

Dubé, J.-P., and S. Misra (2017): "Scalable price targeting," Discussion paper, National Bureau of Economic Research.

Dworczak, P. (2017): "Mechanism Design with Aftermarkets: Cutoff Mechanisms," Discussion paper, Stanford University.

Federal Trade Commission (2014): *Data Brokers: a Call for Transparency and Accountability*. https://www.ftc.gov/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014.

Fong, Y.-F., and J. Li (2016): "Information Revelation in Relational Contracts," *Review of Economic Studies*, 84(1), 277–299.

Freixas, X., R. Guesnerie, and J. Tirole (1985): "Planning under Incomplete Information and the Ratchet Effect," *Review of Economic Studies*, 52(2), 173–191.

Fudenberg, D., and J. M. Villas-Boas (2006): "Behavior-Based Price Discrimination and Customer Recognition," in *Handbook on Economics and Information Systems*, ed. by T. Hendershott. Elsevier.

——— (2015): "Price Discrimination in the Digital Economy," in *The Oxford Handbook of the Digital Economy*, ed. by M. Peitz, and J. Waldfogel. Oxford University Press.

Gerardi, D., and L. Maestri (2016): "Dynamic Contracting with Limited Commitment and the Ratchet Effect," Discussion paper, Collegio Carlo Alberto.

Halac, M. (2012): "Relational Contracts and the Value of Relationships," *American Economic Review*, 102(2), 750–779.

Hart, O., and J. Tirole (1988): "Contract Renegotiation and Coasian Dynamics," *Review of Economic Studies*, 55(4), 509–540.

Heinsalu, S. (2017): "Dynamic Noisy Signalling," *American Economic Journal: Microeconomics*, forthcoming.

Holmström, B. (1999): "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1), 169–182.

Hörner, J., and N. Lambert (2017): "Motivational Ratings," Discussion paper, Yale University.

Kehoe, P. J., B. J. Larsen, and E. Pastorino (2018): "Dynamic Competition in the Era of Big Data," Discussion paper, Stanford University.

Kovbasyuk, S., and G. Spagnolo (2016): "Memory and Markets," Discussion paper, Einaudi Institute for Economics and Finance.

Kuvalekar, A., and E. Lipnowski (2018): "Job Insecurity," Discussion paper, University of Chicago.

Laffont, J.-J., and J. Tirole (1988): "The Dynamics of Incentive Contracts," *Econometrica*, 56(5), 1153–1175.

McAdams, D. (2011): "Discounts For Qualified Buyers Only," Discussion paper, Fuqua School of Business.

Ostrovsky, M., and M. Schwarz (2016): "Reserve Prices in Internet Advertising Auctions: A Field Experiment," Discussion paper, Stanford University.

Pham, H. (2009): *Continuous-Time Stochastic Control and Optimization with Financial Applications*. Springer-Verlag, Berlin.

Platen, E., and N. Bruti-Liberati (2010): *Numerical Solution of Stochastic Differential Equations with Jumps in Finance*. Springer.

Rafieian, O., and H. Yoganarasimhan (2018): "Targeting and Privacy in Mobile Advertising," Discussion paper, University of Washington.

Schmalensee, R. (1981): "Output and welfare implications of monopolistic third-degree price discrimination," *American Economic Review*, 71(1), 242–247.

Schmitz, A. J. (2014): "Secret Consumer Scores and Segmentations: Separating Haves from Have-Nots," *Michigan State Law Review*, pp. 1411–1473.

Shen, Q., and J. Miguel Villas-Boas (2017): "Behavior-Based Advertising," *Management Science*, forthcoming.

Syrgkanis, V., D. Kempe, and E. Tardos (2017): "Information Asymmetries in Common-Value Auctions with Discrete Signals," Discussion paper, Cornell University.

Taylor, C. R. (2004): "Consumer privacy and the market for customer information," *Rand Journal of Economics*, 35(4), 631–650.

Xu, Z., and A. Dukes (2017): "Personalized Pricing with Superior Information on Consumer Preferences," Discussion paper, University of Southern California.