

Consumer Scores and Price Discrimination*

Alessandro Bonatti and Gonzalo Cisternas

August 29, 2019

Forthcoming in *The Review of Economic Studies*

Abstract

We study the implications of aggregating consumers' purchase histories into *scores* that proxy for unobserved willingness to pay. A long-lived consumer interacts with a sequence of firms. Each firm relies on the consumer's current score—a linear aggregate of noisy purchase signals—to learn about her preferences and to set prices. If the consumer is strategic, she reduces her demand to manipulate her score, which reduces the average equilibrium price. Firms in turn prefer scores that overweigh past signals relative to applying Bayes' rule with disaggregated data, as this mitigates the ratchet effect and maximizes the firms' ability to price discriminate. Consumers with high average willingness to pay benefit from data collection, because the gains from low average prices dominate the losses from price discrimination. Finally, hidden scores—those only observed by the firms—reduce demand sensitivity, increase average prices, and reduce consumer surplus, sometimes below the naive-consumer level.

Keywords: price discrimination; purchase histories; consumer scores; persistence; transparency; ratchet effect.

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 Alessandro Acquisti, Rossella Argenziano, Sarah Brayne, Dean Eckles, Guido Friebel, Negin Golrezaei, Renato Gomes, Marina Halac, Johannes Hörner, Jorge Lemus, Pino Lopomo, George Mailath, Hamid Nazerzadeh, Yuliy Sannikov, Dick Schmalensee, Tavneet Suri, Juuso Toikka, John Van Reenen, Glen Weyl, Alex White, Robert Zeithammer, Juanjuan Zhang, and participants at numerous seminars and conferences for many helpful discussions. Finally, we thank Tina Pavlovich for excellent research assistance.

1 Introduction

Consumer scores are metrics that use all available data about individual customers—age, ethnicity, gender, household income, zip code, and purchase histories—to quantify and predict their profitability, health risk, job security, or credit worthiness. The best-known example is the FICO credit score used by every lender in the US. Another prominent, less well-known example is the Customer Lifetime Value (CLV) that many firms assign to their own customers in order to personalize prices, products, advertising messages, and various perks.¹ In addition to being deployed internally, consumer scores are also traded. Several data brokers (e.g., Acxiom, Equifax, and Experian) collect information from a variety of sources, aggregate it into scores, and sell these scores to companies that in turn, use them to refine their market-segmentation strategies. Critically, the transmission of information through such scores creates a link between a consumer’s interaction with one firm and the terms of her future transactions with other firms and industries.²

With the exception of those used in credit markets, consumer scores are neither regulated nor available to consumers. As such, some consumers ignore these links across transactions, and even the most sophisticated consumers cannot perfectly forecast the impact of their actions. Awareness of the mechanisms at play is nonetheless increasing quickly over time, thanks to recent regulatory efforts aimed at improving the transparency of firms’ information (e.g., the European Union’s General Data Protection Regulation). With consumer awareness rising at a fast pace, it is essential to understand how technological and market forces affect incentives. In particular, if the final use of information impacts the distribution of surplus, the mechanisms by which consumer data are collected, aggregated and transmitted can affect the terms of the transactions in which the data are generated and, thus, the informational content of the data itself.³

In this paper, we study the welfare consequences of aggregating purchase histories into scores that are used for third-degree price discrimination.⁴ We examine how these conse-

¹The composition of CLV scores is a rich area of marketing research (Dwyer, 1989; Berger and Nasr, 1998). *The Wall Street Journal* (2018) reports that “every consumer has at least one [CLV score], more likely several.” NPR provides more information in its Planet Money podcast, <https://www.npr.org/sections/money/2018/11/07/665392227/your-lifetime-value-score>.

²For example, the Equifax Discretionary Spending Limit Index is a number 1–1000 that “helps marketers differentiate between two households that look the same in terms of income and demographics but likely have considerably different spending power.” Similarly, information about a consumer’s sporting goods purchases or eating habits can become part of a predictive score for a health insurer (*The Economist*, 2012).

³Many European consumers are already aware of market-segmentation practices. In a recent survey (*European Commission*, 2018), 62% of respondents knew of the personalized ranking of online offers (based on past behavior or contextual information), and 44% knew of instances of personalized pricing.

⁴Price discrimination is implicitly used in a number of consumer markets in the form of personalized coupons, discounts, and fees (Dubé and Misra, 2017). Another related prominent practice is *product steer-*

quences depend on the consumers’ degree of sophistication (do they know they are being scored?) and on the availability of the sellers’ information (can they check their score?).

Our approach embeds a continuous-time model of the *ratchet effect* into an information design framework, which enables us to examine how data aggregation and transparency impact a strategic consumer’s incentives. In our model, a long-lived consumer faces a different monopolist at every instant of time. The consumer’s preferences are quadratic in the quantity demanded and linear in her privately observed willingness to pay, which is captured by a stationary Gaussian process. At any instant of time, an (unmodeled) intermediary observes a signal of the consumer’s current purchase distorted by Brownian noise and updates a one-dimensional aggregate of past signals that we refer to as the *score process*. Only the current value of the score is revealed to each monopolist, who uses it to set prices.⁵

We contrast the cases of naive consumers who ignore the link between the current purchased quantity and future prices (Section 3); strategic consumers who understand how firms react to the score and directly observe their score at all times (Sections 4-6); and strategic consumers who cannot observe their scores (Section 7).

Overview of the Results (1.) *Price discrimination based on purchase histories unambiguously harms naive consumers but can benefit strategic consumers.* When consumers are naive, both consumer and total surplus are decreasing (and producer surplus is increasing) in the precision of the firms’ information.⁶ More strikingly, the aggregation of information into scores does not protect consumers at all: the precision of the information transmitted by the firm-optimal score is the same as if firms observed the entire (disaggregated) history of the consumer’s purchase signals (Proposition 1).

By contrast, a strategic consumer can manipulate future prices by reducing her quantity demanded and hence lower her score. To grasp why the consumer can benefit from the transmission of her score—even if firms ultimately use the information so-gained against her—consider the following two-period example. The consumer interacts with two firms sequentially. Firm 1 sets price P_1 using prior information only, while firm 2 privately observes a signal of the consumer’s first-period quantity before choosing price P_2 . Because the consumer recognizes the impact of her first period quantity choice on the second period price, she reduces firm 2’s signal by adopting a lower demand function than myopically optimal. Firm 1 anticipates the consumer’s manipulation incentives and lowers its price.

ing—changing the order of search results when different consumers search for the same products online to steer high-value consumers towards high-price products (Hannak, Soeller, Lazer, Mislove, and Wilson, 2014).

⁵The Wall Street Journal (2018) reports that “At some retailers, the higher the [CLV score], the less likely you are to receive bigger discounts [...] some stores hold back discounts from higher-value customers until they are at risk of losing them. Why waste a 25% offer when the person is going to buy anyway?”

⁶This is a well-known property of static models with linear demand (Robinson, 1933; Schmalensee, 1981).




Figure 1: Strategic Demand Reduction

Figure 1 illustrates the equilibrium outcome: the consumer buys a lower quantity than she would like ($Q' < Q$) but also pays a lower price ($P' < P$). It is intuitive that a small amount of demand reduction is beneficial in the first period: the consumer gives up the marginal unit of consumption but receives an infra-marginal discount. Countering the impact of lower prices today are the losses from lower consumption today and from tailored prices tomorrow. However, if the consumer's average willingness to pay is sufficiently high, so is the average quantity demanded in the first period (e.g., Q' in Figure 1 above). When firms reduce prices, these discounts are applied to a large number of units. In this case, consumers benefit from lower prices more than they are hurt by personalized pricing.

Having developed some intuition for the effects of strategic demand reduction, we turn to the role of information aggregation and transparency in shaping equilibrium outcomes.⁷

(2.) *Firms can manage the ratchet effect and limit the resulting information loss by using relatively persistent scores.* As we saw in Figure 1, the ratchet effect drives the average equilibrium price and quantity levels down relative to the naive case. Furthermore, higher consumer types (who expect to buy large quantities in the immediate future) have a stronger incentive to reduce their demand to drive prices downward, which reduces the sensitivity of the consumer's actions to her type, the informativeness of the purchase signals and, thus, a firm's ability to price discriminate based on the score.

However, firms can use the aggregation of purchase histories to induce consumers to reveal more information about their preferences. With exponential scores, the question of how to aggregate information reduces to how heavily to discount past quantity signals. The

⁷Relative to the two-period example, our stationary model has two advantages: it eliminates end-game effects that can artificially influence policy implications, and it is considerably more manageable than any finite-horizon version that allows for information aggregation and endogenous learning. Further, continuous time allows a tractable analysis of how the ratchet effect varies across information structures.

score’s discount rate is, in turn, inversely related to its persistence.

We show that a unique score reveals the same amount of information in equilibrium as when firms observe the full history of the signals. Because of the ratchet effect, however, this score does not maximize the firms’ learning in equilibrium: by overweighing past signals, a more persistent score correlates less with the consumer’s current type, thereby incentivizing the consumer to signal her preferences more aggressively. On the margin, the latter effect dominates: learning is optimized by a score that conceals some information about the consumer’s behavior in exchange for more precise purchase signals. If the underlying signal technology is sufficiently precise, such a score is also more profitable (Propositions 4 and 5).⁸

(3.) *Making scores available to consumers makes demand more price sensitive, reduces prices, and increases consumer welfare.* Strategic demand reduction can help a consumer induce lower prices (as in Figure 1) while limiting the amount of information transmitted by her purchases. In fact, if the underlying signal technology is sufficiently precise, a strategic consumer who has access to her score is then better off than a naive one (Proposition 6). For this mechanism to operate successfully, however, score transparency is essential. We make this point by examining the case in which the consumer is strategic but the score is hidden.

When scores are hidden, the firms’ beliefs are private, and prices acquire a signaling value. Specifically, because of the full-support noise present in the score, any observed price is consistent with equilibrium behavior: due to the score’s persistence, therefore, a high price today now tells the consumer that the firms’ beliefs are high and that prices will be high in the near future. Since the consumer then expects to purchase relatively few units, she is less inclined to manipulate her score by reducing her current quantity, everything else being equal. Thus, the consumer’s demand becomes *less price sensitive* relative to the observable case (Proposition 7), where the consumer can identify “abnormally” high prices.

In equilibrium, firms exploit this reduced sensitivity by making prices more responsive to the score. This exacerbates the ratchet effect, resulting in lower average quantities and higher prices (Proposition 8). Moreover, for each level of the score’s persistence, consumer surplus is lower than with observable scores, provided the underlying signal technology is precise enough—having access to their score is therefore beneficial to consumers beyond increasing awareness. Not only that, consumer surplus can even be lower with hidden scores than with naive consumers (Proposition 9). Consequently, these results can inform policy interventions: regulations promoting consumer awareness and score transparency have complementary roles, and one without the other may be detrimental to consumer welfare.

⁸Aggregating purchase histories into scores can improve equilibrium learning, but introducing noise in the original signals cannot: a marginal increase in the score’s persistence yields a second-order loss in information but a first-order gain in the quality of the available signals. Conversely, introducing noise has a first-order negative effect that trumps the associated increase in consumer responsiveness.

Applied Relevance Our model’s policy implications for real-life consumers hinge upon three hypotheses: (i) strategic consumers attempt to manipulate the prices they are offered; (ii) firms use behavior-based data to guide pricing; and (iii) scores compress such data into statistics that exhibit persistence.

We argue that a lack of information by consumers—not a lack of sophistication—is presently the main barrier to observing in practice the ratchet effects we uncover in our model. To substantiate this claim, Section 8.1 presents anecdotal evidence from both business-to-consumer (B2C) and business-to-business (B2B) markets. In particular, we discuss several tactics by which consumers obtain discounts online, including “shopping cart abandonment” for near-complete purchases, and various attempts to receive lower personalized prices from ride-sharing services such as Lyft or Uber. We then turn to the market for online display advertising—a B2B setting in which the sellers are website publishers who set dynamic, personalized prices for advertising space, and the buyers are advertisers looking to reach a targeted audience.⁹ The ratchet effect arises because sellers attempt to learn the true distribution of advertisers’ valuations from their past bids. We describe buyers’ bid-shading behavior, as well as sellers’ strategies to alleviate the ratchet effect.

In all three examples, consumers are aware of the mechanism that links current purchases and future prices, and as such, they are able to anticipate the impact of their behavior. Consistent with our hypothesis, they take costly actions (e.g. delayed purchases, suboptimal routes, and lower current advertising volumes) to misrepresent their true willingness to pay and obtain lower prices.

At the same time, these settings differ from our model because buyers engage repeatedly with the same seller. This discrepancy also correlates with the amount of information available: in the absence of explicit regulation, it is harder for consumers to learn about the underlying links between the terms of trade across different sellers and hence to behave strategically in seemingly unrelated transactions. A notable exception is the case of FICO credit scores, where the link between the borrower’s current behavior and the terms offered by *different* future lenders is regulated, transparent, and hence quite salient. In Section 8.2, we discuss what we can learn from credit scores about the role of transparency policies in other consumer markets. Finally, in Section 8.3, we describe settings where all three our hypotheses may be particularly relevant in the near future.

Related Literature This paper builds on the literature on behavior-based price discrimination (Villas-Boas, 1999; Taylor, 2004), the results of which are surveyed in Fudenberg and Villas-Boas (2006), Fudenberg and Villas-Boas (2015), and Acquisti, Taylor, and Wagman

⁹We describe this market in greater detail in Section S.1 of the Supplementary Appendix.

(2016). Closest to our work is the two-period model of [Taylor \(2004\)](#) with observable actions and stochastic types. [Taylor \(2004\)](#) finds that ratchet forces result in lower equilibrium prices when consumers are strategic, as in our Figure 1 above. Qualitatively, our results differ because the noisy signals in our model imply that the consumer’s actions affect the information available to the firms. Therefore, the score’s persistence and transparency levels affect the firms’ ability to learn in a non trivial way.¹⁰

Our score process is an instance of the linear Gaussian rating introduced by [Hörner and Lambert \(2019\)](#), who study information design in the [Holmström \(1999\)](#) career concerns model. Relative to their setting, we maintain the assumptions of short-lived firms and additive signals but introduce two additional features. First, our consumer is privately informed, which makes the informational content of the score endogenous. Second, we allow for an interaction in the consumer’s payoff function between her action and the firms’ beliefs, which implies that optimal actions depend on the level of the firms’ beliefs. This dependence makes the transparency question critical in our setting because the consumer’s incentives now depend on whether she knows her own score.¹¹

The force driving the dynamics of our model is the ratchet effect (e.g., [Freixas, Guesnerie, and Tirole, 1985](#), [Laffont and Tirole, 1988](#) and, more recently, [Gerardi and Maestri, 2016](#)), which has received experimental validation ([Charness, Kuhn, and Villeval, 2011](#); [Cardella and Depew, 2018](#)). The ratchet effect also underscores the analysis of privacy in settings with multiple principals. [Calzolari and Pavan \(2006\)](#) consider the case of two principals, and [Dworczak \(2017\)](#) that of a single transaction followed by an aftermarket.¹² Relative to all these papers, the presence of noise in our model and the restriction to linear pricing limit the ratchet effect and allow consumers to potentially benefit from information transmission.

The marketing literature ([Lewis, 2005](#)) has already suggested the idea that the dynamics of strategic behavior must be incorporated into consumer valuation methods such as CLV. Finally, the ratchet effect in online advertising markets is the subject of a growing literature. In particular, [Amin, Rostamizadeh, and Syed \(2013\)](#) study mechanisms for inferring a single

¹⁰[Cummings, Ligett, Pai, and Roth \(2016\)](#) and [Shen and Villas-Boas \(2017\)](#) study models where second-period advertising is targeted on the basis of the consumer’s first-period purchase: a similar trade-off arises whereby the value of targeted advertising to the consumer impacts the equilibrium price of the first-period good, and hence the amount of information revealed by the consumer. Similarly, [Ichihashi \(2018\)](#) studies a monopolist who uses the information gathered from consumers to tailor product recommendations as well as to price discriminate. Finally, [Kehoe, Larsen, and Pastorino \(2018\)](#) explore the consequences of dynamic price competition among firms that offer different product varieties.

¹¹The question of transparency in financial markets has received considerable attention. In particular, [Duffie, Dworczak, and Zhu \(2017\)](#) examine how benchmark transparency in over-the-counter markets helps traders to learn their alternative options with other dealers, in the same way that score transparency helps our consumer to learn the prices she will be charged by future firms.

¹²The ratchet effect appears, with a different interpretation or motivation, in relational contracts ([Halac, 2012](#); [Fong and Li, 2016](#)) and in dynamic games with symmetric uncertainty ([Cisternas, 2018](#)).

strategic buyer's value for a good when interacting repeatedly, and [Hummel \(2018\)](#) studies how to design dynamic reserve prices in the presence of the ratchet effect. We discuss this literature more extensively in Section S.1 in the Supplementary Appendix.

2 Model

We develop a continuous-time model with a long-lived consumer and a family of short-run firms. The model is motivated by a discrete-time setting with two key features. First, the consumer faces a different monopolist in every period. Second, within each period, the consumer and the current firm play sequential-move stage game: the monopolist initially posts a unit price for its product based on the current value of a consumer's score; having observed the price, the consumer then chooses which quantity to buy. It is instructive to begin with the *observable* (or transparent) case: the consumer can observe her score directly.

Players, types, and payoffs Consider an infinitely lived consumer who interacts with a continuum of firms in continuous time. The consumer has a discount rate $r > 0$ and, at any instant $t \geq 0$, consuming $Q_t = q$ units of a good at price $P_t = p$ results in a flow utility

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

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

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

where $(Z_t^\theta)_{t \geq 0}$ is a Brownian motion.¹³ In particular, $(\theta_t)_{t \geq 0}$ is Gaussian, and by stationarity,

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

Each firm interacts with the consumer for only one instant, and only one firm operates at any time t ; we refer to the monopolist operating at t simply as *firm t* . Production costs are normalized to zero, and hence firm t 's ex post profits are given by $P_t Q_t$, $t \geq 0$.

Score process and information At any $t \geq 0$, firm t observes only the current value Y_t of a *score process* $(Y_t)_{t \geq 0}$ that is provided by an (unmodeled) intermediary. By contrast, when

¹³Stationarity requires $\theta_0 \sim \mathcal{N}(\mu, \sigma_\theta^2/2\kappa)$ independent of $(Z_t^\theta)_{t \geq 0}$.

scores are *observable*, 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 type realizations.¹⁴ (In the hidden-scores case (Section 7), firm t observes only Y_t , and the score is not directly observed by the consumer.)

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

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

where $(Z_t^\xi)_{t \geq 0}$ is a Brownian motion independent of $(Z_t^\theta)_{t \geq 0}$, Q_t is the realized purchase by the consumer at $t \geq 0$, and $\sigma_\xi > 0$ is a volatility parameter.

The intermediary then aggregates every history of the form $\xi^t := (\xi_s : s < t)$ into a real number Y_t that corresponds to the consumer's time- t score, $t \geq 0$. Building on Hörner and Lambert (2019), we restrict attention to *exponential scores*, i.e., to Ito processes

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

where $\phi \in (0, \infty)$. Under this specification, the consumer's current score is a linear function of the contemporaneous history of recorded purchases, and lower values of ϕ lead to scores processes that exhibit more *persistence*, as past information is discounted less heavily in those cases.¹⁵ In differential form, the score process satisfies

$$dY_t = -\phi Y_t dt + d\xi_t = (Q_t - \phi Y_t) dt + \sigma_\xi dZ_t^\xi, \quad t > 0. \quad (5)$$

Finally, the prior is that (θ_0, Y_0) is normally distributed; the exact distribution is determined in equilibrium so that the joint process $(\theta_t, Y_t)_{t \geq 0}$ is stationary Gaussian along the path of play.¹⁶ In what follows, the expectation operator $\mathbb{E}[\cdot]$ is with respect to such prior, while $\mathbb{E}_0[\cdot]$ is conditioned on the realized value of (θ_0, Y_0) . The former is the relevant operator for studying welfare, while the latter is used in the equilibrium analysis. The conditional expectations of the consumer and firm t are denoted by $\mathbb{E}_t[\cdot]$ and $\mathbb{E}[\cdot|Y_t]$, respectively.

¹⁴We show in Section 5 that firms observing the whole histories Y^t , $t \geq 0$, is one instance of our model.

¹⁵This class of “aggregators” is natural starting point: since the model is Gaussian, Bayesian updating using the whole history of signals leads to a posterior mean that falls within this class.

¹⁶Stationarity requires that Y_0 carries information about θ_0 , i.e., that some data was collected before time zero, as it is only then that a stationary model can be initialized at $t = 0$ —writing the disaggregated recorded purchase histories as $(\xi_s : s < t)$, $t \geq 0$, makes this explicit. Refer to Section 4.1 and to Lemma A.1 in the Appendix for the characterization of a stationary outcome under the type of equilibrium studied.

Strategies and equilibrium concept A strategy for the consumer specifies, for each $t \geq 0$, a quantity $Q_t \in \mathbb{R}$ to purchase as a function of the history of prices, types, and score values, $(\theta_s, P_s, Y_s : 0 \leq s \leq t)$. Instead, firm t chooses a price $P_t \in \mathbb{R}$ that is conditioned on Y_t only, $t \geq 0$. A strategy for the consumer is *linear Markov* if $Q_t = Q(p, \theta_t, Y_t)$ for all $t \geq 0$, where $Q : \mathbb{R}^3 \rightarrow \mathbb{R}$ is linear and p is the current posted price (i.e., $Q(\cdot, \theta_t, Y_t)$ is the demand at the history (θ_t, Y_t)). Similarly for firm t , $P_t = P(Y_t)$ where $P : \mathbb{R} \rightarrow \mathbb{R}$ is linear, $t \geq 0$.

We focus on Nash equilibria in linear Markov strategies. Thus, given a linear pricing rule $P(\cdot)$, an *admissible* strategy for the consumer is any process $(Q_t)_{t \geq 0}$ taking values in \mathbb{R} and 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; (ii) and (iii) are purely technical.¹⁷

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

(i) when firms price using $P(\cdot)$, the policy $(\theta, Y) \mapsto Q(P(Y), \theta, Y)$ maximizes

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

subject to (2) and (5), among all admissible strategies $(Q_t)_{t \geq 0}$; and

(ii) whenever $Y_t = y$, $p = P(y)$ solves $\max_{p \in \mathbb{R}} p \mathbb{E}[Q(p, \theta_t, y) | Y_t = y]$.

A linear pair (Q, P) is a stationary linear Markov equilibrium if, in addition, the type-score process $(\theta_t, Y_t)_{t \geq 0}$ induced by $(\theta, Y) \mapsto Q(P(Y), \theta, Y)$ 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. Examining deviations from a prescribed price is, however, critical for determining the *price sensitivity of demand*, i.e., dQ/dp . In Section 4.2, we select a value for this sensitivity, thus refining our solution concept to provide an analog of Markov perfect equilibrium. Finally, the stationarity notion encompasses two ideas: the (controlled) score must admit a long-run distribution with finite moments, and such a distribution must hold at all times. These two properties allow us to perform a meaningful welfare analysis that is also time-invariant.

¹⁷Under (ii), (5) admits a strong solution given any initial condition; therefore, the consumer's best-response problem is well defined (Section 3.2 in Pham (2009)). Finally, (iii) is a mild strengthening of the condition $\mathbb{E}_0 \left[\int_0^\infty e^{-rt} |u(\theta_t, P(Y_t), Q_t)| dt \right] < \infty$ that is usually imposed in verification theorems (Sections 3.2 and 3.5 in Pham (2009)). In particular, it rules out strategies with the unappealing property of yielding high payoffs by making expenditures, $\int_0^\infty e^{-rt} P_t Q_t dt$, very negative (provided they exist).

3 Naive Consumers

To assess the equilibrium consequences of persistence and transparency, we describe the benchmark case of naive consumers. Specifically, consider a consumer with preferences as in (1) who ignores the link between her current action and future prices. Given a posted price p , maximizing the consumer's flow payoff yields a demand with a unit slope $Q(p) = \theta - p$. Each firm t observes the consumer's score Y_t prior to setting the monopoly price. Letting $M_t := \mathbb{E}[\theta_t|Y_t]$, the equilibrium quantity and price are given by

$$Q_t = \theta_t - M_t/2 \quad \text{and} \quad P_t = M_t/2.$$

The ex ante expected profit and consumer surplus levels are given by

$$\Pi_Y^{\text{static}} = \frac{1}{4} (\mu^2 + \text{Var}[M_t]) \quad \text{and} \quad CS_Y^{\text{static}} = \frac{1}{2} \text{Var}[\theta] + \frac{\mu^2}{8} - \frac{3}{8} \text{Var}[M_t],$$

where the (ex ante) variability of the posterior mean, $\text{Var}[M_t]$, measures the precision of the firms' information.

Because demand is linear, the average price and quantity levels (both equal to $\mu/2$) are independent of the information structure. The welfare consequences of using scores to price discriminate are thus fully determined by the firms' ability to learn from such signals. On the one hand, better information increases firms' profits by allowing them to better tailor the price to the consumer's type. On the other hand, with a constant average quantity, total surplus must fall with greater price discrimination because the correlation between the type and the price reduces the degree of correlation between the type and the quantity purchased. Therefore, the consumer must be unambiguously worse off.

When the consumer is naive, each firm t would like to access the full set of disaggregated purchase signals $\xi^t := (\xi_s : s < t)$, and this benchmark can be attained by an exponential score. Specifically, as we state in Section 5, the (stationary) posterior expectation that arises under the observation of disaggregated data in this case, $\mathbb{E}[\theta_t|\xi^t]$, is an affine function of

$$\int_0^t e^{-\hat{\phi}(t-s)} d\xi_s, \quad t > 0,$$

where $\hat{\phi} > 0$ denotes the (optimal) weight that the Kalman filter uses to discount past information when purchases follow $Q_t = \theta_t - \mathbb{E}[\theta_t|\xi^t]/2$. This, in turn, implies that learning from disaggregated data ξ^t or from the contemporaneous value of a (stationary) score (4) of persistence $\hat{\phi}$ leads to identical beliefs. Moreover, by the definition of the Kalman filter, dis-

counting past recorded purchases with an exponential weight $\hat{\phi}$ maximizes firms' learning.¹⁸ We summarize our findings for the naive case in the next result.

Proposition 1 (Naive Benchmark).

1. *Consumer and total surplus are decreasing in the precision of the firms' information.*
2. *Firm profits are increasing in the precision of the firms' information.*
3. *Firm learning is maximized by observing the disaggregated history of signals.*
4. *There exists a unique $\hat{\phi} > 0$ such that observing the value of a score process with persistence $\hat{\phi}$ is equivalent to observing the corresponding disaggregated history of signals.*

Two corollaries are distilled from this result. First, scores of persistence $\phi \neq \hat{\phi}$ hinder the firms' learning, as information then ceases to be aggregated optimally. This is not qualitatively different from adding noise to the technology $(\xi_t)_{t \geq 0}$: since the behavior of a naive consumer is fixed, the signal-to-noise ratio in $(\xi_t)_{t \geq 0}$ worsens. Second, because the naive benchmark is equivalent to a repetition of static interactions (albeit with varying information on the firms' side), the only channel through which a strategic consumer can benefit from information collection is by changing her demand in a dynamic environment.

In contrast, a strategic consumer understands that larger purchases today lead to higher future prices due to the persistence in the score process. This paves the way for the *ratchet effect*. In Section 4, for any given $\phi > 0$, we describe an equilibrium in which realized purchases follow $Q_t = \alpha\theta_t + \beta M_t + \delta\mu$, which generalizes $Q_t = \theta_t - M_t/2$ in the naive benchmark. The dependence of the tuple (α, β, δ) on $\phi > 0$ captures the strategic effects of the score's persistence, and the difference between (α, β, δ) and the naive values $(1, -1/2, 0)$ measures the strength of the ratchet effect.

The economic implications of consumer sophistication are twofold. First, the precision of the firms' information no longer fully determines the welfare consequences of price discrimination: (α, β, δ) encode demand adjustments in response to scores of different persistence, which make the average quantity purchased and price paid no longer constant (the *level* implications of the ratchet effect). Second, varying $\phi > 0$ not only affects the score's informativeness directly through the way in which information is aggregated: it also does so indirectly by endogenizing the signal-to-noise ratio in the purchase signal since the weight

¹⁸These results are a particular instance of Proposition 3 in Section 5, where we show that given purchase process of the form $Q_t = \alpha\theta_t + \beta M_t + \delta\mu$, the persistence level that maximizes the firms' learning is given by $\nu(\alpha, \beta)$, where the function $\nu(\cdot, \cdot)$ is defined in (18). Thus, $\hat{\phi} = \nu(1, -1/2)$.

on the type (α) depends on ϕ (the *informational* implications of the ratchet effect). In particular, adding noise to the recorded purchases and suboptimally aggregating the available information in a statistical sense cease to have identical implications for learning.

4 Equilibrium Analysis with Observable Scores

In this section, we characterize a stationary linear Markov equilibrium with realized purchases

$$Q_t = \alpha\theta_t + \beta M_t + \delta\mu, \quad (6)$$

where the coefficients depend on $\phi > 0$. We proceed in three steps: (i) we characterize the firms' (stationary Gaussian) beliefs when learning from the score; (ii) we determine the price sensitivity of demand that pins down the firms' monopoly price; and (iii) we solve the consumer's dynamic optimization problem.

4.1 Stationary Beliefs

Stationarity imposes two restrictions in our model. First, each firm t must use the same rule to update its beliefs. Second, the process (θ_t, Y_t) has to admit a long-run distribution with finite moments that is initialized at time 0 (via an appropriate choice of (θ_0, Y_0)) and that is consistent with the firms' updating rule via (6).

If (θ_t, Y_t) is Gaussian, however, the projection theorem for Gaussian random variables yields the following linear updating rule,

$$M_t := \mathbb{E}[\theta_t | Y_t] = \underbrace{\mu}_{=\mathbb{E}[\theta_t]} + \frac{\text{Cov}[\theta_t, Y_t]}{\text{Var}[Y_t]} [Y_t - \mathbb{E}[Y]].$$

We then require that $\text{Cov}[\theta_t, Y_t]/\text{Var}[Y_t]$ and $\mathbb{E}[Y]$ are independent of time; denote them λ and \bar{Y} , respectively. Thus, $M_t = \mu + \lambda[Y_t - \bar{Y}]$ holds at all times.

Lemma A.1 in the Appendix shows that the second restriction reduces to $\phi - \beta\lambda > 0$ and

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

The first condition states that, for a long-run Gaussian distribution with finite variance to exist, the score must have a positive rate of decay when $M_t = \mu + \lambda[Y_t - \bar{Y}]$ enters (6). The second condition, (7), is the consistency requirement: in $\lambda = \text{Cov}[\theta_t, Y_t]/\text{Var}[Y_t]$, the right-hand side, via Y_t , contains past beliefs that depend on λ via $M_s = \mu + \lambda[Y_s - \bar{Y}]$, $s < t$.

The regression coefficient λ measures the responsiveness of beliefs to changes in the score, and it plays a central role in our analysis. In fact, using $M_t = \mu + \lambda[Y_t - \bar{Y}]$, we can recast the problem of controlling the score as one of controlling the firms' beliefs, namely,

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

Thus, the consumer's choice of quantity affects the current firm's belief linearly with a slope of λ , and this effect decays at rate ϕ . The following result underscores a key tension between the short- and long-term responses of beliefs as we vary the persistence of the score:

Lemma 1 (Persistence and Sensitivity). *λ that solves (7) is strictly increasing in ϕ .*

Intuitively, as a score puts more weight on past information, it correlates less with the current type. Beliefs then react less to new information, as captured by λ , and vice versa. In particular, endowing beliefs with persistence (and, a fortiori, prices) by making scores more persistent themselves is not for free: the short-term response of beliefs is diminished.

4.2 Price Sensitivity of Demand and Monopoly Pricing

In this section, we determine the price sensitivity of demand and characterize the firms' monopoly price. Because the score is observed by the consumer and the firms adopt a linear strategy, the consumer can perfectly anticipate the candidate equilibrium price. The price sensitivity of demand is then determined by the (optimal) change in the consumer's quantity demanded in response to a price deviation $p \neq P(Y_t)$. This poses a challenge in continuous time: imposing optimality of the consumer's strategy at such off-path histories does not pin down her response to a deviation, as every firm operates over a zero-measure set.

To overcome this challenge, we refine our stationary linear Markov equilibrium concept by requiring that prices be supported by the *limit sensitivity* of demand along a natural sequence of discrete-time games indexed by their period length. Along such a sequence, as the period length shrinks to zero, the limit demand sensitivity is equal to -1 .

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

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

When Δ is sufficiently small, the missing terms that are affected by q on the right-hand side have only second-order effects on the consumer's payoff; therefore, the impact of quantities on the continuation value becomes asymptotically linear. Furthermore, because firms do not observe past prices, the continuation game is unaffected by the actual choice of p . Using the fact that $e^{-r\Delta} \approx 1 - r\Delta$, the consumer's first-order condition satisfies

$$Q_t = \theta_t - p + \lambda \frac{\partial V_t}{\partial M_t}, \quad (9)$$

where $\partial V_t / \partial M_t$ is independent of the posted price, which leads to a slope of demand of value -1 . In other words, the incentives to manipulate the firms' beliefs affect the intercept but not the slope of the demand function. Henceforth, except for the case of hidden scores, a stationary linear Markov equilibrium is understood to have a unit price sensitivity.¹⁹

Having pinned down this sensitivity, we now characterize the monopoly price process along the path of play of any stationary linear Markov equilibrium.

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

$$P_t = (\alpha + \beta)M_t + \delta\mu, \quad t \geq 0. \quad (10)$$

The intuition is simple: because the demand has a 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 to a price process P_t with parameters (α, β, δ) as a linear-quadratic optimization problem,

$$\begin{aligned} \max_{(Q_t)_{t \geq 0}} \quad & \mathbb{E}_0 \left[\int_0^\infty e^{-rt} \left[(\theta_t - P_t)Q_t - \frac{Q_t^2}{2} \right] dt \right] \\ \text{s.t.} \quad & d\theta_t = -\kappa(\theta_t - \mu)dt + \sigma_\theta dZ_t^\theta \\ & dM_t = (-\phi[M_t - \mu + \lambda\bar{Y}] + \lambda Q_t)dt + \lambda\sigma_\xi dZ_t^\xi \\ & P_t = (\alpha + \beta)M_t + \delta\mu, \end{aligned}$$

where λ satisfies (7).

¹⁹In the Supplementary Appendix (Section S.4), we examine a sequence of discrete-time games that employ the usual discretized version of diffusions in which the noise is scaled by $\sqrt{\Delta}$. We show that, along this sequence, (i) linear best replies on the path of play are also optimal after observing off-path prices and that (ii) the weight that linear best replies attach to the current price converges to -1 as the period length goes to zero.

4.3 Stationary Linear Markov Equilibria

To characterize stationary linear Markov equilibria, we use standard dynamic-programming tools. In a nutshell, we look for a quadratic value function and impose the condition that the firms correctly anticipate the consumer's behavior. This yields a subsystem of equations for the equilibrium coefficients (α, β, δ) that is coupled with equation (7) to pin down the equilibrium sensitivity of beliefs λ ; we then look for a solution that satisfies the stationarity condition $\phi - \beta\lambda > 0$. As we show in the proof of Theorem 1, there is a unique such solution to this system, which in turn allows us to establish the existence and uniqueness of an equilibrium in this class. Furthermore, the equilibrium can be computed in closed form, up to the solution of a single algebraic equation for the coefficient α .

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

$$\alpha = 1 + \frac{\Lambda(\phi, \alpha, B(\phi, \alpha))\alpha B(\phi, \alpha)}{r + \kappa + \phi}, \quad \alpha \in [0, 1]. \quad (11)$$

where the functions B , Λ , and D are defined in (A.7), (A.8), and (A.12), respectively. Moreover, $\beta = B(\phi, \alpha) \in (-\alpha/2, 0)$, $\delta = D(\phi, \alpha) \in \mathbb{R}$, and $\lambda = \Lambda(\phi, \alpha, B(\phi, \alpha)) > 0$.

Figure 2 illustrates the equilibrium coefficients $(\alpha, -\beta, \delta)$, their naive benchmark levels, and the average equilibrium price (and quantity)

$$\mathbb{E}[P_t] = \mathbb{E}[Q_t] = [\alpha + \beta + \delta]\mu. \quad (12)$$




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

4.4 Strategic Demand Reduction: The Ratchet Effect

Recall the first-order condition (9) of the consumer's problem, $Q_t = \theta_t - P_t + \lambda \partial V(\theta, M)/\partial M$. The consumer's strategic behavior is then summarized by $\lambda \partial V(\theta, M)/\partial M$, i.e., by the wedge between her actual behavior and the myopic counterpart. Intuitively, we would expect this derivative to represent a *ratchet effect*. To confirm this intuition, however, we must show that this derivative actually encodes the value of a downward deviation from a natural benchmark due to an adversarial incentive scheme in place.

The next result therefore provides a representation for this derivative in the form of expenditure savings from inducing a lower price: the natural benchmark is given by myopic play, and the incentive scheme by the schedule of future prices faced by the consumer. We further state the level and informational consequences of the ratchet effect.

Proposition 2 (Ratchet Effect).

(i) *Value of future savings: equilibrium prices and quantities satisfy*

$$Q_t = \theta_t - P_t - \Psi_t, \text{ where} \quad (13)$$

$$\Psi_t := \lambda \mathbb{E}_t \left[\int_t^\infty e^{-(r+\phi)(s-t)} (\alpha + \beta) Q_s ds \right], \quad t \geq 0. \quad (14)$$

(ii) *Signaling coefficient and average prices and quantities:*

$$1/2 < \alpha(\phi) < 1 \quad \text{and} \quad \mathbb{E}[P_t] = \mathbb{E}[Q_t] \in (\mu/3, \mu/2).$$

The process Ψ_t is the *value of future (expenditure) savings from a small reduction in today's quantity*. In fact, by the Envelope Theorem, the benefit of a marginal reduction in today's quantity is equal to the net present value of the associated reduction in future prices, holding the future quantities constant. Applied to our case, the price P_{t+dt} falls by $\lambda(\alpha+\beta)$ after such deviation (because M_{t+dt} falls by λ), while the impact on subsequent prices then vanishes at the rate ϕ at which beliefs decay afterwards.²⁰ Comparing the first-order condition with (13)–(14) yields $\lambda \partial V/\partial M = \Psi_t$, and hence a ratchet effect ensues: the larger the savings, the stronger the incentive to deviate downwards from the static equilibrium.

The ratchet effect implies that, on average, all types purchase less, i.e., the average price and quantity fall relative to a static interaction. Specifically, using $Q_s = \alpha \theta_s + \beta M_s + \delta \mu$

²⁰While we prove the result only in equilibrium, (13)–(14) is an optimality condition and thus holds at a greater level of generality. Our proof uses the Envelope Theorem and the Feynman-Kac formula, as in [Abel and Eberly \(1994\)](#). See also [Strulovici \(2011\)](#) for a similar method in a contracting environment.

and taking expectations under the prior distribution of $(\theta_t)_{t \geq 0}$ in (13), we obtain

$$\mathbb{E}[\Psi_t] = \lambda \frac{(\alpha + \beta)(\alpha + \beta + \delta)\mu}{r + \phi}. \quad (15)$$

Because $\alpha + \beta > 0$ (Theorem 1) and $\mathbb{E}[P_t] = \mathbb{E}[Q_t] = (\alpha + \beta + \delta)\mu$, one can easily conclude from (13)–(14) that (15) is strictly positive and that the average quantity demanded contracts below the static level $\mu/2$, as in Figure 2 (right panel).²¹

The incentives for demand reduction are not, however, uniform across consumer types. In particular, the value of future savings in Proposition 2 satisfies

$$\frac{\partial \Psi_t}{\partial \theta_t} = \frac{\lambda \alpha \beta}{r + \kappa + \phi}. \quad (16)$$

Comparing (16) with the right-hand side of equation (11) from Theorem 1, we conclude that $\alpha(\phi)$ is simply the (static) unit weight attached to the type θ , diminished exactly by the *sensitivity of the value of future savings to θ* .²² The benefits of a downward deviation are then greater for higher types because, due to the persistence in the types process, a high θ_t is more likely to buy larger quantities in the future, and hence obtains higher savings from strategically reducing her demand (i.e., $\alpha < 1$).

Finally, the lower bounds on the ratchet effect in (ii) are the consequence of the *strategic substitutability* between the consumer's actual choices and the ones conjectured by the firms. In particular, if the firms believed that the quantity signals were uninformative ($\alpha = 0$), the consumer would have no incentive to deviate from optimal myopic behavior $\alpha = 1$. Several additional properties of the equilibrium coefficients and outcomes that are key technical steps for our subsequent results can be found in Lemma A.4 in the Appendix.²³

5 Equilibrium Learning

The ability of firms to price discriminate depends on their ability to learn from the score. In this section we show that, unlike in the naive benchmark, there is a wedge between maximizing learning and the use of disaggregated data. Specifically, because of the ratchet effect, the firms' learning is maximized by *relatively persistent* scores: scores that discount

²¹A straightforward manipulation of terms reveals that $\mathbb{E}[P_t] = \mu(r + \phi)/[2(r + \phi) + \lambda(\alpha + \beta)]$.

²²The denominator reflects that a marginal change in today's type on all future savings decays at a rate $r + \kappa + \phi$ because the types themselves decay at a rate κ . The numerator indicates that an increase in θ_t positively affects not only future types θ_s but also future belief realizations M_s , $s > t$.

²³These include: tighter bounds as a function of ϕ , asymptotic behavior as $\phi \rightarrow \{0, +\infty\}$; and comparative statics with respect to the noise volatility σ_ξ . We also establish that $\phi \mapsto \alpha(\phi)$ is quasiconvex.

signals too little relative to Bayesian updating based on the histories of observed purchases.

The extent of firms' learning is summarized by the following quantity:²⁴

$$\frac{\text{Cov}[\theta_t, Y_t]}{\text{Var}[\theta_t]\text{Var}[M_t]} = \frac{\alpha\lambda(\phi, \alpha(\phi), \beta(\phi))}{\phi + \kappa - \beta\lambda(\phi, \alpha(\phi), \beta(\phi))} := G(\phi, \alpha(\phi), \beta(\phi)) \in [0, 1]. \quad (17)$$

The function G highlights the two channels through which a score's persistence affects learning: directly via ϕ , which determines the weight attached to past signals (and its resulting impact on λ), and indirectly via the coefficients of the consumer's strategy. The presence of equilibrium effects thus opens the possibility for a score to give up on its optimality as a statistical filter, i.e., as an optimal aggregate of the underlying signals, in exchange for an improvement in the quality of such signals.

The question of how to maximize learning is ultimately one of how to optimally aggregate information accounting for these two channels. It is natural to use the case of disaggregated signals ($\xi_s : 0 \leq s < t$) as a reference point. To this end, we hold *fixed* the consumer's behavior $(\alpha, \beta) \in \mathbb{R}_+ \times \mathbb{R}_-$, and define

$$\nu(\alpha, \beta) := \kappa + \frac{\gamma(\alpha)\alpha(\alpha + \beta)}{\sigma_\xi^2}, \quad (18)$$

where $\gamma(\alpha) > 0$ is the steady-state variance of beliefs when the histories of $(\xi_t)_{t \geq 0}$ are observable.²⁵ We now establish the equivalence between learning from the history of disaggregated signals and from the current level of a score with persistence $\nu(\alpha, \beta) > 0$.

Proposition 3 (Disaggregated Histories). *Consider $(Q_t)_{t \geq 0}$ as in (6) with $\alpha + \beta > 0$.*

1. $\nu(\alpha, \beta) > 0$ is the unique maximizer of $G(\cdot, \alpha, \beta)$.
2. If firms observe the histories of $(\xi_t)_{t \geq 0}$ and their beliefs are stationary, then the posterior mean process is affine in a stationary Gaussian score (4) with $\phi = \nu(\alpha, \beta)$.
3. If firms observe only the current value of a stationary Gaussian score with $\phi = \nu(\alpha, \beta)$, their beliefs coincide with those that arise from observing the histories ξ^t , $t \geq 0$.

The persistence level $\nu(\alpha, \beta)$ aggregates the data generated by a linear strategy Q with coefficients (α, β) without loss of information for a given fixed behavior—it defines what an optimal score in a statistical sense would look like.²⁶ Given a score with generic persistence

²⁴By the projection theorem, $\text{Var}[\theta_t | Y_t] = \text{Var}[\theta_t](1 - \text{Cov}[\theta_t, Y_t]/\text{Var}[\theta_t]\text{Var}[M_t])$.

²⁵ $\gamma(\alpha)$ is the unique positive root of $x \mapsto \alpha^2 x^2 / \sigma_\xi^2 + 2\kappa x - \sigma_\theta^2 = 0$.

²⁶ $\nu(\alpha, \beta) > \kappa$, reflecting the identification problem faced by the firms while observing $(\xi_t)_{t \geq 0}$. The proof of Proposition 3 can be found in Section S.2.2 of the Supplementary Appendix.

$\phi > 0$, however, a strategic consumer need not choose the coefficients $(\alpha(\phi), \beta(\phi))$ for which the score is an optimal filter in a statistical sense, i.e., for which $\phi = \nu(\alpha(\phi), \beta(\phi))$. The importance of these fixed points is clear: if an equilibrium in which the firms have access to disaggregated signals exists, then, by Proposition 3, the weight with which the associated beliefs discount past purchase signals must solve $\phi = \nu(\alpha(\phi), \beta(\phi))$.

Definition 2 (Nonconcealing score). *A score with persistence $\phi > 0$ is nonconcealing if*

$$\phi = \nu(\alpha(\phi), \beta(\phi)). \quad (19)$$

We now introduce the main result of this section: there exists a unique of a solution to (19) and the firms' learning is maximized to the left of it. That is, learning is enhanced by scores that are more persistent than the unique nonconcealing score, despite such scores concealing some information about the consumer's behavior (as $\phi \neq \nu(\alpha(\phi), \beta(\phi))$ in those cases).²⁷ The reason is that the consumer signals her type more aggressively in that case. Without fear of confusion, let $G(\phi) := G(\phi, \alpha(\phi), \beta(\phi))$, denote the *equilibrium gain* function.

Proposition 4 (Uniqueness of a Nonconcealing Score and Signaling).

- (i) *There exists a unique $\phi^* \in \mathbb{R}_+$ solving $\phi = \nu(\alpha(\phi), \beta(\phi))$.*
- (ii) *The coefficient $\alpha(\cdot)$ is strictly decreasing at the fixed point $\phi = \phi^*$.*
- (iii) *The equilibrium gain function $G(\phi) \in [0, 1]$ is maximized in $(0, \phi^*)$.*
- (iv) *The function $G(\phi; \sigma_\xi)$ is decreasing in σ_ξ for all $\phi > 0$.*

By the definition of an optimal filter, changing the persistence of the score has only a second-order effect on learning, holding (α, β) constant.²⁸ Increasing α , however, has a first-order effect on learning, as the score is now more sensitive to the consumer's type. Thus, the indirect effect on the consumer's incentives to reveal information drives the firms' learning around ϕ^* , while the direct effect of ϕ dominates away from the optimal filter ϕ^* . Figure 3 plots G as a function of $\phi - \nu(\phi)$: its maximum is located to the left of the vertical axis.

That a reduction in ϕ from ϕ^* increases α for all discount rates is somewhat surprising: even a very patient consumer finds it optimal to attach a higher weight to her type, despite

²⁷The uniqueness of a nonconcealing score establishes the uniqueness of a stationary linear Markov equilibrium when the histories of $(\xi_t)_{t \geq 0}$ are observable. At the heart of this result is the strategic substitutability between the firms' conjectured actions and the consumer's actual choices, as discussed in Section 4.4.

²⁸Marginally increasing β at $(\phi^*, \alpha(\phi^*), \beta(\phi^*))$, in turn, has no first-order effect on the amount of information transmitted either: β is the coefficient on M_t in the consumer's strategy, and at ϕ^* , the score perfectly accounts for the contribution of the beliefs to the recorded purchases.




Figure 3: $(r, \sigma_\theta, \sigma_\xi, \kappa) = (5, 1, 1/5, 1)$.

the consequences that more persistent scores can have for long-term prices. Opposing this force is the fact that a score that attaches an excessive weight to past signals also correlates less with the consumer's current type (sensitivity-persistence tradeoff). This results in a reduced sensitivity of beliefs (and hence, of prices) to changes in the score. In turn, less sensitive prices make the consumer less concerned about purchasing large quantities.²⁹

To see why the sensitivity effect is relatively stronger for all $r > 0$, recall that the coefficient α reflects the *relative* value of future savings for a marginally higher type θ_t , as derived in (16). From Section 4.4, the sensitivity of the value of future savings to the consumer's current type reflects both the direct impact of a shock to θ_t on future types (which decays at a rate κ) and its indirect impact on future prices (which depreciates at rate $r + \phi$) via the change in the quantity demanded.

With a linear relationship between $(Y_t)_{t \geq 0}$ and $(\theta_t)_{t \geq 0}$, the equilibrium gain function $G(\phi)$ is also akin to an impulse response, where a shock to a past type θ_s , $s < t$, has an impact on the past score Y_s that depreciates at rate ϕ . However, the type shock itself depreciates at rate κ . Loosely speaking then, the gain function is akin to the *undiscounted* impulse response of the marginal value of future savings to a shock to θ_t . The gain function and the marginal value of future savings, then, differ only in that discounting gives the immediate future more relevance in the latter. As a result, the sensitivity-persistence tradeoff is tilted in favor of the sensitivity effect. This, in turn, leads a consumer with *any* degree of time preference to become *more* responsive to her type when the impact of quantities on future prices is *backloaded* relative to the nonconcealing score ϕ^* , which facilitates price discrimination.

Finally, part (iv) shows how noise and persistence have qualitatively different effects on equilibrium learning when consumers are strategic. Lemma A.4 in the Appendix shows that increasing σ_ξ also increases α , just as reducing ϕ below ϕ^* does. However, while distorting

²⁹The trade-off between persistence and sensitivity also arises in signal-jamming models with symmetric uncertainty. See, for example, Cisternas (2017) in the context of career concerns.

persistence away from the optimal filter triggers sufficiently strong equilibrium effects, adding noise to the purchase signals reduces G unambiguously: the intuition is that, unlike moving ϕ around ϕ^* (a choice variable), increasing σ_ξ (a parameter) has a negative first-order effect on learning, which trumps the increase in α . In other words, adding exogenous noise to purchase signals is an inferior means to manage the ratchet effect.³⁰

6 Welfare Analysis

We now turn to the welfare consequences of observable scores. We show that firms can profit from relatively persistent scores, and that such scores do not necessarily hurt consumers relative to a setting without price discrimination. Omitting the dependence of P_t and M_t on ϕ and using the relationship $\mathbb{E}[Q_t|Y_t] = P_t$, firm t 's ex ante profits are given by

$$\Pi(\phi) := \mathbb{E}[P_t Q_t] = \mathbb{E}[P_t^2] = \mathbb{E}[P_t]^2 + \text{Var}[P_t], \quad t \geq 0. \quad (20)$$

A similar calculation yields the ex ante flow consumer surplus,

$$CS(\phi) = \mathbb{E}[P_t] \left(\mu - \frac{3}{2} \mathbb{E}[P_t] \right) + L(\phi) \text{Var}[P_t] + \alpha(\phi) \left(1 - \frac{\alpha(\phi)}{2} \right) \text{Var}[\theta_t], \quad (21)$$

where $L(\phi) := \frac{\alpha(\phi)^2/2 + \beta(\phi)}{(\alpha(\phi) + \beta(\phi))^2} - \frac{3}{2}$ is a negative function.³¹

We know from Proposition 2 that $\mathbb{E}[P_t] > \mu/3$. Therefore, consumer surplus is decreasing, and producer surplus is increasing, both in the expected price level and in the firms' ability to tailor prices based on the information contained in the score, the latter measured by the ex ante variability of the price.³² Moreover, these two moments respond very differently to the score's persistence ϕ and to the average willingness to pay μ . Indeed, we have

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

On the one hand, the expected price is largest for uninformative scores ($\phi = 0$ and $\phi \rightarrow \infty$): the ratchet effect benefits the consumer through a lower expected price for all informative

³⁰In the Supplementary Appendix Section S.2.2 we establish a stronger result for the case of public disaggregated histories: the equilibrium gain $G(\phi^*(\sigma_\xi), \sigma_\xi)$ is decreasing in σ_ξ , even though $\phi^*(\cdot)$ is itself decreasing. Thus, additional noise is not conducive to more learning in a world without scores.

³¹The derivation of the expressions for $CS(\phi)$ and $\Pi(\phi)$, as well as the determination of the sign of L can be found in Section S.2.3 in the Supplementary Appendix.

³²Moreover, $\alpha(\phi) \leq 1$ implies that the third term is increasing in α : by shading down her demand, the consumer moves away from her static optimum, reducing her surplus.

scores. On the other hand, the variance of the price inherits all of the properties of the equilibrium gain function G derived in Proposition 4. In particular, it is maximized by some $\phi \in (0, \phi^*)$, i.e., by an informative and persistent score.³³

The effect of the expected price on consumer surplus is proportional to μ^2 : the benefit of low prices is higher when the average willingness to pay μ is high and discounts are applied to a larger number of units. Instead, the costs of price discrimination for the consumer are independent of her average willingness to pay μ . Thus, consumers with a high μ benefit more from the availability of information than those with a low μ , and firms derive a net benefit from informative scores only if μ is low enough. See Figures 4 and 5.




Figure 4: Consumer Surplus: $(\sigma_\theta, \sigma_\xi, \kappa, r) = (1, 0, 1, 0.8)$ and $\mu \in \{1, 2, 3\}$




Figure 5: Producer Surplus: $(\sigma_\theta, \sigma_\xi, \kappa, r) = (1, 0, 1, 0.8)$ and $\mu \in \{1, 2, 3\}$

The presence of several nonlinear terms in the expressions for consumer surplus and firms' profits makes the full characterization of the associated optimal degrees of persistence as a function of primitives a daunting task. For this reason, we specialize some of our welfare

³³This is also proved in Section S.2.3 in the Supplementary Appendix. The variance of the price can be interpreted as the *value of information* to the firms: since $\text{Var}[P_t] = \mathbb{E}[P_t Q_t] - \mathbb{E}[P_t] \cdot \mathbb{E}[Q_t]$, the variance $\text{Var}[P_t]$ measures the supplemental profits relative to pricing under the consumer's equilibrium strategy, with the knowledge only of the prior distribution. As such, it is a good proxy for a monopolist data broker's profit. See [Bergemann and Bonatti \(2019\)](#) for a model of intermediation in the market for consumer information.

results to the noiseless limit $\sigma_\xi \searrow 0$, for which further insights can be obtained.³⁴

Let ϕ^c and ϕ^f denote the consumer and firm optimal persistence levels, respectively.

Proposition 5 (Optimal Persistence). *For all $\sigma_\xi > 0$:*

(i) ϕ^f is interior for sufficiently low $\mu > 0$, and $\phi^f \in \{0, \infty\}$ for μ above a threshold;

(ii) ϕ^c is interior for μ above a threshold, and $\phi^c \in \{0, \infty\}$ for sufficiently low $\mu > 0$;

In the pointwise limit of $\Pi(\cdot)$ and $CS(\cdot)$ as $\sigma_\xi \searrow 0$:

(iii) for all $\mu > 0$, there exists a firm-optimal $\phi^f < \infty = \lim_{\sigma_\xi \searrow 0} \phi^*(\sigma_\xi)$;

(iv) there is an interval of values of r/κ over which ϕ^f and ϕ^c are continuous and monotone.

Moreover, $\phi^f > 0$ if and only if $\mu < \sqrt{r}/\kappa$ and $\phi^c > 0$ if and only if $\mu > 3\sqrt{r}/\kappa$.

Parts (i) and (ii) establish sufficient conditions for the informative and uninformative optima for the firms and the consumer. Part (iii) formalizes our intuition that firms benefit from aggregating information into (excessively) persistent scores.³⁵ Figure 6 illustrates the result in part (iv): there exists a range of μ over which all market participants prefer uninformative scores, which is intuitive, because information reduces total surplus.




Figure 6: Optimal Persistence, $(\sigma_\theta, \sigma_\xi, \kappa, r) = (1, 0, 1, 0.8)$

Finally, we contrast the expected consumer surplus in the strategic case with the naive-consumer benchmark (Section 3) in the noiseless limit case. Intuitively, a naive consumer, who demands the static quantity as in Section 3, suffers the costs of tailored prices without reaping the benefits of lower average prices. Thus, regardless of whether firms or consumers would prefer informative or uninformative scores, consumers should be better off acting strategically than following the myopic demand $Q_t = \theta_t - p$.

³⁴When $\sigma_\xi = 0$, all coefficients can be solved for in closed form, and the model is of class C^1 around that point (Section S.3 in the Supplementary Appendix).

³⁵When $\sigma_\xi \searrow 0$, we have $\phi^* \rightarrow \infty$ but $G(\phi^*) \rightarrow 1$ due to $\lambda \rightarrow \infty$. Therefore, the nonconcealing score maximizes learning in this case but not profits.

Proposition 6 (Naïve vs. Strategic Consumers). *In the limit as $\sigma_\xi \searrow 0$, consumer surplus is larger when consumers are strategic than when they are naïve for all $\phi > 0$.*

In Section 7, we show that this result is more subtle than it seems: strategic consumers who cannot observe their score would sometimes be better off acting myopically.³⁶

7 Hidden Scores

In this section, we study the case in which the score Y_t is observed by firm t but hidden to the consumer for all $t \geq 0$. The goal of this exercise is twofold: first, to better understand the mechanism by which directly observing her score can help a strategic consumer; second, to predict the welfare implications of growing consumer sophistication (i.e., concerns regarding discriminatory practices) under alternative information structures.

Relative to Section 2, because the consumer no longer directly observes her score, we now suppress the dependence of a consumer's *linear Markov* strategy on Y_t . Instead, a linear strategy for the firms is as in the baseline model. Thus, the objects of interest are

$$Q(\theta, p) = \delta^h \mu + \alpha^h \theta + \zeta^h p \text{ and} \quad (22)$$

$$P(Y) = \pi_0^h + \pi_1^h Y, \quad (23)$$

where the superscript h stands for *hidden*. The corresponding concepts of admissible strategies, equilibrium, and stationarity are all straightforward modifications of those introduced in Section 2.³⁷ We again focus on stationary linear Markov equilibria.

Hidden scores have important strategic implications. First, both the firms and the consumer can now signal their private information. In particular, if $\pi_1^h \neq 0$, the firm's strategy is invertible, and hence the consumer perfectly learns her score along the path of play; the consumer then has the *same information* as in the observable case on the equilibrium path. However, by signaling the level of the consumer's current score, today's price provides information about future firms' beliefs, and hence about future prices. As it turns out, this informational channel deeply affects the consumer's incentives.³⁸

³⁶Proposition 6 can be strengthened to hold in neighborhoods of $\sigma_\xi = 0$ for compact sets of persistence levels. This is because the equilibrium variables are of class C^1 as functions of $(\phi, \sigma_\xi) \in (0, \infty) \times [0, \infty)$, and so they converge uniformly over compact sets as $\sigma_\xi \searrow 0$ (refer to Section S.3.3 in the Supplementary Appendix for details). The same applies to Proposition 9 in the next section.

³⁷For consistency, therefore, we consider admissible strategies that condition on $(\theta_t, P_t)_{t \geq 0}$, rather than on $(\theta_t, Y_t)_{t \geq 0}$, in the consumer's problem. Observe, however, that this choice is innocuous when $\pi_1^h \neq 0$.

³⁸See Cisternas and Kolb (2019) for a model in which a myopic player privately monitors the actions of a long-lived player with private information, but where both players observe actions distorted by noise.

Second, from a technical perspective, the price sensitivity of demand ζ^h is determined along the equilibrium path, unlike in the case of observable scores, where off-path prices are required. The motivation comes from discrete time: with a hidden score that has full-support noise, (i) the consumer is not able to predict the next period's price using today's observation and (ii) any price realization is possible. Thus, the price process induced by a linear strategy exhibits the required intratemporal variation to be able to identify the slope of demand.³⁹

7.1 Equilibrium Analysis with Hidden Scores

Turning to the equilibrium analysis, it is immediately clear that any equilibrium must entail $\zeta^h < 0$. Hence, firm t sets the monopoly 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 purchase process is of the form

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

where $\beta^h := -\alpha^h/2$ and $M_t = \mu + \lambda^h[Y_t - \bar{Y}^h]$ for some λ^h and \bar{Y}^h , $t \geq 0$. In particular, the realized prices and quantities satisfy $P_t = -\mathbb{E}[Q_t|Y_t]/\zeta^h$ along the path of play.⁴⁰

As in Theorem 1, the quest for stationary linear Markov equilibria reduces to a single equation for the coefficient α^h on the consumer's type. This equation is identical to (11) in the observable case, replacing $B(\phi, \alpha) \in (-\alpha/2, 0)$ with $-\alpha/2$.⁴¹

Proposition 7 (Existence and Uniqueness). *There exists a unique stationary linear Markov equilibrium. In this equilibrium, $\alpha^h \in (0, 1)$ and the price sensitivity of demand is given by*

$$\zeta^h = -\frac{2(r + 2\phi)}{2(r + 2\phi) + \alpha^h \lambda^h} \in \left(-1, -\frac{r + 2\phi}{r + 3\phi} \right). \quad (25)$$

Crucially, demand turns out to be *less* price sensitive than in the observable case (i.e., $\zeta^h > -1$). The reason is the informational content of prices when scores are hidden: by

³⁹In continuous time, the price process induced by a linear Markov pricing strategy will have continuous paths, so deviations can be detected. Because with full-support noise this issue arises only in continuous time, we refine our equilibrium in the continuous-time game by assuming that the firms conjecture that the consumer responds to a (intratemporal) deviation with a sensitivity that coincides with the (intertemporal) sensitivity of the quantity demanded along the path of play of a candidate Nash equilibrium. Thus, as in discrete time, the same candidate dynamic policy $Q(\theta, p)$ is used by the firms in their pricing problem.

⁴⁰Since the quantity demanded (24) has the same structure as in (6), the characterization of the stationary beliefs in Section 4.1 applies to the hidden case.

⁴¹Many properties of the baseline model hold here—see section S.2.4 in the Supplementary Appendix.

informing the consumer that her score is high, a high price today is a signal of high prices tomorrow and, hence, of low quantities purchased in the future. This, in turn, diminishes the scope for scaling back current purchases to reduce the price. Conversely, the advantage of reducing prices is greater when prices are low, and the consumer is likely to buy more in the near future. Formally, the marginal value of manipulating future prices downward is decreasing in p due to the convexity of the consumer's value function as shown in Figure 7.




Figure 7: Signaling effect of prices.

This *signaling effect* of prices reduces the incentives for downward quantity deviations relative to the observable case, *at any given price*. However, as we show in Proposition 8 below, the direct effect of the reduced price sensitivity drives down the average quantity traded via higher posted prices. To unify the notation, we rewrite the realized demand and prices along the equilibrium path in the observable case in (θ, P) space, instead of (θ, M) , as follows:

$$Q_t = \delta^o \mu + \alpha^o \theta_t + \zeta^o P_t \quad \text{and} \quad P_t = \pi_0^o + \pi_1^o Y_t. \quad (26)$$

Let (Q^o, P^o) and (Q^h, P^h) denote the average (quantity, price) pairs in the cases examined.

Proposition 8 (Role of Transparency: strategies). *In equilibrium, for all $\phi > 0$:*

- (i) *Sensitivity of price to score:* $\pi_1^h(\phi) > \pi_1^o(\phi) > 0$.
- (ii) *Sensitivity of demand to type:* $1 > \alpha^o(\phi) > \alpha^h(\phi) > 0$
- (iii) *Average prices and quantities:* $\mu/2 > P^h(\phi) > P^o(\phi) = Q^o(\phi) > Q^h(\phi) > \mu/4$.

Facing a demand that is *less sensitive* to price than in the observable case (i.e., $0 > \zeta^h > -1$), each firm charges a price that is *more sensitive* to changes in the score relative to the observable case (i.e., $\pi_1^h(\phi) > \pi_1^o(\phi) > 0$). However, when prices are more sensitive to the

score, the ratchet effect is stronger, i.e., $Q^o > Q^h$. In fact, by the Envelope Theorem, the value of future savings in this hidden case takes the form

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

Since $\pi_1^h > \pi_1^o = \lambda(\alpha + \beta)$, the equilibrium value of a downward deviation increases. These incentives are also stronger for higher types, explaining the ranking of signaling coefficients.

Finally, while the average prices remain below those of the no-information case, they are higher than when scores are observed by the consumer: the reduced price sensitivity more than compensates for the strengthening of the ratchet effect.

7.2 Consumer Surplus: Observable vs. Hidden Scores

The previous properties of prices and quantities strongly suggest that consumers are worse off without transparency. While it is difficult to prove such conjecture at a general level, we are able to confirm these intuitions in the noiseless limit case $\sigma_\xi \searrow 0$, by taking advantage of closed-form expressions in the solutions to both models when $\sigma_\xi = 0$.⁴² Furthermore, we show that, somewhat surprisingly, eliminating consumer naiveté without providing transparent scores is not necessarily beneficial to the consumer.

Proposition 9 (Role of Transparency: Consumer Surplus). *In the limit as $\sigma_\xi \searrow 0$:*

- (i) *for all $\phi > 0$ consumer surplus is larger when scores are observable than hidden;*
- (ii) *let $\rho := r/\kappa < 4$ and $\mu < \min \left\{ \sqrt{\frac{4-\rho}{2\rho}}, \frac{\sqrt{2\rho(2\rho+1)}}{\rho+1} \right\}$. Then, consumer surplus with hidden scores is larger for all ϕ when consumers are naive than when they are strategic.*

The key behind (i) is that observing the score allows consumers to disentangle current from future prices: an abnormally high price today does not imply that future prices will be high. Thus, while transparency does not add to the consumers' information in equilibrium, it enables consumers to eliminate the signaling effect of prices; demand is then more price sensitive, which translates into lower prices and an increased consumer surplus. Intuitively, a consumer buys a considerably lower quantity if the price is too high, in the anticipation of future discounts, a possibility that is simply absent when scores are hidden. Moreover, it is

⁴²We thank an anonymous referee for this suggestion. Observe that the hidden-scores model exhibits a discontinuity at $\sigma_\xi = 0$. Specifically, by keeping track of her purchases, the consumer always knows her score when there is no noise: the hidden and observable models then coincide. However, because the price sensitivity of demand is determined differently in each case, the limits of the (hidden) coefficients as $\sigma_\xi \searrow 0$ differ from their observable counterparts. See Section S.3 of the Supplementary Appendix for more details.

possible to show that firms need not be better off with hidden scores: the stronger ratchet effects (reduced quantities and signaling) can outweigh the benefits of high prices.⁴³

Finally, part (ii) warns against awareness policies that are not combined with transparency: strategic consumers can be hurt by hidden scores relative to the naive case. Intuitively, unless the discount rate is high or the ratchet effect is very costly (high μ), committing to myopic behavior (i.e., ignoring the information content of prices) would benefit consumers, because it preserves the price sensitivity of demand that is lost when scores become hidden.

8 Applications

Our results have two policy implications for the debate on privacy. First, the results on optimal persistence (Section 6) do not support blanket regulations that eliminate data collection and transmission, as adverse uses of information can have positive effects for consumers: when purchase histories are tracked to personalize prices, *strategic* consumers implicitly demand compensation for the information they reveal.⁴⁴ Second, the results in Section 7 suggest the need for joint policy interventions: while transparency policies are always beneficial, they may in fact be *necessary*. In other words, partial regulation that promotes awareness may backfire if it simply makes consumers aware of the informational content of prices.

In this section, we evaluate the credibility of these policy recommendations in light of observed behavior in real-world markets. First, recall that our analysis relies on the assumption that strategic consumers—who are aware of the underlying mechanisms—manipulate their scores. To validate this assumption, we document various consumer responses to behavior-based price discrimination in a number of economically relevant settings. Put differently, a lack of information, and not consumer naiveté, is the main obstacle to score manipulation.

Second, our transparency recommendation exploits an equilibrium phenomenon by which strategic consumers become more price sensitive when they can *independently* verify the information used to set prices. Along these lines, Section 8.2 validates the plausibility of this mechanism by looking at evidence from the use of FICO credit scores in the U.S. Finally, relying on a growing suggestive evidence of personalized pricing that it is argued is happening today, Section 8.3 describes concrete markets where score-based price discrimination is likely to play a critical role in the near future.

⁴³In the proof of this result (Supplementary Appendix Section S.2.4) we show that this always happens for ϕ sufficiently large. This is due to a permanent wedge between the signaling coefficient in the hidden case and the observable counterpart as ϕ grows, which leads to both depressed learning and low quantities.

⁴⁴This is in sharp contrast to any information obtained by the firms from exogenous sources. As in our naive case, this information is bound to benefit firms and to harm consumers if it is later used against them.

8.1 Buyers' Strategic Behavior

Shopping cart abandonment There is widespread evidence of this problem for online merchants that occurs whenever shoppers fail to complete a purchase after selecting a product.⁴⁵ Abandonment often triggers an email by the vendor with a promotional offer, which suggests the profitability of searching for a specific product, and then waiting for a lower price.⁴⁶ Salesforce reports that approximately 40% of consumers open emails from stores about their abandoned shopping carts—a response rate far higher than that for regular marketing material, which is consistent with consumers' attempts to induce discounts.

Rideshare services Most consumers in the US are aware that rideshare platforms such as Uber personalize prices on the basis of individual and market characteristics, e.g., the use of a personal vs. business credit card, but also their past behavior on the app ([The Guardian, 2018](#)). In an attempt to lower their prices, many consumers experiment with various techniques to get around the algorithm. These strategies include requesting and then rejecting quotes for unnecessary rides, so to simulate greater price sensitivity, and changing the destination address mid-route, which leverages discrepancies in the pricing algorithm.⁴⁷

Online display advertising In this market, any publisher or website owner can be a seller of advertising space; and the demand for space comes from advertisers who wish to reach a targeted population of final consumers.⁴⁸ Sellers are able to price discriminate, both across buyers and over time, by choosing personalized reserve prices in real-time auctions for advertising space ([Kanoria and Nazerzadeh, 2017](#)). Information about past auction outcomes is aggregated and distributed by Supply-Side Platforms—technology platforms that help sellers manage their space, and can also set reserve prices on their behalf ([European Commission, 2018](#)).⁴⁹

A growing body of theoretical and empirical work recognizes the importance of the ratchet effect in this market. Furthermore, the ratchet effect influences practical algorithm and

⁴⁵Salecycle reports that approximately 77% of carts go unfinished (for a myriad of reasons). See <https://blog.salecycle.com/infographics/infographic-the-remarketing-report-q3-2018/>.

⁴⁶“Let it be: If you can bear to wait, try leaving your item in the shopping cart for a day. The retailer might send you an email offering a discount on that item. Or you might find a discount in your social feeds,” <https://clark.com/shopping-retail/online-shopping-hacks/>.

⁴⁷See <https://therideshareguy.com/uber-is-ripping-off-frequent-riders-and-heres-how-to-avoid-it/>. Interestingly, consumer perception of personalization may go well beyond the actual extent of price discrimination. For example, a consumer may mistake a surge in price due to market-level demand and supply for a personalized price, and hence take unnecessarily costly actions to reduce her fare.

⁴⁸Worldwide spending on online display ads totaled \$53 billion in 2018, i.e., 20% of all digital advertising revenues (Statista DMO 2019, <https://www.statista.com/outlook/216/100/digital-advertising/worldwide>).

⁴⁹Supplementary Appendix S.1 provides a comprehensive description of this market.

market design by technology platforms. [Lahaie, Munoz Medina, Sivan, and Vassilvitskii \(2018\)](#) develop tests based on bid perturbations to identify the relationship between past bids and future reserve prices. Similarly, the demand-side platform Criteo advertises its ability to reduce bids when (static) second-price auction mechanics are manipulated by the seller ([Abeille, Calauzènes, Karoui, Nedelec, and Perchet, 2018](#)).

All three examples are consistent with the role of information in enabling consumers’ strategic behavior: buyers are aware of the potential for dynamic price discrimination, and they do attempt to manipulate prices. Sellers are also wary of the ratchet effect, although they differ in their approach to mitigating its consequences. Consistent with our model, some sellers view the value of information as sufficiently high to encourage dynamic personalized pricing, while others prefer to limit the value of strategic behavior. For example, the Rubicon Project (a supply-side platform that places almost \$1 billion in advertising spending per year) has recently stopped allowing reserve prices to adjust based on past behavior, specifically citing concerns over buyer manipulation ([The Rubicon Project, 2014](#)).

At the same time, these examples fail to capture some important features of our model. Most notably, consumer scores of the kind just described do not guide the prices set by different sellers. Instead, a single seller interacts with the buyer over time. While this makes it more likely that consumers are aware of behavior-based price discrimination, it also limits the likelihood that purchase histories are aggregated into scores of arbitrary persistence. (In terms of our model, for example, Uber and Google have access to the full history of signals.) Furthermore, repeated interaction with the same seller introduces the potential for non-Markov equilibria and seller reputation effects, which can generate very different dynamics along the path of play. The sellers in our three examples express concerns over buyers’ responses that are consistent with both types of equilibria.⁵⁰

8.2 The Role of Score Transparency

“Consumer scores stand today where credit scores stood in the 1950s: in the shadows” ([Dixon and Gellman, 2014](#)). Not surprisingly, consumers’ experience with credit scores in various countries helps us cast our results on transparency in a concrete setting.⁵¹

⁵⁰The marketing agency ActiveCampaign warns firms, “If you regularly use coupons and discounts to move merchandise, [...] people wait for sales before they buy” (<https://www.activecampaign.com/blog/abandoned-cart-coupon/>). Similarly, [The Rubicon Project \(2014\)](#) claims, “Using buy-side dimensions when setting price floors results in buyer distrust [and] changes in bidding behavior.”

⁵¹Transparency policies have proven effective in other settings, such as in the analysis of overdraft fees ([Grubb, 2015](#)). Although data brokers have made few attempts at improving transparency, recent state laws in California and Vermont, and proposed legislation in New York attempt to hold brokers more accountable.

In the U.S., FICO credit scores are used to determine the borrowing conditions (interest rates, maximum amounts) for consumers. It is well known that consumers can manipulate such scores to improve their borrowing opportunities. Indeed, countless websites explain how to improve scores (e.g., do not apply for credit cards, avoid bad credit, consolidate debt, etc.) and also offer score simulations to check the consequences of a hypothetical default on a bill payment. Most relevant to our results, the FICO Score Open Access program is designed such that the score the customer sees exactly matches a score version the lender uses within their risk management decisions.⁵² Credit bureaus advertise the transparency of their products as enabling consumers to “stop overpaying for credit.”

A similar message applies to consumer scores used to set prices. For example, consider the Equifax Discretionary Spending Index (DSI)—a hidden score from 1 to 1000, whose objective is to “Better understand how much customers can spend [in order to] create targeted promotions that appeal to them.”⁵³ [Ezrachi and Stucke \(2016\)](#) suggest that the inability to access scores such as this one in real time prevents consumers from “assessing their outside options” correctly, and leads them to accept higher prices. This mechanism is quite close to the one at work in our model with hidden scores: the forces that make a strategic consumer more price sensitive cannot operate unless the consumer can correctly estimate the quantity she will demand in the future, and hence the value of manipulating her score.⁵⁴

8.3 Future Trends

The key features of our model—that firms use behavior-based data to guide pricing, and that scores compress such data into low-dimensional statistics that exhibit persistence—are becoming increasingly relevant in a number of consumer markets. In particular, there is growing evidence of discriminatory pricing online, more so than when the practice was first brought to the public’s attention ([Dixon and Gellman, 2014](#); [Council of Economic Advisers, 2015](#)). A recent petition filed to the Federal Trade Commission lists retail and travel websites as sectors where prices and products are often personalized. Although the price differences across individuals are not huge, the business volume of these sellers is such that “it is likely that they are inflicting a *relatively small harm on a large number of consumers* under the FTC’s standard” (Section III. in [Represent Consumers \(2019\)](#)).

In addition to consumer retail, an area of increasing public concern in the US is health insurance. As documented by [Allen \(2018\)](#), insurance companies are using data from various

⁵²See <https://www.fico.com/en/products/fico-score-open-access/>.

⁵³See <https://www.equifax.com/business/discretionary-spending-index/>.

⁵⁴The language of outside options suggests a search model with unit demand. Although our consumer’s utility is additively separable in the quantities purchased at different times, the link between current consumption and future prices makes purchases at different times substitutes in our model.

sources to learn more about their clients and map them into scores (“social determinants of health”); further, experts say that these practices are entirely legal, as the Health Insurance Portability and Accountability Act (HIPPA) restricts only the use of medical information for insurance pricing. As awareness of the linkages between consumer behavior and the information collected by insurance companies is likely to increase over time, we expect the use of consumer metrics for pricing policies to give rise to similar tradeoffs as our model.

The use of score-based price discrimination is also bound to increase dramatically in the market for consumer credit worldwide. This expansion has been driven mostly by the creation of new credit scores in China. In particular, Ant Financial (a division of Alibaba) produces the Zhima (Sesame) Credit score on the basis of several dimensions in addition to the consumer’s financial history. These dimensions include mobile wallet transactions on Alipay, online browsing and shopping behavior, personal characteristics, and even interpersonal relationships on social media.⁵⁵ Tencent developed a similar score with data from WeChat Pay (Wired, 2018).

In addition to these new data sources, what brings these scores closer to our model is that a wide variety of vendors and lenders can access them on the Alipay and Tencent platforms: de la Mano and Padilla (2019) report that “Ant Financial’s Zhao Cai Bao marketplace allows third-party financial institutions to offer loans to small and medium enterprises, and individuals [...]” As further evidence of the use of scores by third parties (i.e., not just Ant Financial), consider that Ant Credit Pay (Alipay’s consumer credit product) has 100 million active users, but Zhima Credit Scores cover more than 300 million Chinese consumers. Finally, both the Alibaba and Tencent scores inform the Chinese government’s Social Credit System. This creates an interesting natural experiment on transparency, whereby the government score is hidden from the public, but the credit scores it is based on can be observed by individual consumers.

9 Conclusions

We have explored the informational and welfare consequences of using purchase histories to price discriminate, with special emphasis on the use of scores and their transparency, a topic of central importance for recent regulatory efforts aimed at protecting consumers. Critically, we have uncovered an *equilibrium* mechanism by which transparency can help consumers.

Our model makes a number of simplifying assumptions, the strongest of which is perhaps the restriction to a continuous score with exponential weights. From an applied standpoint, moreover, consumer scores used to set prices also have other effects. In the case of insurance,

⁵⁵See <https://www.alibabagroup.com/en/news/article?news=p150128>.

a consumer attempting to manage her score might adopt better driving habits. In the case of marketing, scores are partly based on the idea of excluding “bad customers” and targeting the firm’s effort in sales or service. Furthermore, scores clearly have other uses beyond price discrimination. For example, [Brayne \(2017\)](#) describes the role of risk and merit scores in driving law enforcement’s “stratified surveillance” practices. Thus, our model is just a natural benchmark for starting an informed debate on the economics of consumer scores.

Appendix: Proofs

Proofs for Section 3

Proof of Proposition 1. Parts 1 and 2 follow from the expressions for consumer and producer surplus. Part 3 follows directly from the weight on θ_t not responding to the firms’ information structure. Finally, 4 is a special case of Proposition 3. \square

Proofs for Section 4

The statements in Section 4.1 follow from the next

Lemma A.1 (Stationarity and Beliefs). *A process $(\theta_t, Y_t)_{t \geq 0}$ with $(Q_t)_{t \geq 0}$ as in (6) and $M_t = \mathbb{E}[\theta_t | Y_t]$ for all $t \geq 0$ is stationary Gaussian if and only if:*

- (i) $M_t = \mu + \lambda[Y_t - \bar{Y}]$, with $\bar{Y} = \mu(\alpha + \beta + \delta)/\phi$ and $\lambda = \frac{\alpha\sigma_\theta^2(\phi - \beta\lambda)}{\alpha^2\sigma_\theta^2 + \sigma_\xi^2\kappa(\phi - \beta\lambda + \kappa)}$;
- (ii) the score process (4) is mean reverting: $\phi - \beta\lambda > 0$;
- (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.2).

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

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

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

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

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

$$X_t = e^{-A_1 t} X_0 + \int_0^t e^{-A_1(t-s)} A_0 ds + \int_0^t e^{-A_1(t-s)} B dZ_s, \quad t \geq 0, \quad (\text{A.1})$$

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

From the additive structure of (A.1), 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 .⁵⁶ 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 the 2×2 covariance matrix Γ must satisfy the equations

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

where $\text{Var}[\cdot]$ and T denote the covariance matrix and transpose operators, respectively.

Observe that the first condition leads to $\vec{\mu} = A_1^{-1} A_0$ provided A_1 is invertible. This, in turn, 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

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

To guarantee that the previous expressions indeed correspond to the first two moments of stationary Gaussian process, however, we must verify that Γ is both positive semi-definite and finite. Since $\sigma_\theta^2/2\kappa > 0$, positive semi-definiteness reduces to

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

⁵⁶Denote 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.

Because $\phi - \beta\lambda \neq 0$, however, it follows that $\phi - \beta\lambda$ must be strictly positive. As a byproduct, $\Gamma_{22} > 0$ is finite. This proves (ii) and (iii).

To finish the proof, we find λ and \bar{Y} that are consistent with Bayes' rule given a score process that is driven by (6). Using (A.2),

$$\lambda = \frac{\text{Cov}[\theta_t, Y_t]}{\text{Var}[Y_t]} = \frac{\alpha\sigma_\theta^2(\phi - \beta\lambda)}{\alpha^2\sigma_\theta^2 + \kappa\sigma_\xi^2(\phi + \kappa - \beta\lambda)} \quad \text{and} \quad \bar{Y} = \frac{\mu[\alpha + \beta + \delta]}{\phi}$$

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. \square

Proof of Lemma 1. It follows from partially differentiating (7) with respect to ϕ . \square

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 Lemma A.1. Thus, we can write $Q(p, \theta_t, M_t) = q_0 + \alpha\theta_t + q_2M_t - p$ 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.

It is then direct from here that the equilibrium price and quantity are given by

$$P_t = \frac{q_0}{2} + \frac{\alpha + q_2}{2}M_t \quad \text{and} \quad Q_t = \frac{q_0}{2} + \alpha\theta_t + \frac{q_2 - \alpha}{2}M_t,$$

and so 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 $Q(p, \theta_t, Y_t) = 2\delta\mu + [\mu - \lambda\bar{Y}][\alpha + 2\beta] + \alpha\theta_t + \lambda[\alpha + 2\beta]Y_t - p$. This concludes the proof. \square .

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 $V = v_0 + v_1\theta + v_2M + v_3M^2 + v_4\theta^2 + v_5\theta M$ to the HJB equation

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

subject to standard transversality conditions. To find a stationary linear Markov equilibrium, however, (i) we impose the fixed-point condition that the optimal policy is of the form $\delta\mu + \alpha\theta + \beta M$, and (ii) with the use of the equation for λ (equation (7)), find coefficients that satisfy the stationarity condition $\phi - \beta\lambda > 0$ (part (ii) in Lemma A.1).

To this end, observe that the first-order condition of the HJB equation reads

$$\begin{aligned} q &= \theta - [\delta\mu + (\alpha + \beta)M] + \lambda[v_2 + 2v_3M + v_5\theta] \\ &= -\delta\mu + \lambda v_2 + [1 + \lambda v_5]\theta + [2\lambda v_3 - (\alpha + \beta)]M \end{aligned}$$

which leads to the following system matching-coefficient conditions:

$$\delta\mu = -\delta\mu + \lambda v_2, \quad \alpha = 1 + \lambda v_5, \quad \text{and} \quad \beta = 2\lambda v_3 - (\alpha + \beta). \quad (\text{A.3})$$

By the Envelope Theorem, moreover,

$$\begin{aligned} (r + \phi)[v_2 + 2v_3M + v_5\theta] &= -(\alpha + \beta)[\delta\mu + \alpha\theta + \beta M] - \kappa(\theta - \mu)v_5 \\ &\quad + [\lambda(\delta\mu + \alpha\theta + \beta M) - \phi(M - \mu + \lambda\bar{Y})]2v_3, \end{aligned} \quad (\text{A.4})$$

which yields the following system of equations

$$\begin{cases} (r + \phi)v_2 = -(\alpha + \beta)\delta\mu + \kappa\mu v_5 + [\lambda\delta\mu + \phi(\mu - \lambda\bar{Y})]2v_3 \\ (r + 2\phi)2v_3 = -(\alpha + \beta)\beta + 2v_3\lambda\beta \\ (r + \kappa + \phi)v_5 = -(\alpha + \beta)\alpha + 2v_3\lambda\alpha. \end{cases} \quad (\text{A.5})$$

Using that v_2, v_3 and v_5 can be written as a function of α, β and $\delta\mu$, (A.5) becomes

$$\begin{cases} (r + \phi)\frac{2\delta\mu}{\lambda} = -(\alpha + \beta)\delta\mu + \kappa\mu\frac{\alpha-1}{\lambda} + [\lambda\delta\mu + \phi(\mu - \lambda\bar{Y})]\frac{\alpha+2\beta}{\lambda} \\ (r + 2\phi)\frac{\alpha+2\beta}{\lambda} = \underbrace{-(\alpha + \beta)\beta + \beta(\alpha + 2\beta)}_{=(\beta)^2} \\ (r + \kappa + \phi)\frac{\alpha-1}{\lambda} = \underbrace{-(\alpha + \beta)\alpha + \alpha(\alpha + 2\beta)}_{=\alpha\beta}. \end{cases} \quad (\text{A.6})$$

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 A.2. *Any stationary linear Markov equilibrium must satisfy $\alpha \in (0, 1)$.*

Proof. Consider a stationary linear Markov equilibrium with coefficients (α, β, δ) . Straight-

forward integration shows that the consumer's equilibrium payoff is quadratic, and thus the system of equations (A.6) holds.

Suppose that $\alpha = 0$. From (7), $\lambda = 0$, and so $M_t = \mu$ for all $t \geq 0$; but this implies that prices are constant, and hence, it is optimal for the consumer to behave myopically by choosing $Q_t = \theta_t - p$, a contradiction. If instead $\alpha < 0$, the last equation in (A.6) 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 Lemma A.1).

The case $\alpha = 1$ can be easily ruled out too: since $\lambda > 0$ in this case, the last equation in the system (A.6) 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.6) can be used to solve for β and thus to find an expression for λ as a function of ϕ , α , and the parameters r and κ . In addition, from the last equation in (A.6),

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

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, $L > 0$. Since $\alpha > 1$, however, this implies that the left-hand side of this expression is strictly positive, which is a contradiction. Thus, $\alpha \in (0, 1)$. \square

We continue with the proof of the proposition. From the proof of the previous lemma, $\beta \neq 0$. In the system (A.6), 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

$$B(\phi, x) := \frac{-x^2(r + 2\phi)}{2(r + 2\phi)x - (r + \kappa + \phi)(x - 1)} \in \left(-\frac{x}{2}, 0 \right) \quad \text{when } x \in (0, 1). \quad (\text{A.7})$$

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

$$\Lambda(\phi, \alpha, \beta) := \frac{\sigma_\theta^2 \alpha(\alpha + \beta) + \kappa \sigma_\xi^2(\kappa + \phi) - \sqrt{[\sigma_\theta^2 \alpha(\alpha + \beta) + \kappa \sigma_\xi^2(\kappa + \phi)]^2 - 4\kappa(\sigma_\theta \sigma_\xi)^2 \alpha \beta \phi}}{2\beta \kappa \sigma_\xi^2}. \quad (\text{A.8})$$

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.6), we arrive to equation (11): namely, $\alpha \in (0, 1)$ must satisfy $A(\phi, \alpha) = 0$, where

$$A(\phi, x) := (r + \kappa + \phi)(x - 1) - \Lambda(\phi, x, B(\phi, x))xB(\phi, x), \quad x \in [0, 1]. \quad (\text{A.9})$$

We now establish the existence and uniqueness of a solution to this equation, along with regularity properties with respect to $\phi > 0$ and $\sigma_\xi^2 > 0$ that will be used later on.

Lemma A.3. *For every $\phi > 0$ and $\sigma_\xi^2 > 0$, there exists a unique $\alpha \in (0, 1)$ satisfying the previous equation. Moreover, the induced function $\alpha : (0, \infty)^2 \rightarrow (0, 1)$ is of class C^1 .*

Proof: Fix $\phi > 0$ and $\sigma_\xi^2 > 0$; the dependence of A on $\sigma_\xi^2 > 0$ is via Λ , and we omit it until needed. Observe that as $x \rightarrow 1$, $B(\phi, x) \rightarrow -1/2$ and $\lim_{x \rightarrow 1} \Lambda(\phi, x, B(\phi, x)) > 0$. Hence, $\lim_{x \rightarrow 1} A(\phi, x) > 0$. Similarly, as $x \rightarrow 0$, $B(\phi, x) \rightarrow 0$ and $\lim_{x \rightarrow 0} \Lambda(\phi, x, B(\phi, x)) \rightarrow 0$. Hence, $\lim_{x \rightarrow 0} A(\phi, x) < 0$. The existence of $\alpha \in (0, 1)$ satisfying $A(\phi, \alpha) = 0$ follows from the continuity of $x \in [0, 1] \mapsto A(\phi, x)$ 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

$$H(\phi, x) := -\Lambda(\phi, x, B(\phi, x))B(\phi, x) > 0, \quad x \text{ in } (0, 1), \quad (\text{A.10})$$

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

From the previous limits, $\lim_{x \rightarrow 0} H(\phi, x) = 0$ and $\lim_{x \rightarrow 1} H(\phi, x) > 0$; 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 . Also, let $\ell(\phi, x) := \sigma_\theta^2 x(\alpha + B(\phi, x)) + \kappa \sigma_\xi^2(\kappa + \phi)$. At any such \hat{x} ,

$$\ell_x(\phi, \hat{x}) \underbrace{[\ell(\phi, \hat{x}) - (\ell^2(\phi, \hat{x}) - 4\kappa \sigma_\xi^2 \sigma_\theta^2 B(\phi, \hat{x}) \hat{x} \phi)^{1/2}]}_{< 0, \text{ as } B < 0} = 2\kappa(\sigma_\theta \sigma_\xi)^2 [B_x(\phi, \hat{x}) \hat{x} + B(\phi, \hat{x})] \phi. \quad (\text{A.11})$$

Moreover, straightforward algebra shows that

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

so $B_x(\phi, x)x + B(\phi, x) < 0$, $x \in [0, 1]$. Thus, $\ell_x(\phi, \hat{x}) = \sigma_\theta^2[2x + B_x(\phi, \hat{x})\hat{x} + B(\phi, \hat{x})] > 0$; otherwise the left-hand side of (A.11) is positive, while the other 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.11). 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

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

However, 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 $(0, 1) \times (0, \infty)^2 \mapsto A(x, \phi, \sigma_\xi^2)$ is of class C^1 and $\partial A/\partial x > 0$, our existence result allows us to apply the Implicit Function Theorem: namely, around any point $(\phi, \sigma_\xi^2) \in (0, +\infty)^2$ there exists a unique function, $(\phi, \sigma_\xi^2) \in (0, \infty)^2 \mapsto \alpha(\phi, \sigma_\xi^2) \in (0, 1)$ satisfying the equation, and such function is of class C^1 . However, since we already established that existence and uniqueness holds over the whole domain $(0, \infty)$, the local property of continuous differentiability trivially extends globally. This concludes the proof of the lemma. \square

It remains to characterize δ . Recall that the first equation in (A.6) 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

$$D(\phi, x) := \frac{\kappa(\alpha - 1) + [\alpha + 2B(\phi, \alpha)][\phi - (\alpha + B(\phi, \alpha))\Lambda(\phi, \alpha, B(\phi, \alpha))]}{2(r + \phi) + (\alpha + B(\phi, \alpha))\Lambda(\phi, \alpha, B(\phi, \alpha))}, \quad (\text{A.12})$$

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

To conclude the proof, there are two final steps: determining the rest of the coefficients and checking transversality conditions and the admissibility of the candidate equilibrium strategy. Both are verified in section S.2.1 in the Supplementary Appendix. \square

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

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

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 $(\theta_t^\vartheta, M_t^m)_{t \geq 0}$ is the type-belief process starting from $(\theta_0, M_0) = (\vartheta, m) \in \mathbb{R}^2$.

From the proof of Proposition 1, the previous equation admits as solution the function $V_M(\theta, M) = v_2 + 2v_3M + v_5\theta$ 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.4) shows that the previous function satisfies the PDE, while the transversality condition follows directly from $(\theta_t^\vartheta)_{t \geq 0}$ and $(M_t^m)_{t \geq 0}$ being mean reverting 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 (Remark 3.5.6 in Pham 2009) applies: namely, $V_M(\vartheta, m) = -\mathbb{E}_0 \left[\int_0^\infty e^{-(r+\phi)t} (\alpha + \beta) Q_t dt \right]$, $\forall t \geq 0$, where we used that $Q_t = \delta\mu + \alpha\theta_t + \beta M_t$ in equilibrium. The result then follows from $V_M(\theta_t, M_t) = -\mathbb{E}_0 \left[\int_0^\infty e^{-(r+\phi)t} (\alpha + \beta) Q_t dt \right] = -\mathbb{E}_t \left[\int_t^\infty e^{-(r+\phi)(s-t)} (\alpha + \beta) Q_s ds \right]$ if $(\theta_t, M_t) = (\theta_0, M_0) = (\vartheta, m) \in \mathbb{R}^2$.

Part (ii) is proved as part of the next auxiliary lemma. This concludes the proof. \square

Lemma A.4 (Equilibrium Properties).

(i) *Uninformative scores:* $\lim_{\phi \rightarrow 0, \infty} (\alpha(\phi), \beta(\phi), \delta(\phi)) = (1, -1/2, 0)$, $\lim_{\phi \rightarrow 0} \lambda(\phi) = 0$, $\lim_{\phi \rightarrow \infty} \lambda(\phi) = \sigma_\theta^2 / \kappa \sigma_\xi^2$, $\lim_{\phi \rightarrow 0} \lambda(\phi) / \phi = 2\sigma_\theta^2 / [\sigma_\theta^2 + 2\sigma_\xi^2 \kappa^2]$ and $\lim_{\phi \rightarrow 0, \infty} \mathbb{E}[(P_t - \mu/2)^2] = 0$.

(ii) *Bounds on the strength of ratchet effect:* for all $\phi > 0$,

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

(iii) *Strategic demand reduction across types:* $\alpha(\cdot)$ is strictly quasiconvex.

(iv) *Effect of noise:* $\alpha(\phi)$ and $\mathbb{E}[P_t]$ are increasing in $\sigma_\xi / \sigma_\theta$ for all $\phi > 0$.

Proof of Lemma A.4. (i) Limits. 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 (11) (or, equivalently, $A(\phi, \alpha(\phi)) = 0$, where $A(\phi, x)$ is defined in (A.9)), 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 $0 < -4\kappa(\sigma_\xi \sigma_\theta)^2 \beta(\phi) \alpha(\phi) \phi \rightarrow 0$ as $\phi \rightarrow 0$. In addition, because $\alpha(\phi) + \beta(\phi) > 0$, it follows that $\ell(\phi, \alpha) > \kappa^2 \sigma_\xi^2$. Using that $\beta(\phi) = B(\phi, \alpha(\phi))$ then 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 $\lim_{\phi \rightarrow 0} \alpha(\phi)$ 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, $J(\cdot)$ is bounded too. Thus, from $A(\phi, \alpha(\phi)) = 0$, we have $1 - \alpha(\phi) = [\alpha(\phi)J(\phi)/2\kappa\sigma_\xi^2](r + \kappa + \phi)^{-1}$. Because $[\alpha(\phi)J(\phi)/2\kappa\sigma_\xi^2]$ is bounded, it follows that $1 - \alpha(\phi) \rightarrow 0$ as $\phi \rightarrow \infty$.

Regarding the limit values for $\beta(\phi) = B(\phi, \alpha(\phi))$, these follow from the limit behavior of $\alpha(\phi)$ and (A.7). To study the limit behavior of $\delta(\phi)$, we first examine $\lambda(\phi)$.

To this end, $\lim_{\phi \rightarrow 0} \lambda(\phi) = 0$ is direct consequence of the first bound in (A.15) which we establish shortly in the proof of part (ii) of this lemma. Also, letting $\ell(\phi, \alpha) := \sigma_\theta^2 [\alpha +$

$B(\phi, \alpha)] + \kappa\sigma_\xi^2[\phi + \kappa]$, it is straightforward to verify that

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

as $\phi \rightarrow \infty$. Thus, the second limit holds. Finally, it is easy to see that the limit of $\lambda(\phi)/\phi$ as $\phi \searrow 0$ follows directly from the first equality in the previous display.

Returning to $\delta(\phi)$, recall from (A.12) that

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

Using that $\alpha(\phi) \rightarrow 1$, $\lambda(\phi) \rightarrow 0$, and $\alpha(\phi) + 2\beta(\phi) \rightarrow 0$ as $\phi \rightarrow 0$, and that $\alpha(\phi) + \beta(\phi) > 0$, it is direct that $\delta(\phi) \rightarrow 0$ as $\phi \rightarrow 0$. Also, using that $\lambda(\phi) \rightarrow \sigma_\theta^2/\kappa\sigma_\xi^2$ as $\phi \rightarrow \infty$, $[\phi - (\alpha(\phi) + \beta(\phi))\lambda(\phi)]/[2(r + \phi) + (\alpha(\phi) + \beta(\phi))\lambda(\phi)] \rightarrow 1/2$ as $\phi \rightarrow \infty$. The limit $\delta(\phi) \rightarrow 0$ as $\phi \rightarrow \infty$ then follows from $\alpha(\phi) \rightarrow 1$ and $\alpha(\phi) + 2\beta(\phi) \rightarrow 0$ as $\phi \rightarrow \infty$.

To establish the convergence of prices, we first show that $\lim_{\phi \rightarrow 0, +\infty} \text{Var}[\lambda(\phi)Y_t] = 0$. In fact, recall that $(\delta(\phi), \alpha(\phi), \beta(\phi)) \rightarrow (0, 1, -1/2)$ as $\phi \rightarrow 0$ and $+\infty$. Also, from (A.2),

$$\text{Var}[Y_t] = \frac{1}{2(\phi - \beta(\phi)\lambda(\phi))} \left[\sigma_\xi^2 + \frac{\alpha(\phi)\sigma_\theta^2}{\kappa[\phi - \beta(\phi)\lambda(\phi) + \kappa]} \right]. \quad (\text{A.13})$$

We deduce that $0 = \lim_{\phi \rightarrow \infty} \text{Var}[Y_t] = \lim_{\phi \rightarrow \infty} \text{Var}[\lambda(\phi)Y_t]$. Now, writing (A.13) as

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

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

(ii) Bounds. Observe that the bounds for $\beta(\phi)$ were already determined from (A.7) and $\alpha(\phi) \in (0, 1)$. As for the lower bound for α , we will show the stronger result

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

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

To obtain the bound, observe that from (A.8), $\lambda(\phi)$ satisfies

$$\lambda(\phi) = \frac{2\sigma_\theta^2\alpha(\phi)\phi}{\sqrt{\ell^2(\phi, \alpha(\phi)) - 4\kappa(\sigma_\theta\sigma_\xi)^2\alpha(\phi)B(\phi, \alpha(\phi))\phi} + \ell(\phi, \alpha(\phi))} < \frac{\sigma_\theta^2\alpha(\phi)\phi}{\ell(\phi, \alpha(\phi))}, \quad (\text{A.14})$$

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:

$$\lambda < \frac{2\phi}{\alpha(\phi)} \quad \text{and} \quad \lambda < \frac{\sigma_\theta^2}{\kappa\sigma_\xi^2}. \quad (\text{A.15})$$

In fact, since $B(\phi, \alpha(\phi)) \geq -\alpha/2$, then $\ell(\phi, \alpha(\phi)) > \sigma_\theta^2\alpha(\phi)^2/2$; but using this in (A.14) leads to the first inequality in (A.15). Similarly, the second upper bound follows from (A.14) 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}_+ , and it converges to zero as $\phi \rightarrow 0$, as promised.

Consider now the locus $A(\phi, \alpha(\phi)) = 0$. Using the first bound in (A.15) yields

$$\begin{aligned} 0 &= (r + \kappa + \phi)(\alpha(\phi) - 1) + \underbrace{\lambda(\phi)}_{\leq 2\phi/\alpha(\phi)} \underbrace{\alpha(\phi)[-B(\phi, \alpha(\phi))]}_{\in (0, \alpha(\phi)/2)} < (r + \kappa + \phi)(\alpha(\phi) - 1) + \phi\alpha(\phi) \\ \Rightarrow \alpha(\phi) &> \frac{r + \kappa + \phi}{r + \kappa + 2\phi} > \frac{1}{2}, \quad \text{for all } \phi > 0. \end{aligned}$$

Similarly, using the second bound, $0 < (r + \kappa + \phi)(\alpha(\phi) - 1) + [\sigma_\theta\alpha(\phi)]^2/[2\kappa\sigma_\xi^2]$; the desired second bound for alpha follows from imposing that $\alpha^2 < \alpha$ in the previous inequality.

We conclude this part by establishing the bounds for the expected price, omitting the dependence of $(\alpha, \beta, \delta, \lambda)$ on ϕ . Observe that $\mathbb{E}[P_t] = [\delta + \alpha + \beta]\mu$. Now, adding the second and third equation in the system (A.6) yields $(\alpha + 2\beta)(\alpha + \beta)\lambda = (r + 2\phi)(\alpha + 2\beta) + (r + \kappa + \phi)(\alpha - 1) + (\alpha + \beta)^2\lambda$. Using (A.12), straightforward algebra then yields

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

from where it is easy to conclude that

$$\mathbb{E}[P_t] = \mu[\alpha + \beta + \delta] = \mu \frac{r + \phi}{2(r + \phi) + (\alpha + \beta)\lambda}. \quad (\text{A.16})$$

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.14) 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 latter bound in (A.16) leads to $\mathbb{E}[P_t] > \mu/3$ whenever $\mu \neq 0$, as $r > 0$.

(iii) Quasiconvexity of α . To prove this property, it is more useful to solve the last two equations in the system (A.6) for λ and β , namely,

$$\lambda(\phi, \alpha) = -\frac{(\alpha - 1)(\kappa + r + \phi)(\kappa + \alpha(-\kappa + r + 3\phi) + r + \phi)}{\alpha^3(r + 2\phi)}, \quad (\text{A.17})$$

$$\beta(\phi, \alpha) = -\frac{\alpha^2(r + 2\phi)}{\kappa + \alpha(-\kappa + r + 3\phi) + r + \phi}. \quad (\text{A.18})$$

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

$$\tilde{A}(\phi, \alpha) := \frac{(\alpha - 1)(\kappa + r + \phi)(\kappa + \alpha(-\kappa + r + 3\phi) + r + \phi)}{\alpha^3(r + 2\phi)} + \frac{\alpha(\kappa - \alpha(\kappa + r) + r + \phi)}{\alpha^3 + \kappa s(\kappa - \alpha r + r + \phi)} = 0. \quad (\text{A.19})$$

Observe that \tilde{A} is increasing in α whenever $\tilde{A}(\phi, \alpha) = 0$. 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

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

is strictly negative because, by inspection, the first term is nonpositive and the second term is strictly negative. Furthermore, $\phi \mapsto \alpha(\phi)$ is twice continuously differentiable.⁵⁷ 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 point-wise, consider again the locus $\tilde{A}(\phi, \alpha) = 0$ in (A.19), and differentiate with respect to $s := \sigma_\xi^2/\sigma_\theta^2$:

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

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.17)–(A.19), the derivative of the expected price

⁵⁷This follows from $\alpha'(\phi) = \frac{1 - \alpha(\phi) - \alpha(\phi)H_\phi(\phi, \alpha(\phi))}{r + \kappa + \phi + H(\phi, \alpha(\phi)) + \alpha(\phi)H_\alpha(\phi, \alpha(\phi))}$, with H as in (A.10) in Lemma A.3, and the right-hand side of the previous equality being continuously differentiable in ϕ .

$[\alpha + \beta + \delta]\mu$ with respect to α can be written as

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

Moreover, when using (A.17)–(A.19), the expected price does not depend on s directly. Thus, using that $\phi \mapsto \alpha(\phi)$ is increasing in σ_ξ/σ_θ , the expected price is also increasing in s . \square

Proofs for Section 5

The following lemmas (proved in Supplementary Appendix S.2.5) are used in this section.

Lemma A.5. *Suppose $\alpha > 0$ and $\beta < 0$ satisfy $\nu(\alpha, \beta) > 0$, where $\nu(\alpha, \beta)$ is defined in (18). Then, $\phi \mapsto G(\phi, \alpha, \beta)$ has a unique maximizer located at $\phi = \nu(\alpha, \beta)$. 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 derivative of $\Lambda(\phi, \alpha, \beta)$ 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 A.6. $\kappa < \arg \min \alpha < \infty$, and $[\alpha + \beta]'(\phi) < 0$, $\phi \in [\kappa, \arg \min \alpha]$; hence, $\alpha + \beta$ is strictly decreasing at any point satisfying (19). If $r > \kappa$, $[\alpha + \beta]'(\phi) < 0$ for $\phi \in [0, \arg \min \alpha]$.

Proof of Proposition 3. Refer to the Supplementary Appendix section S.2.2. \square

Proof of Proposition 4. We begin the proof by establishing (ii), i.e., that $\alpha'(\phi) < 0$ at any ϕ satisfying (19), 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 A.3 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,

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

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 on $(\phi, \alpha(\phi))$ of H , Δ , ℓ , B , and of their respective partial derivatives,

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

Moreover, since $\ell_\phi = \sigma_\theta^2\alpha B_\phi + \kappa\sigma_\xi^2$ we can write

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

Consider now the first term of the previous expression. In fact,

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

From Lemma A.5, moreover, $\Lambda_\phi(\nu(\alpha, \beta), \alpha, \beta) = \Lambda(\nu(\alpha, \beta), \alpha, \beta)/[\nu(\alpha, \beta) + \kappa]$; therefore, this equality must holds at any ϕ such that $(\phi, \alpha(\phi), \beta(\phi)) = (\nu(\alpha(\phi), \beta(\phi)), \alpha(\phi), \beta(\phi))$.

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

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

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

$$1 - \alpha - \alpha H_\phi = \underbrace{1 - \alpha + \frac{\lambda\alpha\beta}{\phi + \kappa}}_{K_1:=} - \underbrace{\frac{\sigma_\theta^2\alpha B_\phi}{\Delta} [(r + \kappa + \phi)(\alpha - 1) - \phi\alpha]}_{K_2:=}. \quad (\text{A.20})$$

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 (19), provided any such point exists.

We now turn to (i). 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., $\gamma(\alpha)$ is 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, using that $\lim_{\phi \rightarrow 0} (\alpha(\phi), \beta(\phi)) = (1, -1/2)$, we have that $\lim_{\phi \rightarrow 0} \eta(\phi) < 0$. The existence of ϕ s.t. $\eta(\phi) = 0$ follows from the continuity of $\eta(\cdot)$.

For uniqueness, observe first that $\alpha > 0$ and $\alpha + \beta > \alpha/2 > 0$ imply that $\nu(\alpha(\phi), \beta(\phi)) > \kappa$ for all $\phi > 0$. Also, since α is quasiconvex, $\alpha(\phi) \in (1/2, 1)$ if $\phi > 0$, and $\lim_{\phi \rightarrow 0} \alpha(\phi) = 1$, we have that α is decreasing in $[0, \arg \min \alpha]$, and non-decreasing thereafter. Since $\kappa < \arg \min \alpha$ (Lemma A.6) and α is strictly decreasing at any point satisfying $\phi = \nu(\alpha(\phi), \beta(\phi))$, we conclude that any such point must lie in $[\kappa, \arg \min \alpha]$.

Given this observation, it then suffices show that $[\nu(\alpha(\phi), \beta(\phi))]' < 0$ over $[\kappa, \arg \min \alpha]$. 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 $\phi = \nu(\alpha(\phi), \beta(\phi))$ and $[\nu(\alpha(\phi), \beta(\phi))]' > 0$, yielding a contradiction. To this end, write

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

From Lemma A.6, $\alpha(\phi) + \beta(\phi)$ is strictly decreasing over $[\kappa, \arg \min \alpha]$. Since $\alpha + \beta > 0$ and $\alpha\gamma(\alpha(\phi)) > 0$ it suffices to show that $[\alpha(\phi)\gamma(\alpha(\phi))]' < 0$ over the same region. However, $\frac{\alpha\gamma(\alpha)}{\sigma_\xi^2} = \frac{\sigma_\theta^2}{\sigma_\xi^2} \left[\left(\frac{\kappa^2}{\alpha} + \frac{\sigma_\theta^2}{\sigma_\xi^2} \right)^{1/2} + \frac{\kappa}{\alpha} \right]^{-1}$, which is strictly increasing in α . We conclude by using that $\alpha' < 0$ over $[\kappa, \arg \min \alpha]$.

Equipped with (i) and (ii), we now turn to (iii). Recall that $G(\phi) = \alpha(\phi)\lambda(\phi)/[\phi + \kappa - \beta(\phi)\lambda(\phi)] \geq 0$, where $\lambda(\cdot) = \Lambda(\cdot, \alpha(\cdot), \beta(\cdot))$. Since $\lambda(\phi)$ is bounded (second bound in (A.15)), $\lim_{\phi \rightarrow \infty} G(\phi) = 0$. Also $G(0) = 0$. By continuity, G has a global optimum that is interior.

From $\nu(\alpha, \beta)$, $G(\phi) := G(\phi, \alpha(\phi), \beta(\phi)) \leq G(\nu(\alpha(\phi), \beta(\phi)), \alpha(\phi), \beta(\phi))$, with equality only at ϕ^* . Also, from Lemma A.5, $\Lambda(\nu(\alpha, \beta), \alpha, \beta) = \alpha\gamma(\alpha)/\sigma_\xi^2$. Thus,

$$G(\nu(\phi), \alpha(\phi), \beta(\phi)) = \frac{\alpha(\phi)\Lambda(\nu(\phi), \alpha(\phi), \beta(\phi))}{\nu(\phi) + \kappa - \beta(\phi)\Lambda(\nu(\phi), \alpha(\phi), \beta(\phi))} = \frac{\alpha^2(\phi)\gamma(\alpha(\phi))}{\alpha^2(\phi)\gamma(\alpha(\phi)) + 2\kappa\sigma_\xi^2}, \quad (\text{A.21})$$

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

However, by definition of $\gamma(\alpha)$, $\alpha^2\gamma(\alpha) = \sigma_\xi^2[(\kappa^2 + \alpha^2(\phi)\sigma_\theta^2/\sigma_\xi^2)^{1/2} - \kappa]$; thus, from (A.21), $G(\nu(\phi), \alpha(\phi), \beta(\phi))$ 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 interval.

We now 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.17) and (A.18) to substitute for λ and β , we obtain that $G(\phi) = \tilde{G}(\phi, \alpha(\phi))$ where

$$\tilde{G}(\alpha, \phi) := (1 - \alpha) \frac{(\kappa + r + \phi)}{(\kappa + (1 - \alpha)r + \phi)} \frac{(\kappa + \alpha(-\kappa + r + 3\phi) + r + \phi)}{\alpha(r + 2\phi)}. \quad (\text{A.22})$$

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, $\partial\tilde{G}/\partial\phi < 0$. Also, up to a positive multiplicative term, $\frac{\partial\tilde{G}}{\partial\alpha} = -((\rho + f + 1) - \alpha\rho)^2 + \alpha^2(f + 1)(2\rho + 3f - 1)$, where $\rho := r/\kappa$ and $f := \phi/\kappa$. However, $-((\rho + f + 1) - \alpha\rho)^2 + \alpha^2(f + 1)(2\rho + 3f - 1) < -(f + 1)(\alpha^2(2\rho + 3f - 1) + f + 1)$ and the latter term is strictly negative for $\alpha \in [0, 1]$. The quasiconvexity of α then yields that $\frac{d\tilde{G}}{d\phi} = \frac{\partial\tilde{G}}{\partial\alpha}\alpha'(\phi) + \frac{\partial\tilde{G}}{\partial\phi} < 0$ for $\phi \in (\arg \min \alpha, \infty)$.

Part (iv) (i.e., G is decreasing in σ_ξ) follows immediately from the fact that \tilde{G} is decreasing in α for a fixed ϕ , because α is itself increasing in σ_ξ (Lemma A.4). \square

Proofs for Section 6

Proof of Proposition 5. We begin by showing the conditions for interior optima in (i) and (ii). From (A.16), $\alpha + \beta + \delta > 1/3$. Thus, omitting the dependence on ϕ ,

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

On the other hand, from the proof of Proposition 4, $\lim_{\phi \rightarrow 0, \infty} G(\phi) = 0$. Moreover, $\alpha + \beta$ is bounded. Therefore, $\lim_{\phi \rightarrow 0, \infty} (\alpha + \beta)^2 G(\phi) = 0$, and so $\lim_{\phi \rightarrow 0, \infty} \Pi(\phi) = \mu^2/4$.

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 $\Pi(\phi) > \mu^2/4$. Since $\phi \mapsto [\alpha(\phi) + \beta(\phi)]^2 G(\phi)$ is continuous, strictly positive, and converges to 0 as $\phi \rightarrow 0$ and $+\infty$, it has a global maximum; denote it by ϕ^\dagger . Thus, ϕ^f is interior if $\mu < \left[\frac{18\sigma_\theta^2}{5\kappa} (\alpha(\phi^\dagger) + \beta(\phi^\dagger))^2 G(\phi^\dagger) \right]^{1/2}$.

Now, let $CS_\mu(\phi)$ denote her surplus for a given μ and observe that $CS_\mu(\phi) = CS_0(\phi) +$

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

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

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

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

We now prove the corner solutions. Towards a contradiction, suppose that there are sequences $\mu_n \nearrow \infty$ and $\phi_n > 0$, $n \in \mathbb{N}$, 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 $[\alpha(\phi_{n_k}) + \beta(\phi_{n_k})]^2 G(\phi_{n_k})$ 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 $G = \alpha\lambda/[\phi + \kappa - \beta\lambda]$, we obtain

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

However, because α, β and λ are all bounded and $(\alpha, \beta, \lambda) \rightarrow (1, -1/2, \sigma_\theta^2/[\kappa\sigma_\xi^2])$ as $\phi \rightarrow +\infty$, 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 Lemma A.4, $\lim_{\phi \rightarrow 0} (\alpha(\phi), \beta(\phi), \lambda(\phi)) = (1, -1/2, 0)$, and so

$$\frac{4(\alpha + \beta)^2 \alpha \lambda [2(r + \phi) + \lambda(\alpha + \beta)]^2}{[4(r + \phi)(\alpha + \beta)\lambda + (\alpha + \beta)^2 \lambda^2](\phi + \kappa - \beta\lambda)} = \frac{4(\alpha + \beta)^2 \alpha [2(r + \phi) + \lambda(\alpha + \beta)]^2}{[4(r + \phi)(\alpha + \beta) + (\alpha + \beta)^2 \lambda](\phi + \kappa - \beta\lambda)}$$

converges to $2r/\kappa$ as $\phi \rightarrow \infty$ 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 and $L(\phi) := \frac{\alpha(\phi)^2}{2} + \beta(\phi) - \frac{3}{2}(\alpha(\phi) + \beta(\phi))^2 < 0$. 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 does not. Also, in the Supplementary Appendix (section S.2.3) we show that

$$\lim_{\phi \rightarrow 0, +\infty} \frac{[\alpha(\phi) - 1]^2}{R(\phi)} = 0, \quad \text{and} \quad \lim_{\phi \rightarrow 0, +\infty} \frac{G(\phi)}{R(\phi) - 1/8} > 0. \quad (\text{A.23})$$

But since $\lim_{\phi \rightarrow 0, +\infty} L(\phi) = -1/8$, the left-hand side of the same inequality is again violated.

Thus, there exists $\underline{\mu}^c > 0$ such that for all $\mu < \underline{\mu}^c$, $CS_{\mu}(0) > CS_{\mu}(\phi)$ for all $\phi \in (0, \infty)$.

For (iii), by pointwise convergence (section S.3.3 in the Supplementary appendix), we can directly consider the case of noiseless signals ($\sigma_{\xi} = 0$). In this case, it is possible to show that both surplus levels are decreasing as $\phi \rightarrow \infty$. Up to a positive multiplicative constant, the Mathematica file `scores.nb` shows that

$$\lim_{\phi \rightarrow \infty} \phi^2 \Pi'(\phi) = -\left(4\mu^2 r - 3\sqrt{3} + 6\right) - \frac{2r}{\kappa} < 0 \quad \text{and} \quad \lim_{\phi \rightarrow \infty} \phi^2 CS'(\phi) < 0.$$

On the other hand, $\phi^* \rightarrow \infty$ as $\sigma_{\xi}^2 \rightarrow 0$ (section S.3.2 in the Supplementary appendix). Therefore, expected profits in the noiseless case converge to their limiting value from above, which means that the firm-optimal score satisfies $\phi^f < \infty$ for all values of $\mu \geq 0$.

We conclude by addressing (iv). Differentiating CS and Π with respect to ϕ and setting equal to zero yields two loci $\mu^f(\phi)$ and $\mu^c(\phi)$ that describe the critical points of the surplus levels. The expressions for these loci are in the Mathematica file `scores.nb` posted on the authors' websites. In the Mathematica file, we also establish the following properties: (a) $\mu^f(\phi)$ is strictly quasiconcave, and $\mu^c(\phi)$ is strictly quasiconvex; (b) if $r/\kappa > \frac{1}{16}(\sqrt{337} - 7) \approx 0.71$, then $\mu^f(0) < 0$; and (c) if $\rho := r/\kappa$ satisfies $\rho(26 - \rho(24\rho + 31)) + 25 < 0$ (i.e., if $\rho < \bar{\rho} \approx 0.96$), then $\mu^c(0) > 0$. Therefore, when ρ lies in the (approximate) range $[0.71, 0.96]$, the expected surplus levels admit at most one critical point for each μ . Because both surplus levels are decreasing at $\phi = \infty$, this critical point is a local maximum. Furthermore, because

the inverses μ^f and μ^c are respectively quasiconcave and quasiconvex, as well as decreasing and increasing at $\phi = 0$, these loci are also monotone.

Finally, at $\phi = 0$, we have $\Pi'(0) \propto r/\kappa^2 - \mu^2$ and $CS'(0) \propto \mu^2 - 3r/\kappa^2$. Therefore, the conditions in part (iv) of the statement describe the values of μ for which the expected surplus levels admit an interior maximum. \square

Proof of Proposition 6. Fix $\phi > 0$. By the continuity of the equilibrium variables at $\sigma_\xi = 0$ (Supplementary Appendix section S.3.3), we can directly evaluate at $\sigma_\xi = 0$.

The expected price level with strategic consumers when $\sigma_\xi = 0$ is below the naive benchmark for all ϕ .⁵⁸ Thus, it suffices to show that consumer surplus in the strategic case exceeds the naive level when $\mu = 0$. We therefore compare the following expression in the two cases:

$$\text{Var}[\theta] \left(\alpha - \frac{\alpha^2}{2} \right) + \text{Var}[\theta] \left(\frac{\alpha^2}{2} - \frac{3}{2}(\alpha + \beta)^2 + \beta \right) \underbrace{\frac{\alpha\lambda}{-\beta\lambda + \phi + \kappa}}_{=G(\phi, \alpha, \beta)}.$$

In both cases, we let $\lambda = \phi/(\alpha + \beta)$. In the naive case, we further impose $\alpha = 1$ and $\beta = -1/2$, while in the strategic case β satisfies (A.18). Solving for s from (A.19), and imposing that the solution is positive, we obtain the equilibrium restriction $\alpha^2(\rho(2f - 1) + 3f^2 - 1) + \alpha(\rho + 2)(\rho + f + 1) - (\rho + f + 1)^2 > 0$, where $\rho := r/\kappa$ and $f := \phi/\kappa$. We then show (Mathematica file `scores.nb`) that for no triple $(\alpha, f, \rho) \in [1/2, 1] \times \mathbb{R}_+ \times \mathbb{R}_+$ such that the previous condition holds, consumer surplus with naive consumers exceed the one with strategic consumers. \square

Proofs for Section 7

Refer to section S.2.4 in the Supplementary Appendix.

References

- ABEILLE, M., C. CALAUZÈNES, N. E. KAROUI, T. NEDELEC, AND V. PERCHET (2018): “Explicit shading strategies for repeated truthful auctions,” *arXiv preprint*.
- ABEL, A., AND J. EBERLY (1994): “A Unified Model of Investment Under Uncertainty,” *The American Economic Review*, 84(5), 1369–1384.
- ACQUISTI, A., C. TAYLOR, AND L. WAGMAN (2016): “The Economics of Privacy,” *Journal of Economic Literature*, 54(2), 442–92.

⁵⁸ $\mathbb{E}[P_t] = \mu \frac{r+\phi}{2r+3\phi} \in (\mu/3, \mu/2)$; this follows from (A.16) using that $\lambda = \frac{\phi}{\alpha+\beta}$ when $\sigma_\xi = 0$.

- ALLEN, M. (2018): “Health Insurers are Vacuuming Up Details About You,” *Propublica*.
- AMIN, K., A. ROSTAMIZADEH, AND U. SYED (2013): “Learning prices for repeated auctions with strategic buyers,” in *Advances in Neural Information Processing Systems*.
- BERGEMANN, D., AND A. BONATTI (2019): “Markets for Information: An Introduction,” *Annual Review of Economics*, forthcoming.
- BERGER, P. D., AND N. I. NASR (1998): “Customer lifetime value: Marketing models and applications,” *Journal of interactive marketing*, 12(1), 17–30.
- BRAYNE, S. (2017): “Big Data Surveillance: The Case of Policing,” *American Sociological Review*, 82(5), 977–1008.
- CALZOLARI, G., AND A. PAVAN (2006): “On the optimality of privacy in sequential contracting,” *Journal of Economic Theory*, 130(1), 168–204.
- CARDELLA, E., AND B. DEPEW (2018): “Output restriction and the ratchet effect: Evidence from a real-effort work task,” *Games and Economic Behavior*, 107, 182–202.
- CHARNESS, G., P. KUHN, AND M. C. VILLEVAL (2011): “Competition and the Ratchet Effect,” *Journal of Labor Economics*, 29(3), 513–547.
- CISTERNAS, G. (2017): “Career Concerns and the Nature of Skills,” *American Economic Journal: Microeconomics*, forthcoming.
- (2018): “Two-Sided Learning and the Ratchet Principle,” *Review of Economic Studies*, 85(1), 307–351.
- CISTERNAS, G., AND A. KOLB (2019): “Signaling with Private Monitoring,” Discussion paper, MIT Sloan.
- COUNCIL OF ECONOMIC ADVISERS (2015): *Big Data and Differential Pricing*. Obama White House Archives.
- 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.
- DE LA MANO, M., AND J. PADILLA (2019): “Big Tech Banking,” *Journal of Competition Law & Economics*, forthcoming.

- 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.
- DUFFIE, D., P. DWORCZAK, AND H. ZHU (2017): “Benchmark in search markets,” *The Journal of Finance*, 72(5), 1983–2044.
- DWORCZAK, P. (2017): “Mechanism Design with Aftermarkets: Cutoff Mechanisms,” Discussion paper, Stanford University.
- DWYER, F. R. (1989): “Customer lifetime valuation to support marketing decision making,” *Journal of direct marketing*, 3(4), 8–15.
- EUROPEAN COMMISSION (2018): *Consumer market study on online market segmentation through personalised pricing/offers in the European Union*. Consumers, Health, Agriculture and Food Executive Agency.
- EZRACHI, A., AND M. E. STUCKE (2016): “The rise of behavioural discrimination,” Discussion paper, University of Oxford.
- 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.
- GRUBB, M. D. (2015): “Consumer inattention and bill-shock regulation,” *Review of Economic Studies*, 82(1), 219–257.
- HALAC, M. (2012): “Relational Contracts and the Value of Relationships,” *American Economic Review*, 102(2), 750–779.

- HANNAK, A., G. SOELLER, D. LAZER, A. MISLOVE, AND C. WILSON (2014): “Measuring price discrimination and steering on e-commerce web sites,” in *Proceedings of the 2014 conference on internet measurement conference*, pp. 305–318. ACM.
- HOLMSTRÖM, B. (1999): “Managerial Incentive Problems: A Dynamic Perspective,” *Review of Economic Studies*, 66(1), 169–182.
- HÖRNER, J., AND N. LAMBERT (2019): “Motivational Ratings,” *Review of Economic Studies*, forthcoming.
- HUMMEL, P. (2018): “Reserve prices in repeated auctions,” *International Journal of Game Theory*, 47(1), 273–299.
- ICHIHASHI, S. (2018): “Online Privacy and Information Disclosure by Consumers,” Discussion paper, Stanford University.
- KANORIA, Y., AND H. NAZERZADEH (2017): “Dynamic reserve prices for repeated auctions: Learning from bids,” Discussion paper, University of Southern California.
- KEHOE, P. J., B. J. LARSEN, AND E. PASTORINO (2018): “Dynamic Competition in the Era of Big Data,” Discussion paper, Stanford University.
- LAFFONT, J.-J., AND J. TIROLE (1988): “The Dynamics of Incentive Contracts,” *Econometrica*, 56(5), 1153–1175.
- LAHAIE, S., A. MUNOZ MEDINA, B. SIVAN, AND S. VASSILVITSKII (2018): “Testing Incentive Compatibility in Display Ad Auctions,” in *Proceedings of the 2018 World Wide Web Conference on World Wide Web*, pp. 1419–1428.
- LEWIS, M. (2005): “Incorporating strategic consumer behavior into customer valuation,” *Journal of Marketing*, 69(4), 230–238.
- 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.
- REPRESENT CONSUMERS (2019): “Secret Surveillance Scoring: Urgent Request for Investigation and Enforcement Action,” Discussion paper.
- ROBINSON, J. (1933): *The Economics of Imperfect Competition*. London: Macmillan.

- SCHMALENSEE, R. (1981): “Output and welfare implications of monopolistic third-degree price discrimination,” *American Economic Review*, 71(1), 242–247.
- SHEN, Q., AND J. M. VILLAS-BOAS (2017): “Behavior-Based Advertising,” *Management Science*, forthcoming.
- STRULOVICI, B. (2011): “Renegotiation-Proof Contracts with Moral Hazard and Persistent Private Information,” Discussion paper, Northwestern University.
- TAYLOR, C. R. (2004): “Consumer privacy and the market for customer information,” *Rand Journal of Economics*, 35(4), 631–650.
- THE ECONOMIST (2012): “Very personal finance,” *July 2*.
- THE GUARDIAN (2018): “Is your friend getting a cheaper Uber fare than you are?,” .
- THE RUBICON PROJECT (2014): “Maintaining the Equilibrium: How Dynamic Price Floors Preserve the Integrity of the Automated Advertising Ecosystem,” Discussion paper.
- THE WALL STREET JOURNAL (2018): “On Hold for 45 Minutes? It Might Be Your Secret Customer Score,” *November 1*.
- VILLAS-BOAS, J. (1999): “Dynamic competition with customer recognition,” *Rand Journal of Economics*, 30(4), 604–631.
- WIRED (2018): “Inside China’s Vast New Experiment in Social Ranking,” *January*.