Ratings design and Barriers to Entry

Nikhil Vellodi*

[Current Draft]: March 3, 2026

Abstract

I study the impact of consumer reviews on the incentives for firms to participate in the market. I develop a model of endogenous entry in which firms produce goods of heterogeneous, unknown quality. Consumers engage in directed search, and prices are exogenously fixed. A rater uses consumer reviews to construct firm-specific ratings that are informative of quality and help consumers direct their search. The main insight of the paper is that suppressing the reviews of highly-rated firms strengthens entry incentives and improves consumer welfare.

JEL Classification: D21, D82, D83, L11, L15, L86. **Keywords**: Product reviews, information design, firm dynamics, social learning, ergodic analysis, directed search.

1 Introduction

In this paper, I ask whether consumer reviews create barriers to entry for new firms. To answer this question, I develop a model in which firms of unknown quality make entry decisions. Quality is gradually revealed through consumer reviews. Reviews are used by a rater to construct firm-specific ratings. These ratings shape consumer demand. The main insight of the paper is that by stopping sufficiently well-established incumbents from acquiring reviews, the rater can strengthen incentives for new firms to enter and remain active, ultimately leading to gains in consumer welfare.

Product rating systems – raters that aggregate user-generated feedback to help inform consumer choice – are ubiquitous, playing a significant role in shaping choices and transforming the fortunes of all involved. Such raters provide an indispensable source of information, reducing search frictions and informational asymmetries and thereby allowing consumers and producers to engage in

*Paris School of Economics (PSE) and CNRS. I thank Dilip Abreu, Jarda Borovička, Hugo Hopenhayn, Erik Madsen, David Pearce, Steve Tadelis, Joshua Weiss, Jelena Zelenovic, and in particular Alessandro Lizzeri and Debraj Ray for their constant support and guidance, as well as participants at various seminars and conferences. Financial support from a NET Institute Grant and the EUR grant ANR-17-EURE-0001 is gratefully acknowledged. Part of this research was performed during stays at the William S. Dietrich II Economic Theory Center at Princeton University and the Simons Institute for the Theory of Computing at UC Berkeley; I thank both for their hospitality.

profitable trade.[1] Indeed, well-established firms and products often have many hundreds of reviews to their name, affording consumers unprecedented precision when making purchases. But while this stockpile of information might serve incumbent firms to great effect, it might work against a new entrant who unavoidably starts from scratch.[2]

This observation forms the starting point of my analysis — user-generated feedback creates a *barrier to entry* for firms, through the natural informational imbalance that exists between new entrants and incumbents. Such issues relate to the well-documented "cold start" phenomenon, a self-reinforcing link between a lack of demand and a lack of information regarding product quality (Che and Hörner, 2018; Kremer et al., 2014). These studies treat the range of available products as exogenously given, abstracting from producer participation constraints. In contrast, my analysis captures a novel equilibrium feedback channel through which cold-starting endogenously determines the distribution of products on offer, which in turn determines the relative demand and feedback rates for new products.

The model comprises consumers, firms, and a rater. Firms produce output of heterogeneous quality. Output quality is stochastic and depends on a firm's underlying type, which is either high or low. Each firm's type is independent from other firms', initially unknown to itself and all other market participants. Firms pay a fixed cost to enter and are subject to a service capacity constraint. Prices are fixed exogenously. A fixed measure of short-lived consumers choose between all available firms by engaging in directed search, subject to firms' capacity constraints and random rationing. Each consumer truthfully reports the realized output quality of the firm from whom they purchased to the rater, who uses these "reviews" to construct ratings, which the consumers use to direct their search.

I begin by fully characterizing the unique stationary equilibrium of the economy under a regime

[1]According to recent surveys, over 90% of consumers now consult online reviews before making purchase decisions (https://www.brightlocal.com/research/87-of-consumers-aged-18-34-have-reviewed-a-local-business-on-social-media). Displaying reviews can increase purchasing rates by 270% - see https://spiegel.medill.northwestern.edu/how-online-reviews-influence-sales. See Tadelis (2016) for an excellent survey of both supporting theories and evidence.

[2]Zomato, a restaurant rating service, argues that "the penalty from a bad review could have been a death sentence, especially for a new place…as a low rating may prevent new customers from visiting…" (https://www.zomato.com/blog/helping-new-restaurateurs-find-their-feet/). In a similar vein, a recent article describes a single bad review on TripAdvisor as *"…the marketing PR equivalent of a drive-by shooting"*, with the modern consumer labeled as *"…a veritable tyrant, with the power to make or break lives,"* (https://www.theguardian.com/news/2018/aug/17/how-tripadvisor-changed-travel).

of *full transparency*, in which the rater incorporates all available reviews into a firm's rating (Theorem 1). Equilibrium under this regime features rationing of consumers (Lemma 1) who trade off expected quality against the probability of service. This leads to sell-out revenues at only the best firms. The exacting nature of consumer search gives rise to a non-convexity in the firm's problem: its flow profits are discontinuous in its rating, and its continuation value is *convex-concave*. Importantly, from a firm's perspective, consumer feedback when poorly rated comes with *upside gain*; a standard option value effect emerges whereby positive reviews might bring sell-out flow profits. Conversely, for highly-rated firms, feedback entails *downside loss* as the firm experiences diminishing marginal returns from a higher rating due to capacity constraints. As such, struggling firms want rapid feedback, whereas successful firms want minimal feedback. Crucially, under full transparency, precisely the reverse profile obtains, as successful firms attract more customers and thus more reviews.

With these insights in hand, I turn to the optimal design of rating systems. I assume that the rater maximizes consumer welfare but is limited in its choice of instrument: it must respect consumers' desires to sample whichever firms they please, but is able to acquire and pass consumer reviews into firms' ratings, thereby distorting market beliefs and thus demand. Furthermore, while I model the rater's objective as consumer welfare, in equilibrium, the entry incentives firms face inexorably feed into consumer welfare and so optimal design must necessarily account for firms' incentives.

Intuitively, the designer faces a quality-quantity trade-off. On the one hand, they must ensure that *ceteris paribus*, consumers purchase from the highest quality firms. On the other hand, they must ensure sufficient entry to relax capacity constraints and thus satisfy demand. The crucial insight of my analysis is that these two objectives are in tension. To understand why, suppose that consumers were to visit only the highest-rated firms. This would result in entering firms — who are necessarily lower-rated, as their quality is unknown — taking longer to receive profits. This would depress entry incentives and thus the total quantity of firms in the market.

My main result — Theorem 2 — operationalizes this trade-off within a class of policies I term "simple policies". Such policies essentially allow the rater to ignore a fraction of the newly generated reviews for a firm. This fraction must be a function of the firm's current rating but is otherwise general. I show that the optimal simple policy takes a "tenure" form, whereby firms that attain a sufficient rating obtain no further feedback, guaranteeing they both maintain this rating and sell at maximal capacity in perpetuity. This policy permits a measure of low-type firms to continue to sell (as their rating remains unchanged post tenure), and thus entails a loss in quality for consumers. Nevertheless, by doing so, it more than compensates for this loss by stimulating participation incentives, thereby increasing the total quantity of firms in the market.

My baseline model assumes that prices are exogenous and profits thus scale with sales volume. Applications include online labor platforms such as Thumbtack, TaskRabbit, and Upwork, wherein workers face significant upfront costs, prices are heavily regulated, capacity constraints bind, and ratings form an important determinant of demand.[3] Pallais (2014) finds evidence that newly entering workers in a freelance labor platform have trouble attracting interest in the absence of a reputational track record. An alternative application is the market for digital apps.[4] Exogenous prices create two separate inefficiencies. First, firms cannot internalize excess demand, leading to rationing, which is inefficient since consumers weakly prefer consumption at any firm over not consuming. Second, firms cannot internalize the future social value of information generated by consumers. To explore the importance of these inefficiencies, in Section 5 I extend the model to allow for competitive search pricing. At each rating, prices adjust to ensure that markets clear but also that firms do not have an incentive to deviate to off-path prices. Prices alleviate the non-convexity present in the baseline model, rendering full transparency optimal. As such, my findings contribute to our understanding of how price rigidity impacts the optimal design of rating systems in the presence of endogenous entry by firms.

On a technical note, solving for the optimal simple policy involves maximizing over the invariant distribution generated by the equilibrium choices of agents and constitutes a theoretical contribution to the literature on dynamic information design (see section 1.1 for further discussion). Whereas typical information design problems involve maximization subject to distributional constraints that are either static or exogenous, here the rater maximizes over distributions that are themselves

[3] Recent surveys document the importance of customer reviews in the gig economy (https://insight.kellogg.northwestern.edu/article/gig-workers-customer-reviews, https://www.bbc.com/news/uk-england-bristol-63483597). Furthermore, Upwork rewards contractors with better ratings with greater visibility in the form of its “top-rated” badge (see https://support.upwork.com/hc/en-us/articles/211068468-How-to-become-Top-Rated-on-Upwork), while Thumbtack explicitly factors ratings into search rankings (see https://help.thumbtack.com/article/my-rank-in-search-results).

[4] Such markets exhibit a zero lower bound on prices, with apps instead profiting from consumption volume (Bisceglia and Tirole, 2025). As of December 2024, 97% and 95% of apps in Google Play and Apples App Store were freely available.

determined in equilibrium by the dynamic choices of a continuum of agents. I employ techniques from the theory of "mean-field" games over diffusion processes to overcome these challenges.

The paper proceeds as follows. I begin by placing my work within its related literature (Section 1.1). Section 2 introduces the model and characterizes equilibrium under *full transparency*. Section 4 turns to ratings design and delivers the main findings. I extend the baseline analysis to allow for prices in Section 5, while Section 7 concludes by offering avenues for future work. All proofs are contained in Appendix A.

1.1 Related Literature

I study how ratings design can incentivize entry. In my model, social learning disadvantages new and struggling firms, and thereby I connect to the literature on "incentivizing exploration" problem (Che and Hörner, 2018; Kremer et al., 2014), which I discuss in greater detail following Theorem 2. My paper contributes to the literature on dynamic information design. Hörner and Lambert (2021) study the design of ratings in order to incentivize a single firm to exert hidden effort and improve output quality. Their optimal rating scheme has the feature that older information is phased out in order to stimulate current effort. In my optimal rating, older reviews are as informative as newer reviews regarding underlying quality, and thus, such a feature does not play a role.

My paper models information design as optimizing over the stationary equilibrium of a mean-field game. Che and Tercieux (2023) study the design of stationary queuing protocols by maximizing over invariant distributions in a setting with only one seller. Moscarini (2005) develops a mean-field equilibrium model that bears resemblance to my full transparency benchmark, but abstracts from information design.

My analysis shares features of the literature on collective experimentation and social learning. Models of collective learning via Brownian diffusion processes can be found in Bolton and Harris (1999) and Bergemann and Välimäki (1997). Moscarini and Smith (2001) models the sampling rate of information as a direct control variable. Beyond the classic works in the social learning literature (Bikhchandani et al., 1992; Banerjee, 1992), Acemoglu et al. (2022) also study the endogenous speed of learning in an observational learning framework. Campbell et al. (2020) study social learning between consumers connected on a network with endogenous entry. These works all abstract from the implications of information design.

My model also contributes to the literature on firm dynamics (Hopenhayn, 1992; Luttmer, 2007; Jovanovic, 1982; Board and Meyer-ter-Vehn, 2022). Atkeson et al. (2015) study how taxation policies affect initial selection of informed firms, whereas I study how information polices improve social learning about uninformed firms. Hörner (2002) studies how competition boosts endogenous effort incentives, whereas I study how information design boosts endogenous entry incentives.

2 Model

Time is continuous and infinite. The market consists of consumers, firms, and a rater. As I will be studying stationary equilibria, calendar time subscripts are dropped, with t henceforth denoting the age of a firm. For a discussion of various elements of the model, see Section 6.

Firms – A large, infinitely elastic supply of firms is available at any instant to enter the market. Each firm is one of two types, $\theta \in \{0, 1\}$, with a new entrant being of type $\theta = 1$ with probability p_0. Types are fixed throughout the life of a firm and are initially unknown to all market participants, including the firm itself. While active, a firm of type θ is associated with *cumulative review process* $(X_t)_{t\geqslant 0}$ that evolves according to

$$dX_t = r_t\lambda_t\theta\, dt + \sqrt{r_t\lambda_t}\sigma\, dZ_t. \tag{1}$$

The Wiener process $(Z_t)_{t\geqslant 0}$ is independent of θ, and $\sigma \in (0, \infty)$.[5] The process $(\lambda_t)_{t\geqslant 0}$ is defined by $\lambda_t \equiv \alpha\pi_t + \epsilon$, where π_t is the rate at which the firm sells to consumers who use ratings to direct their search, $\alpha > 0$ is an informativeness parameter, and $\epsilon > 0$ represents a constant background learning rate, which represents a stream of consumers who search at random, independently of ratings. Crucially, π_t is determined endogenously. Finally, the process $(r_t)_{t\geqslant 0}$ is determined by the policy of the rater, and is described below.

A natural state variable for a firm's problem is the instantaneous probability p_t that the firm is of high type given the information contained in $(X_t)_{t\geqslant 0}$, which I henceforth refer to as a firm's *rating*. Formally, let $\mathbb{F}^x \equiv \{\mathcal{F}^x_t\}_{t\geqslant 0}$ denote the natural filtration generated by $(X_t)_{t\geqslant 0}$. Finally, let $\mathbb{E}^x_t$ denote the conditional expectation with respect to $\mathcal{F}^x_t$, so that $p_t \equiv \mathbb{E}^x_t(\theta)$ and p_0 is the firm's

[5] See Bergemann and Välimäki (1997), Moscarini and Smith (2001) and Bolton and Harris (1999) for similar approaches.

rating at $t = 0$. Ratings evolve according to:

$$dp_t = \frac{\sqrt{r_t \lambda_t}}{\sigma} p_t(1 - p_t) d\bar{Z}_t, \tag{2}$$

where $d\bar{Z}_t = dX_t - p_t dt$. Since $\left(\bar{Z}_t\right)_{t \geqslant 0}$ is a standard Wiener process with respect to the filtration generated by $(p_t)_{t \geqslant 0}$, ratings follow a martingale diffusion.

Firms pay a lump-sum cost of $K > 0$ to enter the market. Flow profits equal the service rate of consumers using ratings to direct search $\pi_t \in [0, \bar{\pi}]$, with the price of output normalized to 1, and where $\bar{\pi} \in (0, \infty)$ is an exogenous service capacity constraint.[6] Crucially, π_t is endogenously determined by consumer choice, but taken as given by a firm. Finally, firms discount at rate $\rho > 0$ and face a constant hazard rate $\delta > 0$ of exogenous attrition. The present value to a firm with a rating p is thus given by

$$V(p) = \mathbb{E}^x \left[\int_0^\infty e^{-(\rho + \delta)t} \pi_t \, \mathrm{d}t \mid p_0 = p \right], \tag{3}$$

subject to the law of motion (2). I assume that $K < \bar{V} \equiv \frac{\bar{\pi}}{\rho + \delta}$, so that the present value of selling out forever is strictly greater than the entry cost, which is necessary for the existence of an equilibrium with positive entry:

A perfectly elastic supply of firms implies that, were there net positive lifetime profits available to firms upon entry, an infinite mass of firms would enter. Conversely, if there were net negative profits from entry, no firm would enter:[7]

$$\eta = \begin{cases} \infty & \text{if } V(p_0) > K \\ \in [0, \infty] & \text{if } V(p_0) = K \\ 0 & \text{if } V(p_0) < K. \end{cases} \tag{4}$$

Consumers – A unit measure of risk-neutral consumers use ratings to find firms. If a consumer purchases from a firm with type θ, they receive a payoff that is normally distributed with mean θ and variance σ^2. Given risk neutrality, a firm's rating is a sufficient statistic for consumer value.

[6] I assume that consumers not using ratings for search generate no profits and require no capacity, both of which are normalizations that are easily relaxed.

[7] One could impose that $\eta \in [0, \bar{\eta}]$ for sufficiently large $\bar{\eta}$ in order to ensure compactness and without altering the analysis.

To this end, let F denote the measure of available firms over ratings $p \in [0,1]$, henceforth referred to as the *ratings distribution*.[8] I further assume that consumers who do not use ratings to direct search (as captured by the parameter ϵ) receive zero net surplus. Henceforth, I refer to consumers simply as those using ratings to direct search.

I assume that, given the available choices and firms' service capacity constraints, consumers perform directed search subject to random rationing (Guerrieri and Shimer, 2013; Lester, 2011). In words, consumers choose which rating they would like to search at, given the available options. (I will often refer to each rating p as a separate "sub-market".) Given these search choices, firms and consumers at each sub-market p are matched bilaterally and at random. If more firms exist in sub-market p than consumers, then firms are rationed at random, and vice versa. I now describe this matching technology more formally.

A directed search profile is a distribution G on $[0,1]$ that is absolutely continuous with respect to F, and which denotes the allocation of consumers to firms. Specifically, $G(p)$ denotes the measure of consumers attempting to match with firms of rating at most p. Absolute continuity ensures that consumers may only attempt to match with existing firms. The measure F together with the search profile G induce a *market tightness* $\tau(p) = dG(p)/\bar{\pi} dF(p)$ at each rating $p \in \text{supp}(G)$, which computes the ratio of demand ($dG(p)$) to supply ($\bar{\pi} dF(p)$) at each sub-market.[9] This tightness determines a consumer's probability $\Theta(p)$ of purchasing from a firm with rating p, where

$$\Theta^{-1}(p) = \max\left\{\tau(p), 1\right\}. \tag{5}$$

If a consumer chooses to direct their search to firms with rating p, they are served with probability $\Theta(p)$, and thus the expected value of visiting a firm with rating p to a consumer is $p\Theta(p)$. Equation (5) evaluates the market tightness $\tau(p)$ at each sub-market. If this ratio is below one, each consumer is guaranteed to be served. Otherwise, consumers are randomly rationed. Search is costless, and thus consumers must be indifferent between any firms that are visited under G:

[8] The term "distribution" is used in line with the literature on firm dynamics and is a misnomer, as generically the mass of firms will not equal 1 in equilibrium.

[9] The expression dG/dF denotes a Radon-Nikodym derivative.

Definition 1. The directed search strategy G is *incentive compatible (IC)* with respect to F if

$$p\Theta(p) = p'\Theta(p') \text{ for all } p, p' \in \text{ supp}(G). \tag{6}$$

As is common in models of directed search, we extend the definition of the matching function Θ beyond the support of G by imposing the "market utility" constraint, whereby at any $p \notin \text{supp}(G)$, the service rate is computed as the lowest rate that ensures consumers would find it profitable to direct their search toward that sub-market (Eeckhout and Kircher, 2010). That is, for $p \notin \text{supp}(G)$,

$$\Theta(p) = \inf_{\theta \in [0,1]} \{\theta \mid p\theta \geq \max_{p' \in \text{supp}(G)} p'\Theta(p')\}. \tag{7}$$

Rater – I restrict my analysis of ratings design to a class of policies that are both tractable to study and practical to implement:

Definition 2. A *simple (rating) policy* is an $\mathbb{F}^x$-measurable, locally integrable function $r : [0,1] \to [0,1]$ such that the cumulative review process $(X_t)_{t \geqslant 0}$ increments according to:

$$dX_t = r(p_t)\lambda_t \theta dt + \sqrt{r(p_t)\lambda_t}\sigma dZ_t, \quad \text{where } p_t = \mathbb{E}_t^x(\theta).$$

Simple policies prescribe a rate at which reviews are solicited from consumers and passed through into the firm's record $(X_t)_{t \geqslant 0}$.[10] This rate – captured by the function r – is a function only of the firm's current rating p_t and not upon the firm's age t directly. (I provide examples of such policies in Section 4.) More specifically, the law of motion for ratings under a simple policy r,

$$dp_t = \frac{\sqrt{r(p_t)\lambda_t}}{\sigma} p_t(1-p_t) d\bar{Z}_t \tag{8}$$

is still a martingale diffusion. Thus, by choosing $r(p_t) \in [0,1]$, the rater simply controls how rapidly ratings disperse. I endow the rater with the ability to commit to a policy.[11]

Ratings Distribution – The combination of evolving ratings and a continuously churning

[10] The rating p is the only payoff-relevant variable, and is Markovian under simple policies. Thus, firms need not recall past ratings in order to compute present values.

[11] Practically speaking, it would be challenging to update a policy in light of firm-specific or even distributional evidence as it evolves. Yelp! stress that since their algorithm is automated, their staff cannot override the inclusion/exclusion of reviews from a firm's rating. See `https://www.yelp-support.com/article/Why-would-a-review-not-be-recommended`.

positive mass of firms gives rise to a unique ratings distribution F, which admits a density f almost everywhere. The law of motion for this density follows the Fokker-Planck forward equation, with stationarity imposed and subject to various boundary conditions.

Proposition 1. For a fixed IC directed search strategy G, entry rate $\eta > 0$, and simple policy r, the invariant ratings distribution admits a density f almost everywhere, and is unique. At p such that $r(p) > 0$, the density f satisfies the Fokker-Planck forward equation:

$$\frac{\partial^2}{\partial p^2}[\Sigma(p)f(p)] - \delta f(p) = 0, \qquad \Sigma(p) \equiv \frac{r(p)(\alpha\pi(p)+\epsilon)p^2(1-p)^2}{2\sigma^2}, \tag{9}$$

subject to the following conditions:

(D_1) $\Sigma(p_0)[f'_-(p_0) - f'_+(p_0)] = \eta$,

(D_2) $\Sigma(0)f(0) = \Sigma(1)f(1) = 0$,

(D_3) $\frac{\partial}{\partial p}\Sigma(p)f(p) \in \mathcal{C}^1([0,p_0) \cup (p_0,1])$,

(D_4) $\Sigma(p)f(p) \in \mathcal{C}^1([0,1])$,

(D_5) $\delta \int_0^1 f(p)\,\mathrm{d}p = \eta$.

At p such that $r(p) = 0$, F admits an atom $\Omega(p)$ with: (D_6) $\delta\Omega(p) + [\Sigma(p)f(p)]'_- = 0$.

The quantity $\Sigma(p)$ represents the *elasticity* of beliefs (Moscarini and Smith, 2001), and quantifies the rate of change of ratings as a function of their current level p. It depends positively upon residual uncertainty $p(1-p)$, the service rate $\pi(p)$, and the rater's pass-through rate $r(p)$. Intuitively, when $p(1-p)$ is high, Bayes' rule implies that each review is more informative, while when either $\pi(p)$ or $r(p)$ is high, more reviews are obtained by the firm; in all cases, ratings move faster.

The boundary conditions are derived formally by requiring the system to obey a certain conservation of probability mass principle; specifically, the differential operators that define the forward and backward Kolmogorov equations must form a pair of proper adjoint operators. (See Appendix A.1 for a formal treatment.) Though technical in nature, many of these conditions provide economic insight into the dynamical system when properly interpreted. For instance, condition D_1 states that the outflow of incumbent firms in either direction away from the initial rating p_0 must equal the rate of inflow by new entrants. Condition D_5 is an aggregate balance equation, which states that outflows due to attrition must equal inflows due to entry. With these ingredients in hand, I close this section with a formal definition of a stationary equilibrium.

Definition 3. A *stationary equilibrium* consists of a measure F on $[0,1]$, a distribution G on $[0,1]$ and a scalar η such that: 1) given $\{G,\eta\}$, the ratings distribution F satisfies the conditions given in Proposition 1;[12] 2) given $\{F,\eta\}$, the directed search strategy G satisfies equation (6); 3) given $\{F,G\}$, the entry rate η satisfies condition (4).

3 Equilibrium Under Full Transparency

I first characterize stationary equilibria under the assumption that the rater simply includes all consumer feedback into each firm's rating, *full transparency*, so that $r(p)=1$ for all $p\in[0,1]$. I start by characterizing consumers' equilibrium search strategy G. This is a potentially complicated problem, as there exists a non-trivial fixed-point relationship between G and the rating distribution F; intuitively, F is determined by firm entry decisions, which are determined by consumer search as given by G, while G must respect the availability of firms as given by F. However, the solution turns out to take a simple *threshold strategy* form:[13]

Lemma 1. There exists a unique $p^*\in(0,1]$ that solves

$$p^*=\bar{\pi}\int_{p^*}^{1}pf(p)\,\mathrm{d}p \tag{10}$$

and such that

$$g(p)=\begin{cases}\frac{\bar{\pi}pf(p)}{p^*} & \text{if } p\geqslant p^*\\ 0 & \text{if } p<p^*.\end{cases} \tag{11}$$

Equation 11 says that only firms above a certain threshold p^* are visited by consumers. As the rating p increases above p^*, consumer search intensifies, causing a probability of rationing that increases linearly with ratings. Thus, consumers balance expected quality p against likelihood of trade $\Theta(p)$ in order to remain indifferent across all sub-markets visited; notice that the condition $g(p)=\bar{\pi}pf(p)/p^*$ is equivalent to $p\Theta(p)=p^*$, so that expected utility is constant. At sub-market p^*, consumers are served with probability 1 so that they do not have an incentive to deviate and visit a firm with a slightly lower rating (were $\Theta(p*)<1$, the market utility constraint (7) would imply that such a deviation would be profitable). Equation 10 computes p^* from the consumer

[12] Recall that F is called a ratings distribution, despite not in general a distribution function.

[13] Since $r(p)=1$ everywhere, Proposition implies that the rating distribution admits a density everywhere.

market clearing condition

$$\int_0^1 g(p)\,\mathrm{d}p = 1.$$

I turn now to the firms' problem. To ensure that the matching problem is well-defined, it must be that the total quantity of consumer matches equals the total quantity of firm matches (scaled by the capacity $\bar{\pi}$). Thus, the rate at which firms sell $\pi(p)$ must satisfy $\bar{\pi} f(p)\pi(p) = g(p)\Theta(p)$, or equivalently,

$$\pi(p) = \bar{\pi} \min\{\tau(p), 1\}. \tag{12}$$

Lemma 1 implies that:

$$\pi(p) = \begin{cases} \bar{\pi} & \text{if } p \geqslant p^* \\ 0 & \text{if } p < p^*. \end{cases} \tag{13}$$

Firms with low ratings (below p^*) are not visited at all, whereas firms with high ratings (above p^*) sell out. Thus, while rationing is more severe at higher-rated firms in order that consumers remain indifferent, from an individual firm's perspective, equation (12) implies that flow revenues take the form of a step function. With these preliminary lemmas in hand, I now provide a full characterization of equilibrium under fully transparent ratings.

Theorem 1. Under full transparency, there exists a unique stationary equilibrium $E = \{F, G, \eta\}$, where the ratings distribution F admits a density f on $[0, 1]$, the directed search strategy G is as defined in Lemma 1, and featuring:

1. A positive, finite rate of entry ($\eta \in (0, \infty)$);

2. A step flow profit function ($\pi(p) = \mathbb{I}_{p \geq p^*}$);

3. Greater consumer rationing at higher ratings ($p\Theta(p) = p^*, \forall p \geq p^*$);

4. An S-shaped firm value function ($V''(p) > 0$ for all $p \in [0, p^*)$ and $V''(p) < 0$ for all $p \in (p^*, 1]$).

Proving the existence and uniqueness of equilibrium amounts to solving a fixed-point problem with respect to the two scalars p^* and η. I now provide an outline of its proof, as this approach lays the foundation for the normative analysis in Section 4.

Let $R_*(\eta)$ denote the solution to (10) subject to (11) for a given η, and $R_\eta(p^*)$ the solution to (4) for a fixed p^*. That is, $R_*(\eta)$ computes the consumer threshold p^* that satisfies the IC conditions outlined in Lemma 1 for a fixed rate of entry η by firms, while $R_\eta(p^*)$ evaluates the firms' value function $V(p)$ for a fixed consumer threshold p^* and then computes the entry rate η according to the condition (4). By construction, any pair $\{p^*, \eta\}$ that jointly satisfy $p^* = R_*(\eta)$ and $\eta = R_\eta(p^*)$ constitute an equilibrium.

Consider first $R_*(\eta)$. In order to understand this relationship, we must understand how the entry rate η affects the density f and thereby p^*. I show an important property of the density f, namely that it is linearly homogeneous in η and thus higher η increases f pointwise.[14] Thus, as η rises, the measure of firms above any given rating increases, and so the consumer threshold equation (10) then implies that $R_*(\eta)$ must rise continuously with η. Simply put, when there are more firms at the upper end of the distribution, there is greater capacity to serve consumers, and so consumers can afford to increase their quality standards when searching. When η is close to 0, hardly any firms exist in the market, so capacity-constrained service implies that almost all firms must be serving consumers, and so $R_*(\eta)$ is close to 0. Conversely, when η is large, almost all consumers can be served by the top-rated firms, so that $R_*(\eta)$ is close to 1. In summary, R_* is a strictly increasing, continuous function on $[0, \infty)$ with $\lim_{\eta\to 0} R_*(\eta) = 0, \lim_{\eta\to\infty} R^*(\eta) = 1$,

Next, consider $R_\eta(p^*)$. Here we must understand how the consumer threshold p^* affects the value function V of firms, and thereby the entry rate η. I show for fixed $p \in [0, 1]$, the value $V(p)$ is continuous and strictly decreasing as a function of p^*. Intuitively, firms make positive profits on average for a shorter fraction of their lifetimes, as they must achieve a higher rating in order to do so, and thus at every rating p their present-value $V(p)$ is decreasing with p^*. For p^* close to 0, I show that for all $p \in [p^*, 1], V(p)$ is close to $\bar{V}$. This is because firms sell out for almost their entire market duration, and thus are virtually guaranteed their maximal value $\bar{V}$ at every rating. In particular, $V(p_0) \approx \bar{V} > K$ by assumption and so the corresponding entry rate $R_\eta(p^*)$ is infinite. Conversely, for p^* close to 1, firms almost surely make zero profits, and thus $V(p)$ is close to 0 at $p \in [0, p^*]$. In particular, $V(p_0) \approx 0 < K$ and so the corresponding entry rate $R_\eta(p^*)$ is zero. These observations combine to imply that there exists a unique p^{**} such that $V(p_0) = K$, so that for

[14] Hopenhayn (1992) also exhibits linear homogeneity of the invariant firm distribution with respect to the entry rate.

Figure 1: Reaction functions under full transparency.

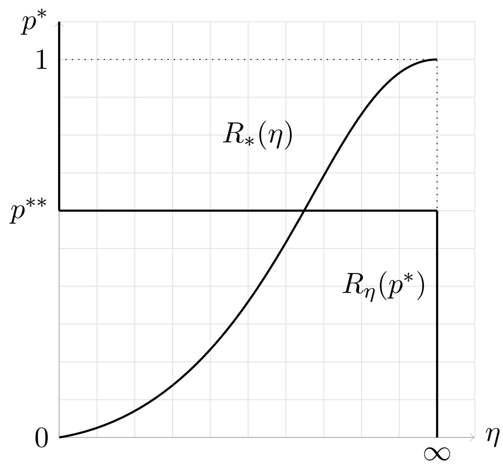

$R_*(\eta)$ is the optimal consumer threshold, given the entry rate η. $R_\eta(p^*)$ is the optimal entry rate, given the consumer threshold p^*. The intersection point corresponds to the unique stationary equilibrium.

$p^* \in [0, 1]$,

$$R_\eta(p^*) = \begin{cases} \infty & \text{if } p^* > p^{**} \\ \in [0, \infty] & \text{if } p^* = p^{**} \\ 0 & \text{if } p^* < p^{**}. \end{cases}$$

Combining these features of the functions R_* and R_η, there exists a unique pair $\{p^*, \eta\}$ that satisfies the equilibrium conditions outlined in definition 3. Figure 1 demonstrates this construction graphically.

To understand part 4 of Theorem 1, note that for $p < p^*$, firms are making zero profits while at $p \geq p^*$ firms enjoy constant sell-out revenues. Since ratings follow the martingale diffusion (2) and thus constitute a mean-zero gamble, firms at $p < p^*$ value information as the upside of selling-out outweighs the downside of continuing to make zero profits, while at $p \geq p^*$ firms dislike information as the downside of losing their current rating and profits outweighs the upside of continuing to sell out. Mathematically, one can show that a firm's continuation value $V(p)$ satisfies:

$$V''(p) = \frac{\rho + \delta}{\Sigma(p)} \left[V(p) - \frac{\pi(p)}{\rho + \delta} \right]. \tag{14}$$

When the firm makes maximal / zero revenues, this quantity is negative / positive.

Put differently, firms prefer a *fast-slow* profile of learning, which would make the climb easier and the fall less precipitous. In equilibrium, however, the profile is precisely the reverse, i.e. *slow-fast*. This insight will prove invaluable when considering the rater's design problem in Section 4.

4 Ratings Design

I now turn to the normative analysis of the paper, wherein the rater chooses a simple policy (Definition 2) in order to maximize consumer surplus. The following are leading examples of simple policies:

- By setting $r(p) = 1$ for all $p \in [0, 1]$, it is clear that the **fully informative** policy is simple. By setting $r(p_0) = 0$, the rater can also implement the **fully uninformative** policy.
- r is a **tenure policy** if $r(p) = \mathbb{I}_{p<\tilde{p}}$ for some $\tilde{p} \in (0, 1)$. These policies "tenure" firms forever once they reach a certain threshold $\tilde{p}$, as they are guaranteed to remain at this rating (if they obtain it) until exogenous attrition. Prior to tenure, the rater acquires and posts all available feedback. Any tenure rating with $\tilde{p} \leq p_0$ is fully uninformative.
- r is **sticky** at $\tilde{p} > p_0$ if $r(\tilde{p}) = 0$ and is not left continuous at p and sticky at $\tilde{p} < p_0$ if $r(p) = 0$ and is not right continuous at p. Thus, tenure policies are sticky at their tenure threshold $\tilde{p}$.

An equilibrium E is *implementable* if there exists a simple policy r such that E is a stationary equilibrium with respect to r (in which case, I say that r implements E). The rater's optimization problem is to maximize consumer welfare across all implementable equilibria through their choice of r.

The following lemma demonstrates that equilibria implementable under simple policies retain the tractable structure obtained under full transparency.

Lemma 2. Let the simple policy r implement the stationary equilibrium $E = \{F, G, \eta\}$. Then there exists a unique consumer threshold $p^*_E \in (0, 1)$ such that

1. A positive, finite rate of entry $\eta \in (0, \infty)$.
2. $\lim_{p \uparrow p^*_E} G(p) = 0$, and $\Theta(p^*_E) = 1, \Theta(p) < 1$ for all $p > p^*_E$.

3. A firm's continuation value $V(\cdot)$ is convex on $[0, p_E^*)$ and concave on $(p_E^*, 1]$ wherever $r(p) > 0$.

Furthermore, consumer welfare is given by p_E^*.

Part 1) states that consumers still follow a threshold consumption rule, with rationing above the associated cut-off p^* and guaranteed service at p^*. This is a result of simple policies inducing the same stationary structure as in Section 2. For part 2), note that since consumers are served at rating p^* with probability $\Theta(p^*) = 1$, and are indifferent at all other ratings in the support of the search strategy G, p^* is the expected utility of each consumer and thus must equal overall consumer welfare.

To understand part 3), note that as parts 1) and 2) verify, under any simple policy, the firm again makes either maximal or zero revenues, depending on their current rating p. By equation (8), ratings still follow a martingale under simple policies. Thus, just as in full transparency, in the latter region, the firm has positive option value from learning and so the value function is convex, and vice versa.[15] Finally, part 4) is straightforward; if $\eta = 0$, then consumers would purchase from any firm that enters the market ($p^* = 0$), providing positive entry profits and thus violating the entry condition (4). If $\eta = \infty$, then consumers would only visit the best firms ($p^* = 1$), which provides entering firms zero profits almost surely, thus again violating (4).

The rater faces a stark tension. On the one hand, they seek to maximize consumer welfare (p^*). On the other hand, a higher p^* lowers $V(p_0)$ and thus the entry rate η and the quantity of active firms. How then should the rater balance this trade-off through optimal ratings design? Within the context of simple policies, Theorem 2 provides a definitive answer and constitutes my main finding.

Theorem 2. There exists a unique $q \in (p_0, 1)$ such that the optimal simple policy is the tenure policy $r(p) = \mathbb{I}_{p<q}$. Furthermore, q is such that in equilibrium:

1. Consumers only visit tenured firms ($G(p) = \mathbb{I}_{p \geq q}$);

2. Firms sell out if and only if tenured ($\pi(p) = \bar{\pi}\mathbb{I}_{p \geq q}$);

The proof of Theorem 2 begins by fully characterizing implementable equilibria within the class of tenure policies. For a given tenure policy, equilibrium is unique. Furthermore, ranging

[15] The firm's continuation value $V(p)$ satisfies equation (14) almost everywhere. Note that if $r(p) = 0$, then p is an absorbing state, and thus $V''(p)$ is not defined.

across all tenure policies, consumer welfare is strictly single-peaked in the tenure rating $\tilde{p}$ at $\tilde{p} = q$, where q is precisely the tenure rating such that 1) firms sell out if and only if they achieve tenure ($\pi(p) = \bar{\pi}\mathbb{I}_{p \geq q}$), and 2) firm value functions are such that $V(p_0) = K$, with the entry rate η ensuring that there are precisely enough firms at the tenure rating q to serve all consumers.[16]

To understand this result, take first a tenure threshold $\tilde{p} > q$. In this case, I show that if firms were to sell out only once tenured so that the consumer threshold $p^* = \tilde{p}$, then participation incentives would be too weak to permit entry ($V(p_0) < K$). Intuitively, it takes so long to acquire tenure and profits that entry is not profitable. Thus, to sustain entry, it must be that in equilibrium, firms start selling out prior to tenure, so that $p^* < q$ and consumers are rationed on the interval $(p^*, \tilde{p}]$. I then show constructively that there exists a $p^\dagger \in (p^*, \tilde{p}]$ such that the tenure policy $r(p) = \mathbb{I}_{p<p^\dagger}$ implements strictly higher equilibrium consumer welfare. Under this policy, the rater maintains sufficiently strong entry incentives ($V(p_0) = K$) to ensure enough firms exist at rating $p^\dagger$ to serve all consumers, ensuring consumer welfare $p^\dagger > p^*$. The right panel of Figure 2 demonstrates this argument graphically.

If instead $\tilde{p} < q$ and firms were guaranteed to sell out once tenured at $\tilde{p}$, then tenure would be achieved so fast that participation incentives would be too strong ($V(p_0) > K$), contradicting the boundedness of entry established in Lemma 2 part 4). Thus, in order that $V(p_0) = K$, equilibrium must necessarily involve firm rationing at $\tilde{p}$, so that $\pi(\tilde{p}) < \bar{\pi}$. Specifically, an excess measure Ω of firms must exist at rating $\tilde{p}$ in order that $\bar{\pi}\Omega > 1$ and $\pi(\tilde{p}) < \bar{\pi}$ be small enough that $V(p_0) = K$. Intuitively, in order to balance entry incentives, early tenure for firms must be offset by rationing, so that firms are not guaranteed to sell out once tenured. Notice that consumers do not benefit from this excess supply of firms, as consumer welfare is $\tilde{p}$ whether or not firms are rationed. The rater can thus do better by tenuring at any level $p^\dagger \in (\tilde{p}, q]$. Any such policy would still ensure enough firms exist at the tenure threshold to avoid consumer rationing and sustain consumer welfare $p^\dagger > \tilde{p}$. The left panel of Figure 2 demonstrates this case graphically. In particular, that $q > p_0$ follows from this latter reasoning and demonstrates that a policy of no information is never optimal.

The same intuition can be cast in terms of the reaction function analysis offered in Theorem 1. Recall that the function $R_\eta(\cdot)$ traces the optimal entry rate as a function of the consumer threshold. If $\tilde{p} > q$, then the above argument implies that the associated equilibrium consumer threshold p^*

[16] See Lemma A.7 for a formal proof.

must be below q in order for the entry condition $V(p_0) = K$ to hold and hence $R_\eta(p^*) \in (0, \infty)$ to hold. Consumer welfare in this case is thus lower than q. If $\tilde{p} < q$, then firm rationing at $\tilde{p}$ occurs in order to ensure that $R_\eta(\tilde{p}) \in (0, \infty)$, so that again consumer welfare is lower than q. Figure 3 demonstrates these two cases.

The proof of the Theorem concludes by showing that any non-tenure policy is strictly improved upon by a tenure policy. To see this, observe first that for a fixed p^*, the tenure policy $r(p) = \mathbb{I}_{p<p^*}$ maximizes firm continuation values V point-wise. To see why, note that Lemma 2 demonstrated that V is convex below p^* and concave above p^*. Since ratings still follow a martingale, V is maximized through a policy of maximal feedback on its convex part and minimal feedback on its concave part.[17] Starting from an equilibrium under an arbitrary non-tenure policy r with consumer threshold p^*, the previous argument combined with the entry condition that $V(p_0) = K$ then implies that there exists a tenure policy supporting a higher equilibrium p^*. If not, the condition that $V(p_0) = K$ cannot be satisfied under the constructed tenure policy, contradicting Lemma 2 part 4).

More generally, the proof of Theorem 2 demonstrates that R_η is maximized pointwise by the optimal tenure policy across all simple policies, and in this sense it stimulates the greatest incentives for entry and thus sustains the highest level of consumer welfare. Note that this reasoning is only in partial equilibrium (it is silent on how the optimal policy affects R_*), and thus it is not necessarily true that the equilibrium entry rate η under the optimal policy is maximal across all policies. Since the rater only cares about η indirectly through its impact on consumer welfare p^*, verifying this additional claim is not necessary for optimality.

Managing rationing – Theorem 2 parts 1 and 2 combined imply that the optimal tenure policy calibrates incentives so as to avoid rationing; in equilibrium, the measure $\Omega(q)$ of firms achieving the maximal rating q satisfies $\bar{\pi}\Omega(q) = 1$, ensuring supply and demand are met. That being said, many policies exist that also obviate rationing but that are strictly dominated by the optimal tenure policy:[18]

Corollary 1. For any $a \in (0, 1)$, there exists $p_a \in (0, q)$ such that the policy $r(p) = a\mathbb{I}_{p<p_a}$

[17] Kolotilin et al. (2022) provide a model in which an S-shaped value function induces optimal persuasion that takes a similar upper censorship form, resulting from the same intuition. In their model, the value function is exogenous, whereas here, it is endogenously determined via the variable p^*, while the dynamic nature of my model requires using alternative techniques (Escudé and Sinander, 2023). Gindin and Shimko (2023) study a dynamic political economy setting wherein the underlying value function is convex, concave due to asymmetric information and agency frictions.

[18] The proof follows precisely the steps in Lemma A.6 used in proving Theorem 2 and is thus omitted.

Figure 2: Policy Improvement

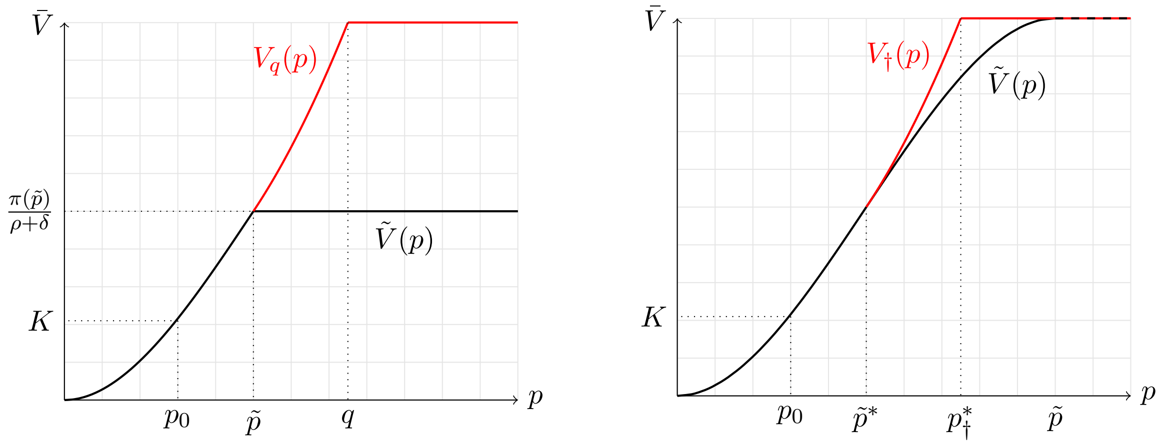

Left panel: $\tilde{p} < q$. $\tilde{V}(p)$ is the value function under policy $r(p) = \mathbb{I}_{p<\tilde{p}}$, $V_q(p)$ is the value function under policy $r(p) = \mathbb{I}_{p<q}$. Early tenure leads to excess entry, over-supply of firms and hence firm rationing at $\tilde{p}$. Right panel: $\tilde{p} > q$. $\tilde{V}(p)$ is the value function under policy $r(p) = \mathbb{I}_{p<\tilde{p}}$, $V_\dagger(p)$ is the value function under policy $r(p) = \mathbb{I}_{p<p^\dagger}$. Late tenure leads to insufficient entry, under-supply and hence consumer rationing above $\tilde{p}^*$.

Figure 3: Reaction Functions, Tenure Policies

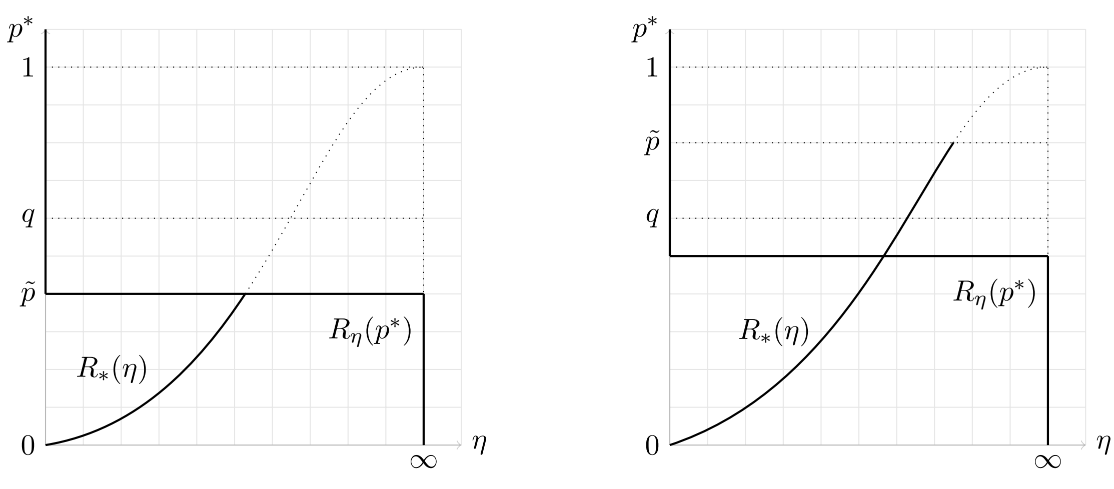

Left panel: $\tilde{p} < q$. Consumer welfare is bounded above by $\tilde{p}$. Right panel: $\tilde{p} > q$. Insufficient entry induces consumer welfare $p^* < q$.

implements a unique equilibrium wherein $G(p) = \mathbb{I}_{p \geq p_a}$ and $\pi(p) = \bar{\pi}\mathbb{I}_{p \geq p_a}$.

The policies in Corollary 1 are sticky policies and similar to tenure policies, with the added feature that they also restrict information at ratings below the tenure threshold p_a. By Lemma 2 part 3, restricting information prior to tenure reduces the value function due to the convexity of $V(\cdot)$ on this region. In order that $V(p_0) = K$, it must be that the tenure level $p_a < q$, so that such policies support lower consumer welfare than the optimal tenure policy. That is, among those policies that perfectly mitigate rationing, the optimal tenure policy provides the strongest incentives for entry, and thus through a boost in the quantity of active firms, supports the highest level of consumer welfare.[19]

Incentivizing exploration – As previously discussed, my paper connects to the literature on incentivizing exploration (Che and Hörner, 2018; Kremer et al., 2014). These papers find that redirecting demand toward lower-rated products can help support more efficient social learning, which in turn maximizes the likelihood of discovering high-quality products. By contrast, the optimal tenure policy in Theorem 2 shuts down social learning entirely (consumers only consume at tenured firms, for whom new reviews are entirely discarded) and thus precludes the possibility of discovering high-quality firms post-tenure. This distinction highlights the role of endogenous entry incentives in shaping optimal rating design. In their setting, consumers gain solely from discovering the highest quality products, whereas in my setting, consumers also gain from having a greater quantity of firms in the market to relax capacity constraints. It further highlights the importance of the background learning rate ϵ in our respective results. In their setting, background learning is important to generate initial information that allows the designer to send false recommendations, whereas in my setting, it is the sole source of reputational dynamics under the optimal tenure policy.

5 Prices

The baseline model assumes that prices are fixed. This creates two separate inefficiencies. First, firms cannot internalize excess demand, leading to rationing, which is inefficient since consumers

[19] In a previous version of the paper (available on my website), I explore a variant of the baseline model wherein voluntary entry as well as rationing are mechanically shut down, and find conditions under which full transparency dominates the optimal tenure policy identified in Theorem 2.

weakly prefer consumption at any firm over not consuming. Second, firms cannot internalize the future social value of information generated by consumers. To explore the role of these inefficiencies, I now introduce flexible pricing into the model by employing the approach of the competitive search literature (Eeckhout and Kircher, 2010). The underlying matching technology remains unchanged, to facilitate comparison with the baseline model. I will demonstrate that as continuation values for firms are now convex, full transparency is optimal.

The extensive form at each instant is now as follows. First, firms simultaneously choose which prices $q \in \mathbb{R}$ to supply at, if at all. Let the function $s : [0,1] \times \mathbb{R} \to [0, \bar{\pi}]$ denote this supply function, where $s(p, q)$ specifies the supply rate given a rating p and price q.[20] Whereas in the baseline model, the supply choice was trivial, as firms bore no cost of service, here the supply choice will depend on prices, and thus the problem must be carefully specified. Capacity-constrained service requires that at each rating p,

$$\int s(p, q)\,\mathrm{d}q \leq \bar{\pi}. \tag{15}$$

Sub-markets now comprise of pairs (p, q) specifying the rating of the firms and the trading price.

Next, observing the choices of firms, consumers simultaneously chose at which sub-markets (p, q) to direct their search. A directed search strategy G is a distribution on $[0, 1] \times \mathbb{R}$ that is absolutely continuous with respect to $s \cdot F$, so that consumers can only direct their search to sub-markets where there exist firms of rating p that choose to supply at price q. Market tightness $\tau(p, q) = dG(p, q)/(s(p, q)dF(p))$ now comprises the ratio of demand $dG(p, q)$ for firms with rating p charging price q with supply $s(p, q)dF(p)$. Consumers are served in sub-market (p, q) with probability $\Theta(p, q)$, where:

$$\Theta^{-1}(p, q) = \max\left\{\tau(p, q), 1\right\}.$$

Conditional upon being served within sub-market (p, q), consumers receive an expected payoff p but pay price q, so that incentive compatibility of the directed search strategy G thus requires

$$(p - q)\Theta(p, q) = (p' - q')\Theta(p', q') \text{ for all } (p, q), (p', q') \in \text{ supp}(G), \tag{16}$$

[20]It is without loss to assume firms set supply as function of current rating p alone, since each firm is small and consumers are short-lived with payoffs fully defined by p.

with the $\Theta(p,q)$ extended beyond $\text{supp}(G)$ as before by the market utility condition

$$\Theta(p,q) = \inf_{\theta\in[0,1]}\{\theta \mid (p-q)\theta \geq \sup_{(p',q')\in\text{supp}(G)} (p'-q')\Theta(p',q')\}. \tag{17}$$

Returning to the supply choice of firms, the rate that a firm with rating p trading at price q and supplying at rate s serves at is

$$\pi(p,q) = s\min\{\tau(p,q),1\},$$

analogous to equation (12) in the baseline model. This service rate is now jointly determined by the firm's supply choice s, as well as the strategies of consumers and other firms, which enter through the market tightness $\tau(p)$. Let

$$\hat{\pi}(p) = \int q\pi(p,q)\,\mathrm{d}q$$

denote the total flow profit a firm with rating p receives. Optimality of the supply function then requires that it maximize the present value of a firm:

$$V(p) = \sup_{s\in\mathcal{S}} \mathbb{E}^x\left[\int_0^\infty e^{-(\rho+\delta)t}\hat{\pi}(p_t)\,\mathrm{d}t \mid p_0 = p\right], \tag{18}$$

such that $(p_t)_{t\geqslant 0}$ follows the law of motion (2) and where $\mathcal{S}$ is the set of $[0,\bar{\pi}]$-valued integrable functions on $[0,1]\times\mathbb{R}$. Thus, firms make trading decisions based on their entire present value rather than simply their flow profit. The equilibrium market-clearing condition now ensures that overall supply meets demand:

$$\int \pi(p,q)\,\mathrm{d}G(p,q) = 1. \tag{19}$$

Definition 4. A stationary equilibrium with prices consists of distributions F on $[0,1]$, G on $[0,1]\times\mathbb{R}$, a function $s:[0,1]\times\mathbb{R}\to[0,\bar{\pi}]$ and a scalar η such that: 1) given $\{G,s,\eta\}$, F satisfies the conditions given in Proposition 1; 2) given $\{F,s,\eta\}$, G satisfies equation (16); 3) given $\{F,G,\eta\}$, s satisfies equations (15) and (18), and; 4) given $\{F,G,s\}$, η satisfies condition (4).

A simple modification to the assumption that $K<\bar{V}$ ensures a positive rate of entry under full transparency:

Assumption 1. $K < \frac{p_0 \bar{\pi}}{\rho + \delta}$.

To understand the sufficiency of this condition, consider the present value of an entering firm guaranteed to sell out forever at the highest feasible prices. Given that consumers will never pay more than the firm's expected quality of output, prices are bounded above by the firm's current rating at each instant. Since ratings follow a martingale, so too do these maximal prices, and hence this present value is bounded above by $p_0 \bar{\pi}/(\rho + \delta)$.

5.1 Equilibrium Under Full Transparency

To begin, the following result shows that rationing cannot occur in equilibrium, and further provides a sharp characterization of the prices that trade occurs at.

Lemma 3. The equilibrium consumer trading probability $\Theta(p, q) = 1$ for all $(p, q) \in \text{supp}(G)$. Furthermore, there exists $w \geqslant 0$ such that if $(p, q) \in \text{supp}(G)$ then $q = p - w$. Consumer welfare equals w.

The market utility (MU) assumption, combined with the piecewise-linear matching, tells us that were a firm with rating p trading at price q and facing excess demand ($\Theta(p, q) < 1$), they would deviate to trade at a higher price q' that partially alleviates consumer rationing in a manner that keeps consumers happy, while clearly benefiting the firm. As such, rationing will cease to exist in equilibrium. An immediate implication of the consumer indifference condition (16) is that in equilibrium, on-path trading prices are affine in rating.

Taking these prices $q(p) \equiv p - w$, standard methods allow me to write the value function of a firm $V(p)$ as a function of their rating alone and in differential form:

$$V(p) = \max_{s \in [0, \bar{\pi}]} \left\{ \frac{1}{\rho + \delta} \left[(p - w)s + (s + \epsilon) \frac{p^2 (1 - p)^2}{2\sigma^2} V''(p) \right] \right\}, \tag{20}$$

from which it is clear that optimal service rates are bang-bang, i.e. there exists a $\tilde{p}$ such that $s(p, q) = \bar{\pi}$ if $p \geqslant \tilde{p}$ and $q = q(p)$ and $s(p, q) = 0$ otherwise. I now provide the analogous result of Theorem 1 by fully characterizing the unique stationary equilibrium.

Proposition 2. There exists a unique stationary equilibrium with prices $E = \{F, G, s, \eta\}$, featuring:

1. A positive rate of entry ($\eta > 0$);
2. No rationing ($\Theta(p, q(p)) = 1$ for all $p \in [0, 1]$);
3. Prices that increase linearly with ratings ($q(p) = p - w, w > 0$);
4. Negative prices for the lowest-rated active firms ($w > \tilde{p}$);
5. Globally strictly convex firm value function ($V''(p) > 0$ for all $p \in (0, 1)$).

As discussed previously, at ratings where firms sell ($p \geq \tilde{p}$), prices are linear and firms sell out, so that flow profits are linear. The threshold $\tilde{p}$ is precisely the rating at which firms are indifferent between selling and not. This profile generates a piecewise linear flow profit function $\hat{\pi}(p) = (p - w)\mathbb{I}_{p \geq \tilde{p}}$, which in turn generates a strictly convex value function $V(p)$. Thus, firms at all ratings have positive option value of learning, and thus would rather acquire new reviews. To understand why negative prices occur for low-rated firms, note that from equation (20), the rating $\tilde{p}$ at which a firm is indifferent between supplying (at price $\tilde{p} - w$) and not satisfies

$$-(\tilde{p} - w) = \bar{\pi} \frac{p^2(1-p)^2}{2\sigma^2} V''(\tilde{p}).$$

Since $V''(\tilde{p}) > 0$, it must be that $\tilde{p} - w < 0$, so that the price at which a firm stops selling is negative. Intuitively, the equation above equates the flow loss from negative prices against the strictly positive option value of learning. In short, flexible prices eliminate the non-convexity that drives the economics of both Theorems 1 and 2, and thus within the space of simple policies, full transparency is thus optimal. Taken together, Proposition 2 and Theorem 2 further our understanding of how price rigidity impacts the optimal design of rating systems in the presence of endogenous entry by firms.

6 Model Discussion

I now provide some discussion of the model's elements. The parameter ϵ represents a mass of consumers who do not use ratings to direct their search. While I normalize the profits generated by these consumers to zero and ignore their impact on capacity constraints, the model could easily be extended to allow these consumers to generate positive profits and take up capacity without changing the results. In particular, if profits below p^* were equal to ϵ, and the capacity constraint was relaxed to include the flow of consumers ϵ, firms' value functions would still be S-shaped and the results would remain unchanged. From a technical perspective, a strictly positive ϵ is required to ensure that the entry decision is well-behaved. In short, were $\epsilon = 0$, $V(p)$ would be discontinuous at p^* taking only the two values $\{0, \bar{V}\}$, which would in turn make the entry condition $V(p_0) = K$ undefined, and thus would make the existence and uniqueness of stationary equilibrium hard to verify. The parameter δ captures exogenous attrition (exit) by firms. It is required to ensure existence of a bounded steady-state; were $\delta = 0$, an infinite measure of firms would attain the ratings $0, 1$ with the density f not supported at any other rating. By separating the informativeness of feedback $\alpha > 0$ and the capacity constraint $\bar{\pi}$, my analysis is able to capture settings wherein capacity constraints are weak (high $\bar{\pi}$) so that only a few firms are needed to cover the market, but where individual reviews are sufficiently uninformative (low α) that learning is still not immediate (for instance, markets for digital apps).

My baseline analysis assumes that prices are exogenous, with rationing serving instead as a market-clearing mechanism. By extending the model to allow for competitive prices, I show in Section 5 that this assumption is a crucial ingredient underlying the mechanism behind my main findings. That said, the endogenous variable p^* can be viewed as capturing the equilibrium level of competition in two separate ways. First, higher p^* transfers surplus from firms to consumers. Second, a greater number of firms in the market (through higher entry η) results in a higher value of p^*. Thus, p^* can be viewed as an inverse price index, and the rater's objective as maximizing equilibrium competition.

I assume a specific form of matching known as Leontief matching in the directed search literature. With this approach, the rate of service at firms takes one of two values in equilibrium, allowing for an explicit characterization of firm value functions. Nevertheless, it preserves the fundamental trade-off between probability of service and value present in directed search. See Eeckhout and Kircher (2010) for a survey of commonly used alternative matching functions.

Finally, simple policies offer a tractable yet rich method to explore ratings design by allowing me to retain the ergodic analysis of Section 2 while allowing the rater a broad scope to control the diffusion of information within the market. In this sense, designing simple policies bears close resemblance to the "control of variance" approach to dynamic information acquisition (Moscarini and Smith, 2001), with a key difference being that rather than being costly to acquire, information is itself endogenously constrained by consumer choice. That said, they entail a restriction on the scope of ratings design, in so far as they require 1) the rater to not acquire information privately, and 2) the choice variable r to only depend upon the public record $(X_t)_{t \geqslant 0}$ via the rating p. They thus rule out policies that store and release acquired reviews with any form of delay.

7 Conclusion

In this paper, I study whether consumer reviews can pose a barrier to entry for new firms. To this end, I build an equilibrium model in which firms of heterogeneous quality make entry choices, and whose quality is gradually revealed via consumer reviews. A rater acquires these reviews and uses them to create ratings for firms, which in turn determine consumer demand. The central insight of the paper is that policies that limit reviews for well-established firms can stimulate entry incentives and ultimately boost consumer welfare.

In an attempt to shed light solely on the interplay between entry and ratings design, my analysis shuts down various features that might both interact with my results and play an important role in applications. Firms often voluntarily exit the market due to fixed operating costs and dwindling revenues.[21] Firms often exert effort or invest in quality, which I assume is fixed. I also assume that firms do not know their quality prior to entry. Were this not the case, entry would exhibit ex-ante selection effects (Atkeson et al., 2015). Simple policies rule out a rich class of policies wherein the rater can privately acquire information (Hörner and Lambert, 2021).[22] My analysis is performed in steady state, in order to utilize techniques from ergodic analysis and mean-field games. It would no doubt be instructive to perform a non-stationary analysis. Finally, I assume that all consumers freely provide unbiased feedback, abstracting from the challenge of soliciting reviews faced by online review platforms.[23] Undoubtedly, adding these various elements would be an essential step in order to bring my analysis closer to various applications of interest, and thus I view my analysis as taking a first step toward a greater understanding of ratings design at the industrial level.

[21] I explore this variant in an earlier version of the paper available on my website. The main result, Theorem 2, remains unchanged.

[22] I analyze a special case of such policies, called *private certification* policies, in an earlier version of the paper available on my website.

[23] See Bénabou and Vellodi (forthcoming) for recent work that explores this idea.

A Proofs

A.1 Proposition 1

First, note that the process governing the evolution of ratings for a firm can be re-cast as a *resetting* process by adding an additional state in which the firm enters upon exit, and transitions out of (entry) at rate η (see Luttmer (2007) for an identical approach). In this way, it is (strongly) Markovian, positive recurrent, and irreducible, and thus admits a unique invariant distribution. Next, I derive the boundary conditions for the Fokker-Planck equation (9), drawing upon techniques from the adjoint theory of differential operators to derive boundary conditions. See Gabaix et al. (2016) for an economic application of this approach, and (Gardiner, 2009, Chapter 5) more generally.

Let $X \in \mathbb{B}(\mathbb{R})$ and for two functions $u, v \in \mathcal{L}^2(X)$, define their inner product as $\langle u, v \rangle = \int_X u(x)v(x)dx$. Further, for an operator $\mathcal{A}$, the *adjoint operator* is defined as $\mathcal{A}^*$ such that $\langle u, \mathcal{A}v \rangle = \langle \mathcal{A}^* u, v \rangle$. For a diffusion process Y satisfying $dY_t = a(Y,t)dt + b(Y,t)dW_t$ for an appropriately defined Wiener process W with constant hazard-rate of death δ, the *infinitesimal operator* is given by

$$\mathcal{A}_b\{u\}(x,t) = a(x,t)\frac{\partial u}{\partial x}(x) + \frac{1}{2}b(x,t)\frac{\partial^2 u}{\partial x^2}(x) - \delta u(x).$$

Finally, the operator

$$\mathcal{J}\{u\}(x,t) = a(x,t)u(x) - \frac{\partial}{\partial x}(b(x,t)u(x))$$

denotes the *mass flux*, i.e. for $S \subset \mathbb{R}$, the integral $\int_{\partial S} \mathcal{J}\{f\}(x,t)$ measures the total mass crossing the boundary of S per unit time (to see this, integrate the non-stationary version of (9) directly and use the Fundamental Theorem of Calculus).

I begin with the case where the policy r is such that $r(p) > 0$ for all $p \in [0,1]$. Thus, for the ratings process defined by the SDE (2),

$$(\mathcal{A}_b\{u\})(p,t) = \frac{1}{2}\Sigma(p)\frac{\partial^2 u}{\partial p^2} - \delta u(p) \tag{A.1}$$

$$\mathcal{J}\{u\}(p,t) = -\frac{\partial}{\partial p}\left[\Sigma(p)u(p)\right]. \tag{A.2}$$

Standard results in statistical mechanics then imply that the transition measures form a root of the adjoint operator to (A.1).

Remark 1. The stationary distribution satisfies $\mathcal{A}_f f = 0$, where $\mathcal{A}_f = \mathcal{A}_b^*$.

This verifies that f must solve (9). It remains to derive the boundary conditions for f. To do so, we state the relevant boundary conditions for the equation $\mathcal{A}\{u\} = 0$. Standard results tell us that the operator $\mathcal{A}$ is well-behaved, i.e. that solutions to $\mathcal{A}\{u\} = 0$ lie in $\mathcal{C}^2$. We use this and the adjoint relation in Remark 1 to transform these into conditions on f.

Lemma A.1. Let $\tilde{\mathcal{D}}$ denote the set of discontinuities of f, and let $\mathcal{D} = \tilde{\mathcal{D}} \cup \{\underline{p}, p_0, 1\}$. Then

$$\int_{\mathcal{D}} \left[u(p)\mathcal{J}\{f\}(p) + f(p)\Sigma(p)\frac{\partial u}{\partial p} \right] dp = 0 \tag{A.3}$$

for all $u \in \mathcal{C}^2([0,1])$.

Proof.

$$\begin{aligned}
\langle \mathcal{A}^* f, u \rangle &= -\int_0^1 u \left[(\Sigma f)'' - \delta f \right] \mathrm{d}p \\
&= \int_0^1 \left[(u(\Sigma f)')' - \Sigma f u' - \delta f u \right] \mathrm{d}p + \int_{\mathcal{D}} u(\Sigma f)' - f\Sigma u_p \, \mathrm{d}p \\
&= \int_0^1 f \left[\Sigma u'' - \delta u \right] \mathrm{d}p + \int_{\mathcal{D}} u \left[(\Sigma f)' - f\Sigma u_p \right] \mathrm{d}p \\
&= \langle f, \mathcal{A}u \rangle - \int_{\mathcal{D}} u \left[\mathcal{J}\{f\} + f\Sigma u_p \right] \mathrm{d}p,
\end{aligned}$$

where the second equality follows by the Fundamental Theorem of Calculus. The result then follows by Remark 1. □

Armed with Lemma A.1, I now derive conditions 1 - 5 of Proposition 1 (condition 6 is derived later separately from an aggregate conservation of probability principle).

1. $\Sigma(1)f(1) = \Sigma(0)f(0) = 0$. We have that $f(p)\Sigma(p)u' = 0$ for all u. This is a "natural boundary" condition, as $p = 1$ cannot be attained in finite time.

2. $\Sigma(p_0)\left[f'_-(p_0) - f'_+(p_0)\right] = \eta$, $\frac{\partial}{\partial p}\Sigma(p)f(p) \in \mathcal{C}^1([0,p_0) \cup (p_0,1])$, $\Sigma(p)f(p) \in \mathcal{C}^1([0,1])$. An inflow rate of η yields the boundary condition $[\Sigma(p_0)u(p_0)]_-^+ = \eta$ on any u that solves $\mathcal{A}_f\{u\} = 0$.

Furthermore, since u and u' are continuous, this implies that:

$$
\begin{aligned}
[\mathcal{J}\{f\}(p_0)]_-^+ &= \eta \\
-[\Sigma(p_0)f'(p_0) + \Sigma'(p_0)f(p_0)]_-^+ &= \eta \\
\Sigma(p_0)[f'_-(p_0) - f'_+(p_0)] &= \eta.
\end{aligned}
$$

The condition states that the total outflow of mass from p_0, given by $[\mathcal{J}\{f\}(p_0)]_-^+$, must equal the total inflow η. Note that as a consequence, it also implies that f is continuous at p_0. The continuity conditions on Σf and $(\Sigma f)'$ also follow from the fact that u and u' are both continuous.

Finally, condition 6 is derived by imposing that the total mass of firms is constant in a stationary equilibrium, i.e.

$$\frac{d}{dt}\int_{\underline{p}}^{1} f(p,t)dp = 0. \tag{A.4}$$

Direct substitution of the Fokker-Planck equation into equation (A.4) implies that

$$
\begin{aligned}
[(\Sigma f)']_{\mathcal{D}} - \delta \int f \mathrm{d}p &= 0 \\
\eta - \delta \int f \mathrm{d}p &= 0.
\end{aligned}
$$

where all other terms in $[(\Sigma f)']_{\mathcal{D}}$ vanish due to continuity of $(\Sigma f)'$.

Next, suppose that $r(p) = 0$ for some $p \in (0,1)$. The ratings process attains such points in finite time, with such p being termed *sticky boundaries*. It is clearly without loss to consider only ratings for which there are at most two such points $p_1, p_2 \in (0,1)$ with $p_1 < p_0 < p_2$.

In this case the ratings distribution admits atoms at p_1, p_2, as firms are trapped until they experience exogenous death at rate δ. Since $\delta \in (0,\infty)$, such atoms have bounded measure. By using Lemma A.1, we derive an additional boundary condition to account for such atoms (Gardiner, 2009, Chapter 5.2). Formally, the ratings distribution is defined by the pair $\{f, \Omega_1, \Omega_2\}$ that satisfy:

(F_1) $\Sigma(p_0)[f'_-(p_0) - f'_+(p_0)] = \eta$,

(F_2) $\frac{\partial}{\partial p}\Sigma(p)f(p) \in \mathcal{C}^1([p_1, p_0) \cup (p_0, p_2])$,

(F_3) $\Sigma(p)f(p) \in \mathcal{C}^1([p_1, p_2])$,

(F_4) $\delta[\int_{\underline{p}}^1 f(p)\,\mathrm{d}p + \Omega_1 + \Omega_2] = \eta$,

(F_5) $\delta\Omega_1 + [\Sigma(p_1)f(p_1)]'_- = 0.$

(F_6) $\delta\Omega_2 + [\Sigma(p_2)f(p_2)]'_- = 0.$

A.2 Lemma 1

I first prove a preliminary lemma:

Lemma A.2. If G is IC with respect to F, then there exists a unique $p^* \in (0, 1]$ such that $\Theta(p^*) = 1$ and $\Theta(p) < 1$ for all $p \in (p^*, 1]$.

Proof. Suppose to the contrary that there exist p_1, p_2 in the common support of G and F, such that $p_1 \neq p_2$ and $\Theta(p_1) = \Theta(p_2) = 1$. The consumer then obtains an expected payoff of p_i by choosing to consume at p_i, and hence cannot be indifferent, since $p_1 \neq p_2$. Simply put, to maintain indifference across actively visited firms, rationing must occur in a manner that offsets differences in expected quality. The market utility assumption (7) guarantees this logic holds even outside the common support of G and F. □

To obtain indifference in (6), it must be that for any rating visited in equilibrium, rationing must (weakly) occur, i.e. $\frac{g(p)}{\bar{\pi} f(p)} \geqslant 1$ for all $p \in \text{supp}(g)$. For if not, take $p_1, p_2 \in \text{supp}(g), p_1 \neq p_2$ such that $\frac{g(p_1)}{\bar{\pi} f(p_1)} < 1$, $\frac{g(p_2)}{\bar{\pi} f(p_2)} < 1$. The consumer then obtains an expected payoff of p_i by choosing to consume at p_i, and hence cannot be indifferent, since $p_1 \neq p_2$. This verifies that $\pi(p) \in \{0, \bar{\pi}\}$. A similar argument verifies that $\pi(p)$ is increasing, and hence there exists $p^* \in [0, 1]$ such that $\pi(p) = 0$ for $p < p^*$ and $\pi(p) = \bar{\pi}$ for $p \geqslant p^*$. This also establishes the conjectured form of $g(p)$. Finally, $p^* > 0$ to maintain indifference.

A.3 Theorem 1

The outline of the proof is as follows:

1. Fix p^*, η. Compute f.

2. Define $R_*(\eta)$ as the solution to (10) subject to (11) for a given η, and $R_\eta(p^*)$ as the solution to (4) for a fixed p^*.

3. Argue that both $R_*(\eta)$ and $R_\eta(p^*)$ have precisely one intersection in (p^*, η)-space.

A.3.1 Computing the Ratings Distribution

The density f is solved for separately over three regions, the boundaries of which depend on the equilibrium value of p^*. In all cases, conditions 4 and 5 boil down to:

$$\left[\frac{\partial}{\partial p}\Sigma(p^*)f(p^*)\right]_-^+ = 0$$
$$[\Sigma(p^*)f(p^*)]_-^+ = 0$$

Case 1: $p^* > p_0$

The general solution to the Fokker-Planck equation (9) is:

$$f(p) = \begin{cases} c_0^1 p^{-1-\gamma_0^f}(1-p)^{\gamma_0^f-2} + c_0^2 p^{\gamma_0^f-2}(1-p)^{-1-\gamma_0^f} & \text{for } p \in [0, p_0] \\ c_1^1 p^{-1-\gamma_0^f}(1-p)^{\gamma_0^f-2} + c_1^2 p^{\gamma_0^f-2}(1-p)^{-1-\gamma_0^f} & \text{for } p \in [p_0, p^*) \\ c_2^1 p^{-1-\gamma_1^f}(1-p)^{\gamma_1^f-2} + c_2^2 p^{\gamma_1^f-2}(1-p)^{-1-\gamma_1^f} & \text{for } p \in [p^*, 1] \end{cases}$$

where γ_0^f, γ_1^f are as per the statement of the proposition. For $x, y \in [\underline{p}, 1]$, define:

$$\begin{aligned}
\alpha_i^1(x) &= x^{-1-\gamma_i^f}(1-x)^{\gamma_i^f-2} \\
\alpha_i^2(x) &= x^{\gamma_i^f-2}(1-x)^{-1-\gamma_i^f} \\
\psi_i^1(x,y) &= \int_x^y p^{-1-\gamma_i^f}(1-p)^{\gamma_i^f-2}\,\mathrm{d}p \\
&= \int_{\frac{x}{1-x}}^{\frac{y}{1-y}} \left(\frac{z}{1+z}\right)^{-1-\gamma_i^f}\left(\frac{1}{1+z}\right)^{\gamma_i^f-2} dz \\
&= \int_{\frac{x}{1-x}}^{\frac{y}{1-y}} z^{-1-\gamma_i^f}(1+z)dz \\
&= \frac{1}{\gamma_i^f}\left[\left(\frac{x}{1-x}\right)^{-\gamma_i^f} - \left(\frac{y}{1-y}\right)^{-\gamma_i^f}\right] + \frac{1}{\gamma_i^f-1}\left[\left(\frac{x}{1-x}\right)^{1-\gamma_i^f} - \left(\frac{y}{1-y}\right)^{1-\gamma_i^f}\right] \\
\psi_i^2(x,y) &= \int_x^y p^{\gamma_i^f-2}(1-p)^{-1-\gamma_i^f}\,\mathrm{d}p \\
&= \frac{1}{1-\gamma_i^f}\left[\left(\frac{x}{1-x}\right)^{\gamma_i^f-1} - \left(\frac{y}{1-y}\right)^{\gamma_i^f-1}\right] - \frac{1}{\gamma_i^f}\left[\left(\frac{x}{1-x}\right)^{\gamma_i^f} - \left(\frac{y}{1-y}\right)^{\gamma_i^f}\right]
\end{aligned}$$

Since $\gamma_1^f, \gamma_0^f > 1$, condition 3 immediately implies that

$$c_2^2 = c_0^1 = 0.$$

Some (omitted) algebra then allows me to transform the remaining boundary conditions into a system of independent equations for the remaining undetermined coefficients from the boundary conditions.

Lemma A.3. The coefficients $\mathbf{c} = [c_0^2\ c_1^1\ c_1^2\ c_2^1]^T$ solve the linear algebraic system $\mathbf{Ac} = \mathbf{b}$, where

$$\mathbf{A} = \begin{bmatrix} -\frac{\partial \alpha_0^1}{\partial p}(p_0) & -\frac{\partial \alpha_0^2}{\partial p}(p_0) & 0 & 0 \\ 0 & \frac{\partial}{\partial p}\Sigma_0(p^*)\alpha_0^1(p^*) & \frac{\partial}{\partial p}\Sigma_0(p^*)\alpha_0^2(p^*) & -\frac{\partial}{\partial p}\Sigma_1(p^*)\alpha_1^1(p^*) \\ 0 & \Sigma_0(p^*)\alpha_0^1(p^*) & \Sigma_0(p^*)\alpha_0^2(p^*) & -\Sigma_1(p^*)\alpha_1^1(p^*) \\ \delta\psi_0^2(\underline{p}, p_0) & \delta\psi_0^1(p_0, p^*) & \delta\psi_0^2(p_0, p^*) & \delta\psi_1^1(p^*, 1) \end{bmatrix}$$

$$\mathbf{b} = \begin{bmatrix} \eta \\ 0 \\ 0 \\ \eta \end{bmatrix}$$

Case 2: $p^* < p_0$

The general solution to the Fokker-Planck equation (9) now becomes

$$f(p) = \begin{cases} c_0^1 p^{-1-\gamma_0^f}(1-p)^{\gamma_0^f - 2} + c_0^2 p^{\gamma_0^f - 2}(1-p)^{-1-\gamma_0^f} & \text{for } p \in [0, p^*] \\ c_1^1 p^{-1-\gamma_1^f}(1-p)^{\gamma_1^f - 2} + c_1^2 p^{\gamma_1^f - 2}(1-p)^{-1-\gamma_1^f} & \text{for } p \in [p^*, p_0) \\ c_2^1 p^{-1-\gamma_1^f}(1-p)^{\gamma_1^f - 2} + c_2^2 p^{\gamma_1^f - 2}(1-p)^{-1-\gamma_1^f} & \text{for } p \in [p_0, 1] \end{cases}$$

As before, $c_2^2 = c_0^1 = 0$, and so the coefficients solve $\mathbf{Ac} = \mathbf{b}$, where

$$\mathbf{A} = \begin{bmatrix} 0 & \frac{\partial \alpha_1^1}{\partial p}(p_0) & \frac{\partial \alpha_1^2}{\partial p}(p_0) & -\frac{\partial \alpha_1^1}{\partial p}(p_0) \\ \frac{\partial}{\partial p}\Sigma_0(p^*)\alpha_0^2(p^*) & -\frac{\partial}{\partial p}\Sigma_1(p^*)\alpha_1^1(p^*) & -\frac{\partial}{\partial p}\Sigma_1(p^*)\alpha_1^2(p^*) & 0 \\ \Sigma_0(p^*)\alpha_0^2(p^*) & -\Sigma_1(p^*)\alpha_1^1(p^*) & -\Sigma_1(p^*)\alpha_1^2(p^*) & 0 \\ \delta\psi_0^2(\underline{p}, p^*) & \delta\psi_0^1(p^*, p_0) & \delta\psi_0^2(p^*, p_0) & \delta\psi_1^1(p_0, 1) \end{bmatrix}$$

and $\mathbf{b}$ is as in the previous case.

A.3.2 Computing the Firm's Value Function

Standard results (Rüschendorf and Urusov, 2008) confirm that the firm's value function belongs to $\mathcal{C}^1([0,1])$ and satisfies the differential equation

$$V(p) = \frac{1}{\rho+\delta}\left[\bar{\pi}\mathbb{I}_{p\geq p^*} + \Sigma(p)V''(p)\right]$$

on $[0,1]$. The boundary conditions reduce to:

$$V(0) = 0, \quad V_-(p^*) = V_+(p^*), \quad V(1) = \frac{\bar{\pi}}{\rho+\delta}$$

The general and particular solutions to this equation are:

Proposition A.1. Let

$$\gamma_0^v = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2\sigma^2(\rho+\delta)}{\epsilon}}; \quad \gamma_1^v = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2\sigma^2(\rho+\delta)}{\epsilon+\alpha\bar{\pi}}}.$$

Then a firm's value function is given by

$$V(p) = \begin{cases} c_1^0 p^{1-\gamma_1^v}(1-p)^{\gamma_1^v} + \frac{\bar{\pi}}{\rho+\delta} & \text{if } p \geqslant p^* \\ c_0^1 p^{\gamma_0^v}(1-p)^{1-\gamma_0^v} & \text{if } p < p^*. \end{cases} \tag{A.5}$$

where the coefficients $\begin{bmatrix} c_0^1 & c_1^0 \end{bmatrix}' = \mathbf{c_v}$ solve the linear algebraic system $\mathbf{\Theta_v c_v} = \mathbf{b_v}$, where

$$\mathbf{\Theta_v} = \begin{bmatrix} p^{*\gamma_0}(1-p^*)^{1-\gamma_0} & p^{*1-\gamma_1}(1-p^*)^{\gamma_1} \\ \left[\frac{\gamma_0}{p^*} - \frac{1-\gamma_0}{1-p^*}\right] p^{*\gamma_0}(1-p^*)^{1-\gamma_0} & -\left[\frac{1-\gamma_1}{p^*} - \frac{\gamma_1}{1-p^*}\right] p^{*1-\gamma_1}(1-p^*)^{\gamma_1} \end{bmatrix}$$

and

$$\mathbf{b_v} = \begin{bmatrix} \frac{\bar{\pi}}{\rho+\delta} \\ 0 \end{bmatrix}.$$

A.3.3 Reaction Functions

Denote by $R_*(\eta)$ the solution to (10) subject to (11) for a given η, and $R_\eta(p^*)$ the solution to (4) for a fixed p^*. That is, $R_*(\eta)$ computes optimal consumer thresholds for fixed η, while $R_\eta(p^*)$ computes optimal entry rates η for fixed p^*.

Lemma A.4. 1. $R_*(\eta) \in \mathcal{C}^1([0,\infty))$ and strictly increasing.

2. $\lim_{\eta\to 0} R_*(\eta) = 0$, $\lim_{\eta\to\infty} R_*(\eta) = 1$.

Proof. 1. By Lemma A.3, the coefficients determining f solve $\mathbf{Ac} = \eta\mathbf{1}$. Hence, f is linearly homogeneous in η. Denote by $f(p;\eta)$ the explicit dependence of the density on η. By equation (11),

$$R_*(\eta) = \eta\bar{\pi}\int_{p^*}^{1} pf(p;1)\,\mathrm{d}p. \tag{A.6}$$

Implicit differentiation yields

$$\frac{dR_*(\eta)}{d\eta} = \underbrace{\left[\bar{\pi}\int_{p^*}^{1} pf(p;1)\mathrm{d}p\right]^{-1}}_{>0}\underbrace{\left[1 - \eta\bar{\pi}\frac{d}{dp^*}\left[\int_{p^*}^{1} pf(p;1)\mathrm{d}p\right]\right]}_{>0}.$$

2. By linear homogeneity of $\eta \mapsto f(p;\eta)$, $\lim_{\eta\to 0} f(p;\eta) = 0$ for all $p \in [0,1]$ and $\lim_{\eta\to\infty} f(p;\eta) = \infty$. The result then follows from equation (A.6).

□

Lemma A.5. There exists a unique p_F^* such that

$$R_\eta(p^*) = \begin{cases} \infty & \text{if } p^* > p_F^* \\ \in [0,\infty] & \text{if } p^* = p_F^* \\ 0 & \text{if } p^* < p_F^*. \end{cases} \tag{A.7}$$

Proof. Let $V(p;p^*)$ denote the firm value function with explicit dependence on p^*. By Proposition A.1, $p^* \mapsto V(p;p^*)$ is continuous and strictly increasing, with $\lim_{p^*\to 0} V(p;p^*) = \frac{\bar{\pi}}{\rho+\delta}$ and $\lim_{p^*\to 1} V(p;p^*) = 0$ for all $p \in [0,1]$. The result then follows since $K < \frac{\bar{\pi}}{\rho+\delta}$. □

Combining Lemmas A.4 and A.5 proves the Theorem.

A.4 Lemma 2

First, since simple ratings again induce strongly Markovian, positive recurrent, and irreducible ratings processes, the ratings distribution exists and is unique. As such, the reasoning in Lemma 1 applies, verifying the first part.

Given this structure, the firm's continuation value will again be of the form:

$$V(p) = \frac{1}{\rho + \delta}\left[\bar{\pi}\mathbb{I}_{p \geqslant p_E^*} + r(p)\Sigma(p)V''(p)\right].$$

Note that again $\frac{\bar{\pi}}{\rho+\delta}$ forms an upper bound for $V(p)$ and so for $p \in (0, p_E^*))$,

$$V''(p) = \frac{\rho + \delta}{r(p)\Sigma(p)}V(p) \geqslant 0,$$

since $V(p) \geqslant 0$ by construction, implying V is convex for $p < p_E^*$. On the other hand, for $p \in [p_E^*, 1]$,

$$V''(p) = \frac{\rho + \delta}{r(p)\Sigma(p)}\left[V(p) - \frac{\bar{\pi}}{\rho + \delta}\right] \leqslant 0,$$

implying V is concave.[24]

Finally, to prove that $\eta \in (0, \infty)$, suppose that $\eta = 0$. Then equation (11) implies that $p^* = 0$, so that $V(p_0) = \bar{V}$ and thus $\eta = \infty$ by equation (4), a contradiction. Similarly, if $\eta = \infty$, then equation (11) implies that $p^* = 1$, so that $V(p_0) = 0$ and thus $\eta = 0$ by equation (4), a contradiction. Finally, to see that p_E^* equals consumer welfare, note that at this rating, consumers are guaranteed service, and thus their expected value is simply p_E^*, while the indifference condition in equation (11) ensures that expected values are identical across all sub-markets visited in equilibrium. This argument holds for an arbitrary ratings distribution, and thus in particular at any equilibrium for any simple rating.

A.5 Theorem 2

I begin by showing that tenure policies admit a unique equilibrium, and moreover that there is a unique $q \in (0, 1)$ such that tenuring at q leads to a stationary distribution f with an atom of mass

[24] Note that $r(p) \neq 0$ almost everywhere by the discussion in Proposition 1.

$1/\bar{\pi}$ at q.

Lemma A.6. All tenure policies admit a unique equilibrium. Furthermore, there exists a unique $q \in (p_0, 1)$ such that if $\tilde{p} > q$ then $p^*(r) < \tilde{p}$ and firms sell out ($\pi(p) = \bar{\pi}$) at $p \geqslant p^*(r)$ whereas if $\tilde{p} \leqslant q$ then $p^*(r) = \tilde{p}$ and firms do not sell out $\pi(p) < \bar{\pi}$ at $p \geqslant p^*(r)$.

Proof. I will explicitly demonstrate the existence of a unique equilibrium under tenure policies. Consider first policies with $\tilde{p} \leqslant p_0$, which are clearly fully uninformative since $r(p_0) = 0$. The ratings distribution in this case involves a point mass at p_0 of endogenous size $\Omega > 0$, say. Market tightness is then given simply by

$$\Theta = \min\{\bar{\pi}\Omega, 1\},$$

and firms' value functions satisfy

$$V = \frac{\min\{\Omega^{-1}, \bar{\pi}\}}{\rho + \delta}.$$

The free entry condition (4) and the assumption that $K < \frac{\bar{\pi}}{\rho+\delta}$ imply then that $\bar{\pi}\Omega > 1$, and thus that $\Theta = 1$ and so consumer welfare is p_0. The mass Ω is pinned down by the entry condition:

$$\frac{\Omega^{-1}}{\rho + \delta} = K.$$

Finally, the entry rate η is pinned down by the balance equation for Ω:

$$\eta = \delta\Omega.$$

Next, consider tenure policies with $\tilde{p} > p_0$. In such cases, firms' value functions $V(p)$ are given by Equation (A.5) for $p \in [0, \tilde{p}]$, and

$$V(\tilde{p}) = \frac{\min\{\Omega^{-1}, \bar{\pi}\}}{\rho + \delta},$$

with continuity at $\tilde{p}$ and where Ω is now the mass of firms at $\tilde{p}$. The ratings distribution follows the expressions in Appendix A.3.1 on $[0, \tilde{p})$ and

$$\delta\Omega = [\Sigma(\tilde{p})f(\tilde{p})]'_{-}. \tag{A.8}$$

Consider the function V that solves

$$V(p) = \frac{1}{\rho+\delta}\left[\Sigma(p)V''(p)\right] \tag{A.9}$$

on $[0, p_0]$ with $V(0) = 0, V(p_0) = K$. Then V has unique solution as defined in the first clause of equation (A.5), and is thus strictly increasing and convex. Consider the analytic continuation V_{ext} of V on $[0, 1]$. Then there exists a unique $q > p_0$ such that $V_{ext}(q) = \bar{V}$. We proceed in two cases:

Case 1: $\tilde{p} > q$ – By construction of q, it must be that in any equilibrium, $p^* < \tilde{p}$, since otherwise, the value function is such that $V(p_0) < K$. Thus, V solves (3) on $[0, \tilde{p}]$ with $V(\tilde{p}) = \bar{V}$. Thus, as per the same argument in Proposition A.1, V is strictly decreasing in p^*. Hence, the analysis of the reaction function $p^*_E(r) \mapsto R_\eta(p^*_E(r))$ is unchanged from Lemma A.5.

To check that $\eta_E(r) \mapsto R^*(\eta_E(r))$ is increasing, note that the additional boundary condition $\delta\Omega + [\Sigma(\tilde{p})f(\tilde{p})]'_- = 0$ does not involve η, and hence both f on $[0, \tilde{p})$ and Ω are linearly homogeneous in η by the same argument as in Lemma A.3. Thus $p^*_E(r) \mapsto R_\eta(p^*_E(r))$ is increasing by the same argument as in Theorem 1. The analysis proceeds then as in Theorem 1, and equilibrium is unique. (Note that since f is not supported above $\tilde{p}$, $R_\eta(p^*_E(r)) = \tilde{p}$ for η sufficiently high, but this domain is irrelevant since the intersection must occur at $p^* < \tilde{p}$.)

Case 2: $\tilde{p} < q$ – I first claim then that $V(\tilde{p}) < \bar{V}$. For if not, by construction of q, we would have that $V(p_0) > K$. Since $V(\tilde{p}) < \bar{V}$, it must be that $\pi(\tilde{p}) < \bar{\pi}$, and hence that $p^* = \tilde{p}$ and furthermore that the mass of firms Ω at $\tilde{p}$ is such that $\bar{\pi}\Omega > 1$ by equation (12).

I claim that for fixed $\tilde{p} < q$, there exists a unique $\tilde{\pi} \equiv \pi(\tilde{p}) \in (0, \bar{\pi})$ such that $V(p_0) = K$. To see this, note that the value function V solves equation (A.9) on $[0, \tilde{p}]$ with $V(0) = 0, V(\tilde{p}) = \frac{\pi(\tilde{p})}{\rho+\delta}$. Using the derivation in Proposition A.1, simple algebra confirms that V is strictly increasing in $\pi(\tilde{p})$, and hence there exists a unique $\tilde{\pi} \in (0, \bar{\pi})$ such that $V(p_0) = K$, and furthermore that $\tilde{\pi}$ is strictly increasing in $\tilde{p}$ for $\tilde{p} < q$. This then implies that

$$R_\eta(p^*) = \begin{cases} \infty & \text{if } p^* > \tilde{p} \\ \in [0, \infty] & \text{if } p^* = \tilde{p} \\ 0 & \text{if } p^* < \tilde{p}, \end{cases} \tag{A.10}$$

analogous to Lemma A.5. Finally, by the linear homogeneity of f and Ω in η, there exists a unique $\eta \in (0, \infty)$ such that $\Omega^{-1} = \tilde{\pi}$, completing the proof of uniqueness.

□

Next, I argue that the policy $r(p) = \mathbb{I}_{p<q}$ is optimal within the class of tenure policies, by proving explicitly that consumer welfare is strictly single-peaked in the tenure rating $\tilde{p}$ at q.

Lemma A.7. The policy $r(p) = \mathbb{I}_{p<q}$ is the optimal tenure policy.

Proof. Note that by Lemma A.6, any tenure policy with $\tilde{p} \leqslant q$ leads to consumer welfare equal to $\tilde{p}$, and thus in the region $\tilde{p} \in [p_0, q]$, consumer welfare is linearly increasing with $\tilde{p}$.

I now argue that p^* is strictly decreasing in $\tilde{p}$ for $\tilde{p} > q$. By Lemma A.6, such policies admit a unique equilibrium in which $p^* < \tilde{p}$. Suppose to the contrary that there exist $\tilde{p}_1, \tilde{p}_2 \in (q, 1)$ such that $q < \tilde{p}_1 < \tilde{p}_2$ and $p_1^* \leq p_2^*$. By Lemmas A.1 and A.7, the value function V_1 is strictly decreasing in both p^* and $\tilde{p}$, and similarly for V_2. Thus, it must be that $V_1(p) > V_2(p)$ for all $p \leq \tilde{p}_2$, and in particular $V_1(p_0) > V_2(p_0)$, since $p_0 < \tilde{p}_2$. Since $p_1^* < 1$, it cannot be that $V_1(p_0) > K$ by Lemma A.5, and hence $V_1(p_0) = K, V_2(p_0) < K$, in which case $\eta_2 = 0$ and thus $p_2^* = 0$, a contradiction. □

Finally, I argue that any non-tenure policy is strictly dominated by $r(p) = \mathbb{I}_{p<q}$. Take an equilibrium $E = \{p^*, \eta, f\}$ with associated value function V and a non-tenure policy r such that r implements E. By Lemma 2 part 4, $V(p_0) = K$ in any implementable equilibrium, and in particular under E. Furthermore, $r(p) = \mathbb{I}_{p<q}$ supports consumer welfare $q > p_0$ by Lemma A.6 and so it is without loss to assume that $p^* > p_0$.

Suppose to the contrary that $p^* \geq q$, and let V_q denote the value function under $r(p) = \mathbb{I}_{p<q}$. I argue that $V_q > V$ pointwise on $(0, 1)$. To see this, assume first that $p^* = q$. Lemma 2 implies that both V and V_q are strictly convex and increasing on $[0, p^*]$. Consider the function $V_\dagger$ that solves

$$V_\dagger(p) = \frac{1}{\rho + \delta} \left[\Sigma(p) V''(p) \right] \tag{A.11}$$

on $[0, p^*]$, with $V_\dagger(p^*) = V(p^*)$. Then setting $u(z) = (\pi(z))/(\rho + \delta)$ in Escudé and Sinander, 2023, Corollary 2 implies that $V_\dagger \geq V$ on $[0, p^*]$, while $V_q > V_\dagger$ pointwise on $(0, p^*]$ since $V_q(p^*) = \bar{V} > V(p^*)$ and both solve (A.11) on $[0, p^*]$. Finally, $V_q > V$ on $[p^*, 1)$ since $V < \bar{V} = V_q$ on this region.

Since $V(p_0) = K$, it must then be that $V_q(p_0) > K$ and thus $\eta_q = \infty$, a contradiction to Lemma 2 part 4.

A.6 Lemma 3

Suppose there exists $(p, q) \in \text{supp}(G)$ such that $\Theta(p, q) < 1$. Let $\Theta(p, q) = 1 - \psi$ for some $\psi > 0$. I claim there exists $q' > q$ such that $(p - q') > (p - q)\Theta(p, q)$, i.e. such a price offers consumers a strictly higher expected value. To see this, let $q' = q + \chi$. Then:

$$\begin{aligned} (p - q') &> (p - q)\,\Theta(p, q) \\ p - (q + \chi) &> (p - q)(1 - \psi) \\ \chi &< (p - q)\,\psi, \end{aligned}$$

which holds for sufficiently small χ (note that $(p, q) \in \text{supp}(G)$ implies $p - q > 0$). Thus $\Theta(p, q) < \Theta(p, q') \leqslant 1$. Finally, note that such a price offer would constitute a profitable deviation for a firm with rating p, since $\Theta(p, q') \leqslant 1$ implies that $\pi(p, q') = \bar{\pi}$, and hence:

$$\begin{aligned} V(p, q') &= \frac{1}{\rho + \delta}\left[q'\pi(p, q') + (\pi(p, q') + \epsilon)\frac{p^2(1-p)^2}{2\sigma^2}V''(p)\right] \\ &= \frac{1}{\rho + \delta}\left[q'\bar{\pi} + (\bar{\pi} + \epsilon)\frac{p^2(1-p)^2}{2\sigma^2}V''(p)\right] \\ &> \frac{1}{\rho + \delta}\left[q\bar{\pi} + (\bar{\pi} + \epsilon)\frac{p^2(1-p)^2}{2\sigma^2}V''(p)\right] \\ &= V(p, q). \end{aligned}$$

Finally, that prices take the claimed form is now a simple corollary of consumer indifference, noting that at all p, p' in the support of g, we now must have $p - q = p' - q$.

A.7 Theorem 2

The details of the proof are similar to that of Theorem 1. The outline is as follows:

1. Fix the entry rate η and equilibrium price level w.

    (a) Solve explicitly for $(V, \tilde{p})$.

(b) Solve explicitly for f for arbitrary $\tilde{p}$.

2. Prove there exists a unique $w > 0$ such that $V(p_0) = K$.

3. Prove there exists a unique $\eta > 0$ such that $\tilde{p}$ solves (19).

Standard arguments yield that the HJB equation (20) has a unique solution and associated policy. Given that the problem is a linear control, the candidate policy takes the threshold form:

$$s(p) = \begin{cases} \bar{\pi} & \text{if } p \geqslant \tilde{p} \\ 0 & \text{if } p < \tilde{p}, \end{cases} \tag{A.12}$$

where $\tilde{p}$ satisfies the equation:

$$p - w + \bar{\pi}\frac{p^2(1-p)^2}{2\sigma^2}V''(p) = 0. \tag{A.13}$$

Arguments taken from Kuvalekar and Lipnowski (2020) imply that V has a two-sided derivative at $\tilde{p}$ so that smooth pasting occurs. Under this candidate, the value function takes the form:

$$V(p) = \frac{1}{\rho+\delta}\left[\bar{\pi}(p-w)\mathbb{I}_{p\geq\tilde{p}} + (\bar{\pi}\mathbb{I}_{p\geq\tilde{p}} + \epsilon)\frac{p^2(1-p)^2}{2\sigma^2}V''(p)\right] \tag{A.14}$$

The general solution to equation (A.14) is as in the baseline analysis, while the particular solution is now affine. Thus, the full solution is:

$$V(p) = \begin{cases} c_1^0 p^{1-\gamma_1^v}(1-p)^{\gamma_1^v} + \frac{\bar{\pi}(p-w)}{\rho+\delta} & \text{if } p \geqslant \tilde{p} \\ c_0^1 p^{\gamma_0^v}(1-p)^{1-\gamma_0^v} & \text{if } p < \tilde{p} \end{cases} \tag{A.15}$$

where the exponents γ_0^v, γ_1^v and coefficients c_0^0, c_0^1, c_1^0 solve the matrix as detailed in Proposition A.1. Simple algebra confirms the verification step. Furthermore, it is readily verified that $V''(p) \geqslant 0$ for all $p \in [0,1]$, and hence that $\tilde{p} \leqslant w$ (clearly $w - w + \bar{\pi}\frac{p^2(1-p)^2}{2\sigma^2}V''(w) \geqslant 0$). The solutions for the distribution f are identical to the derivations provided in Section A.3.1. To prove the claim that there exists a unique $w > 0$ such that $V(p_0) = K$, note that by expression (A.15), $V(p)$ is continuous and strictly increasing taken as a function of w. As $w \to 0$, $\tilde{p} \to 0$ by the squeeze theorem, and hence $V(p_0) \to \frac{p_0\bar{\pi}-c}{\rho+\delta} > K$, and as $w \to 1$, $V(p_0) \leqslant V(1) \to 0$, thus the claim follows from the Intermediate Value Theorem. Finally, that there exists a unique $\eta > 0$ such that $\tilde{p}$ solves (19) follows from the linear homogeneity of f in η and the argument detailed in Lemma A.4.

References

Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar, "Learning From Reviews: The Selection Effect and the Speed of Learning," *Econometrica*, 2022, *90* (6), 2857–2899.

Atkeson, A., C. Hellwig, and G. Ordoñez, "Optimal Regulation in the Presence of Reputation Concerns," *Quarterly Journal of Economics*, 2015, *130* (1), 415–464.

Banerjee, A. V., "A simple model of herd behavior," *Quarterly Journal of Economics*, 1992, *107* (3), 797–817.

Bénabou, R. and N. Vellodi, "(Pro)-Social Learning and Strategic Disclosure," *American Economic Journal: Microeconomics*, forthcoming.

Bergemann, D. and J. Välimäki, "Market Diffusion with Two-Sided Learning," *RAND Journal of Economics*, 1997, *28* (4), 773–795.

Bikhchandani, S., D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, 1992, *100* (5), 992–1026.

Bisceglia, M. and J. Tirole, "Fair Gatekeeping in Digital Ecosystems," *working paper*, 2025.

Board, S. and M. Meyer-ter-Vehn, "A Reputational Model of Firm Dynamics," *American Economic Journal: Microeconomics*, 2022, *14* (2).

Bolton, P. and C. Harris, "Strategic Experimentation," *Econometrica*, 1999, *67* (2), 349–374.

Campbell, A., C.M. Leister, and Y. Zenou, "Word-of-Mouth Communication and Search," *RAND Journal of Economics*, 2020, *51* (3), 676–712.

Che, Y-K. and J. Hörner, "Recommender Systems as Mechanisms for Social Learning," *Quarterly Journal of Economics*, 2018, *133*, 871–925.

_ **and O. Tercieux**, "Optimal Queue Design," *working paper*, 2023.

Eeckhout, J. and P. Kircher, "Sorting and Decentralized Price Competition," *Econometrica*, 2010, *78* (2), 539–574.

Escudé, M. and L. Sinander, "Slow persuasion," *Theoretical Economics*, 2023, *18* (1), 129–162.

Gabaix, X., J-M. Lasry, P-L. Lions, and B. Moll, "The Dynamics of Inequality," *Econometrica*, 2016, *84* (6), 2071–2111.

Gardiner, C., *Stochastic methods: A handbook for the natural sciences* 2009.

Gindin, A. and E. Shimko, "Dynamic Political Investigations: Obstruction and the Optimal Timing of Accusations," *working paper*, 2023.

Guerrieri, V. and R. Shimer, "Dynamic Adverse Selection: A Theory of Illiquidity, Fire Sales, and Flight to Quality," *American Economic Review*, 2013, *104* (7), 1875–1908.

Hopenhayn, H., "Entry, Exit, and firm Dynamics in Long Run Equilibrium," *Econometrica*, 1992, *60* (5), 1127–1150.

Hörner, J., "Reputation and Competition," *American Economic Review*, 2002, *92* (3), 644–663.

— **and N. Lambert**, "Motivational Ratings," *Review of Economic Studies*, 2021, *88* (4), 1892–1935.

Jovanovic, B., "Selection and the Evolution of Industry," *Econometrica*, 1982, *50* (3), 649–670.

Kolotilin, A., T. Mylovanov, and A. Zapechelnyuk, "Censorship as Optimal Persuasion," *Theoretical Economics*, 2022, *17* (2), 561–585.

Kremer, I., Y. Mansour, and M. Perry, "Implementing the Wisdom of the Crowd," *Journal of Political Economy*, 2014, *122* (5), 988–1012.

Kuvalekar, A. and E. Lipnowski, "Job Insecurity," *American Economic Journal: Microeconomics*, 2020, *12* (2), 188–229.

Lester, B., "Information and Prices with Capacity Constraints," *American Economic Review*, 2011, *101* (4), 1591–1600.

Luttmer, E.G., "Selection, Growth, and the Size Distribution of Firms," *The Quarterly Journal of Economics*, 2007, *122* (3), 1103–1144.

Moscarini, G., "Job Matching and the Wage Distribution," *Econometrica*, 2005, *73* (2), 481–516.

— **and L. Smith**, "The Optimal Level of Experimentation," *Econometrica*, 2001, *69* (6), 1629–1644.

Pallais, A., "Inefficient Hiring in Entry-Level Labor Markets," *American Economic Review*, November 2014, *104* (11), 3565–99.

Rüschendorf, L. and M.A. Urusov, "On A Class of Optimal Stopping Problems for Diffusions with Discontinuous Coefficients," *The Annals of Applied Probability*, 2008, *18* (3), 847–878.

Tadelis, S., "Reputation and Feedback Systems in Online Platform Markets," *Annual Review of Economics*, 2016, *8*, 321–340.