

RATINGS DESIGN AND BARRIERS TO ENTRY

Nikhil Vellodi*

[*Current Draft*]: April 11, 2019 (click [here](#) for the latest version)

Abstract

I study the impact of consumer reviews on the incentives for firms to participate in the market. Firms produce goods of heterogeneous, unknown quality that is gradually revealed via consumer reviews, and face both entry and exit decisions. A platform combines past reviews to construct firm-specific *ratings* that help guide consumer search. When the platform integrates all reviews into ratings — *full transparency* — consumers form queues at the highest-rated firms. This demand cliff induces an S-shaped continuation value for firms as a function of ratings, generating both low entry rates as well as unwanted selection effects — high-quality firms exit early. Whereas firms prefer more feedback when starting out and less feedback when established, equilibrium induces precisely the reverse profile. I then study the design of ratings systems. The platform must balance the need to provide consumers with accurate information against the need to encourage high-quality firms to enter and remain active. The key insight is that optimal rating systems involve *upper censorship*, i.e. the suppression of reviews from highly-rated firms’ ratings, as a means of incentive provision. This makes the task of “climbing the ratings hill” less daunting, stimulating participation. An exploratory calibration using data provided by Yelp! estimates a consumer welfare gain of roughly 7% from adopting the optimal policy.

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

*Department of Economics, New York University. I thank Dilip Abreu, Jarda Borovička, Keith O’Hara, Elliot Lipnowski, David Pearce, Ennio Stacchetti, Joshua Weiss and Leeat Yariv as well as various seminar participants. In particular, I thank Sylvain Chassang, Erik Madsen, Guido Menzio and Alessandro Lizzeri for their invaluable support. Finally, I thank Debraj Ray for his tremendous encouragement and uncompromising guidance. Financial support from a NET Institute Grant is gratefully acknowledged.

1 INTRODUCTION

Product ratings systems - platforms 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 platforms provide an indispensable source of information, reducing search frictions and informational asymmetries and thereby allowing consumers and producers to engage in 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 whilst this stockpile of information might serve incumbent firms to great effect, it might intimidate a new entrant who unavoidably starts from scratch.

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. Consider a recently-opened restaurant. Initially, their meal quality is unknown, but to make matters worse, a nearby restaurant has a reasonable rating on Yelp! with hundreds of reviews. Given the two options, consumers will more likely choose the latter, simply through judicious application of Bayes' Rule - the newcomer's quality is highly uncertain in the absence of informative signals.² In light of this exacting consumer behaviour, the new entrant could shut down prematurely, perhaps after a few poor reviews. Indeed, they might not enter the market in the first place, given the severity of the initial conditions.³

¹Examples include Amazon, Yelp!, Google Reviews, Zomato, TripAdvisor and RateMDs to list just a few. According to recent surveys, over 90% of consumers now consult online reviews before making purchase decisions. Displaying reviews can increase purchasing rates by 270% - see http://spiegel.medill.northwestern.edu/_pdf/Spiegel_Online%20Review_eBook_Jun2017_FINAL.pdf. Various empirical studies document the importance of consumer reviews in determining firm revenue (e.g. Andersen and Magruder (2012), Lewis and Zervas (2016), Luca (2016)). Fradkin (2018) and Farronato and Fradkin (2018) provide evidence that such platforms reduce trading frictions and significantly improve consumer welfare. Tadelis and Zettelmeyer (2015) propose a theory and supporting evidence on how the information provided by such platforms helps trading counter-parties to match efficiently, increasing overall surplus. Finally, see Tadelis (2016) for an excellent survey of both supporting theories and evidence.

²The problem might be exacerbated in reality by heuristic updating procedures that give undue weight to a greater mass of reviews regardless of perceived quality (Powell et al. (2017)).

³The following quote from the website of Zomato, a restaurant review platform, echoes this concern: "*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 the restaurant.*". See <https://www.zomato.com/blog/helping-new-restaurateurs-find-their-feet>. A recent article describes a single bad review on TripAdvi-

One may not see this as a problem. After all, the incumbent had a reasonable rating, and so was likely of good quality. The problem is that the demoralized entrant might have had a superior quality product - they were simply not given the chance to prove their worth. That is, consumer choices give rise to undesirable *selection effects*, whereby high-quality firms opt out of the market. These characteristics could fit a variety of markets, such as medical services, hotels and of course traditional product markets, as well as labor markets.

Such a story relates to the well-documented “cold start” phenomenon, whereby new entrants to a marketplace struggle due to a lack of sampling. Existing research has viewed this as a purely informational problem - a self-reinforcing link between a lack of sampling and a lack of information regarding product quality (Che and Hörner (forthcoming), Kremer et al. (2014)). As such, these studies treat the range of available products as exogenously given, abstracting from producer participation constraints. My analysis captures an important equilibrium feedback channel through which cold-starting endogenously determines the distribution of product quality, which in turn determines the relative demand for new products.

The model comprises three groups of interacting agents: consumers, firms and a *ratings platform*. Firms produce output of heterogeneous quality. Output is stochastic, and depends on the firm’s underlying type, which is either high or low. Each firm’s type is unknown to all market participants, and is gradually revealed through a *rating* publicly provided by the platform to all market participants. Firms pay a fixed cost to enter, subsequently incurring a constant flow cost of operating, and are subject to a service capacity constraint. Once active, firms decide if and when to irreversibly exit the market.⁴ A fixed measure of myopic consumers choose between all available firms by using the rating provided by the platform regarding firms’ quality to engage in frictionless, directed search, subject to firms’ capacity constraints and random rationing. Each consumer truthfully reports their experience to the platform, who uses a firm’s history of such *reviews* to form its rating.

sor as “... *the marketing PR equivalent of a drive-by shooting*”, with the modern consumer labelled as “... *a veritable tyrant, with the power to make or break lives.*”

⁴From an individual firm’s perspective, the problem is a standard optimal stopping problem (see Bar-Isaac (2003)). I extend the analysis to allow firms to make effort and pricing decisions later.

I begin by fully characterizing the unique stationary equilibrium of the economy under a regime of *full transparency*, in which all consumer reviews received by the platform are incorporated into a firm’s rating (Theorem 1). Equilibrium under this regime features congestion by consumers (Lemma 3) who trade off expected quality against the probability of service and thus visit the highest-rated firms and queue for their services. The exacting nature of consumer choice gives rise to a powerful non-convexity in the firm’s problem: its continuation value from remaining operational is S-shaped in its rating (Lemma 4). Importantly, from a firm’s perspective, consumer feedback when poorly-rated comes with *upside gain*; a standard option value effect emerges whereby the gains from a positive review are potentially large, but the losses are ameliorated by the possibility of exit. Conversely, feedback once highly-rated 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. This informational misallocation depresses firms’ incentives, resulting in reduced entry and excessive exit rates.

From this benchmark, a number of testable predictions are derived. The model shares many predictions with well-known models of firm dynamics. For instance, it predicts that average quality increases with age, that exit hazard rates are hump-shaped in age, and that the full invariant distribution of firms over ratings is right-skewed due to selection effects. (see Jovanovic (1982), Luttmer (2007)) However, it produces several predictions that are novel: due to rapid feedback for highly-rated firms, the ratings distribution has a fatter right tail than left (Corollary 1). The sell-out rate and rate of feedback are both increasing functions of a firm’s rating. In particular, the fraction of consumers unable to secure a purchase is an increasing, strictly convex function of the firm’s rating (Lemma 3).

I turn next to design. In equilibrium, it is immediate that difficulties faced by firms will ultimately feed into consumer welfare. Thus, while I model the platform’s objective as being consumer welfare, its choice of policy should evidently account for firms’ incentives. I assume that

the platform is limited in its choice of instrument: it must respect consumers' desires to sample whichever firms they please, but is able to control the inclusion of consumer reviews into firms' ratings. Thus, rather than simply report the entire history of reviews for each firm, the platform can commit to a filtering of this history, in the spirit of Hörner and Lambert (2018). Ratings design in this context then has the ability to shape the industry profile by providing firms with incentives to participate.

The central result from this normative analysis is that the optimal rating system involves *upper censorship*, i.e. excluding reviews from the ratings process of highly-rated firms. In the benchmark model, this feature is stark: the optimal system involves maximal feedback for low-rated firms, and minimal feedback for high-rated firms (Theorem 2). The intuition is simple yet compelling. In order to maximize their welfare, the platform wants consumers to buy from only the highest rated firms. This harms firms, as they enjoy profits only if they achieve this high standard. Without further intervention, entry rates would be depressed and exit would be rapid, ultimately adversely affecting consumers themselves. Since the platform cannot offer direct transfers, nor can they send consumers to firms they are unwilling to visit, they must shore up incentives for firms through careful design of the ratings system. Throwing away reviews for well established, highly rated firms not only prevents their rating from soaring further, but also prevents their rating from sliding. For struggling firms, this makes climbing the ratings ladder both easier and more rewarding once conquered, thus providing the necessary encouragement. To test the robustness of this result, I explore several extensions, the most important of which is allowing firms to post prices. Prices provide an important intensive margin through which firms can not only calibrate their current terms of trade, but also their future prospects. For instance, new firms might price at a loss in order to attract consumers and feedback, while established firms might increase prices in order to soak up the extra demand generated by higher ratings. I show that both effects emerge in equilibrium. If price caps or floors were to exist, the qualitative features of my analysis would carry through. However, in the limiting case where prices are fully flexible and with respect to the policy

space I consider in this paper, full transparency emerges as optimal.⁵

The paper concludes with a simple calibration of the theoretical model. I use data publicly provided by Yelp!⁶ Specifically, I proceed under the assumption that Yelp! is using a policy best approximated by the full transparency regime, and calibrate the parameters of the model using first and foremost the ratings distribution. To discipline parameters further, I utilize data provided by the Bureau of Labor Statistics (BLS) on firm hazard rates to match the empirical expected lifetime of firms to the model prediction (Luo and Stark (2014)). I use the calibrated model to perform various counterfactual comparisons (Table 1). I find that, were the platform to adopt the optimal rating design proposed in section 3, consumer welfare would increase by roughly 7%. Of course, given both the simplicity of the theoretical model and the exploratory nature of the exercise performed here, these results should be interpreted more as back-of-the-envelope calculations.

1.1 CONTRIBUTION AND RELATED LITERATURE

The main contribution of the paper is to two strands of research - the design of **recommendation/ratings systems** and **platform design**. It is the first analysis to study the role of information design in shaping industry dynamics through endogenous participation. Hörner and Lambert (2018) studies the design of ratings in order to incentivize a single firm to exert hidden effort and improve output quality. In their analysis, the arrival rate of information is independent of the firm’s current rating, and thus abstracts from the cold start constraint central to the current paper. Furthermore, their model comprises a single firm, and thus the distributional concerns central to my analysis are absent. Hörner (2002) does study the interaction of competitive forces with firms effort choices and exit decisions, showing how competition can mitigate the inefficiencies that plague settings with career concerns and moral hazard. There however, consumers do not learn socially, and the paper also abstracts from the implementation problem I study here. Che and Hörner (forthcoming) and Kremer et al. (2014) examine the intertemporal informational

⁵In ongoing work where I consider a broader class of ratings policies, suppression still has a role to play in generating incentives: pooling high-rated firms allows the lowest-rated amongst these to charge higher prices, effectively front-loading incentives.

⁶See: <https://www.yelp.com/dataset>.

externality that consumer choices generate, and thus also identify policies that can help alleviate the cold-start problem. Crucially, they treat the range of products as exogenous, and thus abstract from firms' incentives. Also, by focusing on the single-product case, their consumers necessarily have an exogenous incentive-compatibility constraint, whereas in my setting this constraint is endogenous and determined itself by firm's entry/exit choices. Finally, Goel and Thakor (2015) argues that coarse credit ratings might balance the need for transparency in financial markets against the need to alleviate moral hazard. The idea of diverting consumers from their optimal product choices in order to redistribute market power amongst competing producers is present in Hagiu and Jullien (2011), Yang (2018), Kovbasyuk and Spagnolo (2018) and Romanyuk and Smolin (forthcoming). The latter argues how congestion can occur naturally as a result of excess information, and thus argue how restricting information via upper censorship can alleviate this inefficiency, whereas in Kovbasyuk and Spagnolo (2018), discounting past observations from firm's ratings can be welfare-enhancing in the presence of evolving hidden quality. Both papers abstract from the entry/exit choices of firms. Indeed, that upper censorship plays a role in my analysis even in the absence of congestion (see section 4.1) highlights the importance of participation constraints in determining the value of information. The optimality of *upper censorship* disclosure policies is present in a number of recent papers in the persuasion literature (Romanyuk and Smolin (forthcoming), Yang (2018), Bloedel and Segal (2018), Kolotilin et al. (2017)). In Bloedel and Segal (2018), the result obtains due to costly information processing and thus the need to trade-off accuracy against fidelity. Finally, the provision of a coarse information policy by a monopolist echoes results from the **certification** literature (Lizzeri (1999), Biglaiser (1993)). More recently, Marinovic et al. (2018) extend these results to a dynamic setting with moral hazard, demonstrating that restricting the ability to voluntarily certify can improve incentives to invest in quality.

The analysis shares features of the literature on **collective experimentation** and **social learning**, wherein the central inefficiencies center on dynamic free-riding effects. 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) also cast 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. (2018) also study in recent work the endogenous speed of learning in an observational learning framework. They argue that consumer selection drives both the accuracy and speed of feedback, and abstract from firms' incentives entirely. A related recent paper is Campbell et al. (2018), which to the best of my knowledge, is the only other paper to endogenize the underlying quality distribution in an otherwise standard social learning setting. Their model features social learning between consumers connected on a network, rather than via a platform, and firm entry, and abstracts from information design question my normative theory addresses.

A classic literature in industrial organization describes how informational asymmetries can pose a **barrier to entry** for late arriving firms (Schmalensee (1982), Bagwell (1990), Grossman and Horn (1988)). These papers highlight how superior information regarding product quality can endow incumbent firms with a first-mover advantage, leading to inefficient entry choices, even with price setting. Schmalensee (1982) outlines a two-period model, in which the incumbent and entrant are of identical quality. Bagwell (1990) extends this analysis by allowing the incumbent to be of lower quality than the entrant, showing how inefficiencies still prevail. Grossman and Horn (1988) focusses on a moral hazard margin in the choice of quality in a trade setting. However, none of these papers incorporate the social learning feature of information diffusion that my analysis hinges on, nor do they consider ratings design as a method for stimulating entry.

My model also contributes to the literature on **firm dynamics**, in which firms make entry and exit decisions that are governed by an evolving state process. Classical models such as Hopenhayn (1992) and Luttmer (2007), Jovanovic (1982) and Board and Meyer-ter-Vehn (2014) also center around predictions regarding cross-sectional firm distributions and dynamic hazard-rates. Technically, Luttmer (2007) also builds upon the theory of *resetting processes*, whereby firms exit at some state z and arrive at some other state $z' > z$. Closely related is the work by (Atkeson et al., 2015), who study a competitive market with adverse selection and firm dynamics. Again, the asymmetry in learning rates is not present, nor is the information design problem – their policy

instrument is the entry cost. Furthermore, the market structure is quite different: there, a firm’s demand function is linear in its current rating and increasing in the overall mass of active firms. Bar-Isaac (2003) involves a firm that solves a similar stopping problem, the main differences being the endogenously generated flow profit function as well as variable learning rates. In a recent paper, Kuvalekar and Lipnowski (2018) study the firing decision of a firm over a worker with no monetary transfers, and show how the worker’s equilibrium choices involves signal jamming as a means of avoiding rapid exit. Also related is Hörner (2002), who shows how market forces can alleviate standard moral hazard incentive concerns by embedding a standard career-concerns problem into a competitive market with entry and exit decisions. There, the assumptions that consumers are forward-looking and learn privately as well as that one firm can serve the entire market gives rise to severe switching effects not present in my analysis.

The baseline model itself contributes to the theory of **mean-field games**, models which comprise a backward equation (here, the ODE governing firms’ continuation values) and a forward equation (here, the Fokker-Planck equation governing the ratings distribution). Gabaix et al. (2016) and Moscarini (2005) both involve the derivation of boundary conditions similar to those outlined in Proposition 1. Mine is the first study to explicitly embed an information design problem into such an environment. This confluence poses new technical challenges, which I solve using elements of both stopping and control theory, as well as the theory of so-called *oscillating diffusion processes* - see le Gall (1985) and Keilson and Wellner (1978).

Finally, the exploratory calibration offered in Section 5 follows the **empirical** literature on consumer reviews. Previous studies that attempt to back out quality from reviews include Dai et al. (forthcoming), (Andersen and Magruder, 2012), Li and Hitt (2008), Chevalier and Mayzlin (2006). For instance, Dai et al. (forthcoming) assumes that the underlying process governing firm quality is mean zero, capturing as a reduced-form the martingale property of belief updating under Bayesian inference. Insofar as it seems plausible that consumers choose firms based on expected quality, rather than simply the average rating, this is an important transformation in and of itself.

Horton (2018) and Fradkin (2018) document the presence of severe congestion externalities

due to capacity constraints in online platforms. Relatedly, Lewis and Zervas (2016) document how increased demand due to reputational effects does not fully pass through into higher prices, implying the co-existence of congestion and price discrimination (see Section 4.1 for a discussion in the present context). In recent work, Luca and Luca (2018) document a negative correlation between a firm’s average rating and the probability of exit. Hui et al. (2018) demonstrate how, following a tightening of the certification standards at an online platform, entry initially increases but then levels out in the long-run, while the quality distribution of new entrants exhibits a higher mean and thicker right tail. However, theirs is a setting with severe adverse selection concerns. Of particular note is a recent working paper by Li et al. (2018), who document an interesting policy intervention that temporarily allowed sellers at an online trading platform to pay for reviews via rebates. They demonstrate how new entrants used the service more than established players, consistent with the results presented here.

2 BENCHMARK MODEL - FULL TRANSPARENCY

Time is continuous and doubly infinite. The economy consists of positive measures of firms and consumers, as well as a platform. As I will be working with stationary equilibria, time subscripts are dropped, with t henceforth denoting the age of a firm. I begin by abstracting from the design question by assuming that the platform simply includes all consumer feedback into each firm’s rating.

Firms - A large, infinitely elastic supply of firms can potentially participate in market production. Each firm can be one of two types, $\theta \in \{0, 1\}$. Types are fixed throughout the life of a firm and are initially hidden, with a new entrant being of type θ_h with probability p_0 . While active, a firm of type θ and age t is associated with stochastic process $(X_t)_{t \geq 0}$ that evolves according to

$$dX_t(p_t) = \lambda(p_t)\theta dt + \sqrt{\lambda(p_t)}\sigma dZ_t \tag{1}$$

where $(Z_t)_{t \geq 0}$ is a Wiener process independent of θ , $\sigma \in (0, \infty)$, the function λ satisfies⁷:

Assumption 1. $\lambda(p) \neq 0, \forall p \in [0, 1], \quad \frac{1}{\lambda^2(p)} \in L_{loc}^1([0, 1]),$

and p_t - henceforth referred to as a firm's *rating* - is the instantaneous probability that the firm is of high type given the information contained in X_t . Formally, let $(\Omega, \Sigma, \mathbb{P})$ be a probability space rich enough to admit Z , and let \mathbb{E} denote the unconditional expectation operator under \mathbb{P} . Let $\mathbb{F}^x = \{\mathcal{F}_t^x\}_{t \geq 0}$ denote the natural filtration generated by $(X_t)_{t \geq 0}$. Finally, \mathbb{E}_t^x denote the conditional expectation under \mathbb{P} with respect to \mathcal{F}_t^S , so that $p_t = \mathbb{E}_t^x(\theta)$.

In this section, I take $\lambda(p) = \pi(p) + \epsilon$, for some $\epsilon > 0$. This specification admits a simple, established micro-foundation.⁸ $(X_t)_{t \geq 0}$ can be interpreted as *cumulative review process*, where $\pi(p)$ denotes the rate at which consumers that use the platform are served by the firm, and ϵ represents background learning generated by un-modelled consumers that do not use the platform to guide their search, visiting firms at random while leaving feedback nonetheless (see Che and Hörner (forthcoming) for a similar approach). As such, $\pi(p)$ is endogenously determined by consumer choice, but taken as given by a firm, while in this case X is called *cumulative output*.

Firms simply choose whether to enter and subsequently exit the market. Firms pay an entry cost of $K > 0$. Once entered, firms pay a flow cost $c > 0$ to remain active. A firm with current rating p makes flow revenues equal to $\pi(p)$.⁹ Firms are subject to a service capacity constraint $\bar{\lambda}$. Finally, firms discount at rate ρ and face a constant hazard-rate δ of exogenous attrition. I make the following minimal assumptions on firms' costs:

Assumption 2. $\bar{\lambda} > c$

Assumption 3. $K < \frac{\bar{\lambda} - c}{\rho + \delta}$

Were either of these assumptions violated, no firm would ever enter the market.¹⁰ The exit

⁷See Engelbert and Schmidt (1991)

⁸See Bergemann and Välimäki (1997) for details.

⁹While it is in principal possible to allow λ to depend directly on the age t , this will turn out not to be necessary under full transparency. I further assume that background consumers generate zero revenue. This is without loss - one could simply incorporate this gain into c .

¹⁰The expression $\frac{\bar{\lambda} - c}{\rho + \delta}$ is the present-discounted value of selling out forever, and thus forms an upper bound on incumbent firms' continuation values.

decision of firms takes the form of a standard optimal stopping problem, in which its rating p_t forms a natural state variable. Applying Theorem 9.1 in Lipster and Shryaev (1977), noting that the capacity constraint ensures that λ is a locally bounded function and so Itô's Lemma can be applied.

Lemma 1. *Ratings evolve according to the SDE:*

$$dp_t = \frac{\sqrt{\lambda(p_t)}}{\sigma} p_t(1 - p_t) d\bar{Z}_t \quad (2)$$

where $d\bar{Z}_t = \frac{1}{\sigma\sqrt{\lambda(p_t)}} [dX_t - \lambda(p_t)p_t dt]$ is a standard \mathbb{F}^x -adapted Wiener process.

Thus, more generally, $\lambda(p)$ can be thought of as measuring the rate at which reviews pass into the stock X and so controls the speed at which a firm's rating evolves. For instance, if $\tilde{\lambda}(p) = 0$, no new reviews are added to the firm's profile, and consequently the firm's rating remains frozen. The present value to a firm with rating p is

$$v(p) = \sup_{\tau^x} \mathbb{E}^x \left[\int_0^{\tau^x} e^{-(\rho+\delta)t} (\pi(p_t) - c) dt | p_0 = p \right] \quad (3)$$

where the supremum is taken over all \mathbf{F}^x -measurable stopping times.

Standard verification theorems exist for this setting (Rüschendorf and Urusov (2008)), yielding that the firm's exit strategy takes the form of a *rating threshold* $\underline{p} \in (0, 1]$, combined with a second-order ODE that expresses the firm's continuation value $V(p)$ from remaining operational, given a current rating of p .

Lemma 2 (Incumbent's Problem). *Suppose there exists an $\omega < p_0$ such that $\pi(p) < c$ for all*

$p \in [0, \omega]$. Let the pair $\{u(\cdot), \underline{p}\}$ for $u \in \mathcal{C}^1([0, 1])$ and $\underline{p} \in [0, 1]$ denote the variational inequality:

$$\begin{aligned}
\mathcal{A}\{u\} &= 0 \\
u(\underline{p}) &= 0 \\
u'(\underline{p}) &= 0 \\
u(p) &\geq 0 \quad \forall p \in [\underline{p}, 1] \\
u(p) &= 0 \quad \forall p \in [0, \underline{p}],
\end{aligned} \tag{4}$$

where

$$\mathcal{A}\{u\} = \pi(p) - c + \Sigma(p)u''(p) - (\rho + \delta)u(p), \quad \Sigma(p) = \frac{1}{2\sigma^2}p^2(1-p)^2\lambda(p) \tag{5}$$

Suppose $\{V(\cdot), \underline{p}\}$ solves the problem (4), (5), and that with respect to the process $(p_t)_{t \geq 0}$ and that $\tau = \inf_s \{s \geq 0 : p_s \leq \underline{p}\}$. Then the function $V(\cdot)$ coincides with the value $v(\cdot)$ in (3), τ attains the supremum in (3), and $\{V(\cdot), \underline{p}\}$ are unique.

Finally, an infinitely elastic supply of firms leads to a free entry condition on the present value to an entrant:

$$V(p_0) = K, \tag{6}$$

that in turn pins down the equilibrium rate of entry η .

The Ergodic Ratings Distribution - The combination of evolving ratings and a continuously churning positive mass of firms gives rise to an equilibrium firm distribution f . Since ratings are described by a diffusion process, they are Markovian and strongly recurrent. As such, the invariant distribution F^∞ for the ratings process is also ergodic, admitting a density almost everywhere with support $[\underline{p}, 1]$. Denote this density by f^∞ . Since the mass of firms endogenously determined, and thus generically not equal to unity, f is a re-scaling of f^∞ determined by aggregate flow conditions detailed below. We abuse terminology and refer to f itself as the *ratings distribution*. The law of motion for this distribution follows the Fokker-Planck forward equation, with stationarity imposed and subject to various boundary conditions.

Proposition 1. *Let $f(p)$ denote the ratings distribution. Then $f(p) = 0$ for all $p \in [0, \underline{p}]$. For almost all $p \in [\underline{p}, 1]$, f satisfies the Fokker-Planck forward equation:*

$$\frac{\partial^2}{\partial p^2} \Sigma(p) f(p) - \delta f(p) = 0 \quad (7)$$

subject to the following conditions:

- | | |
|--|---|
| 1. $f(\underline{p}) = 0$ | 4. $\frac{\partial}{\partial p} \Sigma(p) f(p) \in \mathcal{C}^1([\underline{p}, p_0] \cup (p_0, 1])$ |
| 2. $\Sigma(p_0)[f'_-(p_0) - f'_+(p_0)] = \eta$ | 5. $\Sigma(p) f(p) \in \mathcal{C}^1([\underline{p}, 1])$ |
| 3. $\Sigma(1) f(1) = 0$ | 6. $\delta \int_{\underline{p}}^1 f(p) dp + \Sigma(\underline{p}) f'_+(\underline{p}) = \eta$ |

While technical, many of these conditions provide economic insight into the dynamical system when properly interpreted. Condition 1 is the classical “attainable boundary” condition, stating that firms spend no time at the exit threshold \underline{p} . Condition 2 states that the rate at which incumbent firms move away from the initial rating p_0 must equal the rate of inflow by new entrants. Condition 6 is an aggregate balance equation, analogues of which can be found in models of labor search and matching, e.g. Moscarini (2005). It states that outflows due to either attrition or voluntary exit must equal inflows.

Consumers - A unit measure of consumers use the platform 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 the available choices $f(p)$, consumers perform frictionless directed search subject to random rationing (see Guerrieri and Shimer (2013), Lester (2011)). That is, if a consumer chooses to direct their search to firms with rating p , they are served at a rate:

$$\Theta(p) = \min \left\{ \frac{\bar{\lambda} f(p)}{g(p)}, 1 \right\} \quad (8)$$

where $g(p)$ is the density of consumers also searching within this submarket. Combining this matching technology with the assumption of frictionless search and risk-neutrality yields the

well-known indifference condition on the equilibrium value J achieved by consumers:

$$J = \max_{p \in [p, 1]} [p\Theta(p)], \quad (9)$$

subject to the market clearing condition:

$$\int_p^1 g(p) dp = 1 \quad (10)$$

Discussion - The benchmark model was constructed to balance the usual objectives of allowing the key economic forces to speak clearly while incorporating sufficient richness to render the analysis robust. However, given its novelty, a brief defence of its main features is certainly warranted.

The assumption that prices are fixed and exogenous is certainly restrictive. In certain settings, such as the market for medical services, prices are fixed or highly inflexible from the viewpoint of consumers. It may well be that allowing new entrants to engage in introductory pricing would provide the informational advantage they desire. Arguably the most natural model of price formation in such a setting would be competitive search price posting e.g. Eeckhout and Kircher (2010). I fully solve this extended model; see Section 4.1 for details.

The optimal stopping problem faced by firms is well-studied. See Bar-Isaac (2003) for the leading example, as well as Board and Meyer-ter-Vehn (2014). Section 4.2 considers the case where firms know their type, and hence their actions serve to signal quality.

Modeling consumers as performing frictionless directed search is both tractable and realistic, capturing several intuitive features of platform search; a consumer opens Yelp!, and is greeted by a list of all available restaurants, along with a fully informative rating for each one. They then decide which one to go to, trading off quality against congestion - better restaurants typically involve longer service times. Such patterns echo recent empirical findings - Horton (2018) and Fradkin (2018) document pervasive congestion externalities in the context of an online platform. The idea that consumers trade off congestion against quality seems natural, and indeed has empirical support in the context of restaurants (Andersen and Magruder (2012)). That said, the central

findings of the benchmark analysis would remain in tact if several alternative specifications for consumer behaviour were used. I detail these in the Appendix.

2.1 STATIONARY EQUILIBRIUM

I look for stationary equilibria of the above model:

Definition 1. A stationary equilibrium is a collection of functions $\{V, f, g\}$ defined on $[0, 1]$ and scalars $\{J, \underline{p}, \eta\}$ such that equations (4), (6), (7), (9), (10) and the relevant conditions listed in Lemma 2 and Proposition 1 all hold.

I start by solving for consumers' equilibrium choices $g(p)$. This is a potentially complicated problem, as there exists a non-trivial fixed-point relationship between g and the distribution f - note that equation (7) governing f depends on $\lambda(p)$, which in turn depends on $g(p)$ and thus on f through equation (9). However, the solution turns out to be describe by a simply *threshold strategy*:

Lemma 3. *In equilibrium,*

$$g(p) = \begin{cases} \frac{\bar{\lambda} p f(p)}{p^*} & \text{if } p \geq p^* \\ 0 & \text{if } p < p^* \end{cases} \quad (11)$$

for some $p^* \in (\underline{p}, 1]$.

While firms with a rating that is too low are simply not visited at all, firms with high ratings enjoy queues that ensure they sell out. Moreover, these queues are longer at higher-rated firms in order that consumers maintain indifference in equilibrium. Thus, from an individual firm's perspective, profits $\pi(p) = \min \left\{ \frac{g(p)}{f(p)}, \bar{\lambda} \right\}$ take the form of a step function:

$$\pi(p) = \begin{cases} \bar{\lambda} & \text{if } p \geq p^* \\ 0 & \text{if } p < p^* \end{cases} \quad (12)$$

This simple structure means that, all told, only three scalars - the thresholds \underline{p}, p^* and the entry rate η - are required to fully characterize a stationary equilibrium. In particular, proving the existence and uniqueness of equilibrium boils down to a pair of simple, scalar fixed-point problems. While details are relegated to the Appendix, I outline the basic idea here. Fix the entry rate η . Trace the locus $R^*(\underline{p})$ that yields the optimal p^* given a fixed \underline{p} , as well as the mirror image $R_-(p^*)$. By the uniqueness of both \underline{p} in Proposition 2 and p^* in Proposition 9, these loci are both functions. It remains to be shown that they intersect precisely once.

First, an increase in p^* leaves firms worse off - on average, they make positive profits for a shorter fraction of their life-cycle. Hence, firms exit earlier and thus \underline{p} must rise - $R_-(p^*)$ is upward sloping. The reaction of consumers to firms' exit is both more subtle and more illuminating. The key insight of my proof is to show that, when \underline{p} rises, the ratings distribution $f(p)$ is lower at every rating p . In particular, fewer high quality firms remain active. Here then, we see the negative effect that exit induces on selection. In response, consumers are simply forced to lower their standards, as fewer firms remain in the right tail of the ratings distribution. Thus, $R^*(\underline{p})$ is downward sloping. Finally, it remains to show that the value to an entrant $V(p_0)$ is decreasing in the entry rate, which follows from the fact that higher entry allows consumers to raise their standards (p^* rises) and in turn depress firms' values.


Theorem 1. *There exists a unique stationary equilibrium, featuring:*

- *Positive rates of both entry and exit.*
- *Sell-out profits and rapid feedback at established firms.*
- *Losses and slow feedback at struggling firms.*

2.2 EQUILIBRIUM FEATURES

A Tough Climb and a Fear of Falling - Learning from consumer feedback transforms the step profit function into a smooth, S-shaped value function for firms, as shown in Figure 1. Intuitively,

Figure 1: Firms' Value Function, Ratings Distribution



Left panel: incumbent value function $V(p)$. Right panel: ratings distribution $f(p)$. Parameter values: $\rho = 2, \delta = 0.2, \sigma = 1.5, \epsilon = 0.2, \bar{\lambda} = 0.8, c = 0.15, p_0 = 0.55$

below the consumer threshold p^* , firms derive option value from their rating potentially rising via background learning, whereas above p^* , feedback simply leaves firms fearful that their rating might drop, losing their consumer base. This polarized preference for learning between firms manifests in a value function that is convex for $p < p^*$ and concave for $p > p^*$, a fundamental non-convexity that stems from the severe equilibrium behaviour of consumers.

Lemma 4. *A firm's equilibrium value function $V(p)$ is S-shaped. Specifically, $V''(p) > 0$ for all $p \in [\underline{p}, p^*)$ and $V''(p) < 0$ for all $p \in [p^*, 1]$.*

It stands to reason that rapid exit stems from both the daunting task of climbing the ratings ladder and the fearful prospect of slipping down the ladder once climbed. Put differently, if it were feasible, the following profile of signal precision would maximize incumbent firm's value functions; set $\sigma = 0$ for firms below p^* , and $\sigma = \infty$ above p^* , a fast-slow profile of learning. Alas, in equilibrium, the profile is *entirely the reverse*, i.e. slow-fast. Here then, we begin to see the extent of the information misallocation problem that plagues the full transparency regime, at least from firms' viewpoint, an insight that will prove invaluable when considering the platform's design problem.

Fat Tails and Power-Law Ratings Distribution - Classical models of industry dynamics combine firm entry, exit and a stochastically evolving state variable to generate an endogenous

equilibrium distribution of firms.¹¹ Detailed empirical work has uncovered a robust finding, namely that the right tail of this distribution follows a power law. Such a feature is shared by the ratings distribution $f(p)$ generated in this paper. In fact, the tractability of the model allows me to say more:

Proposition 2.

Let

$$\gamma_0^f = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2\sigma^2\delta}{\epsilon}}, \quad \gamma_1^f = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2\sigma^2\delta}{\epsilon + \bar{\lambda}}}$$

Then for $[p, p^*]$, $f(p)$ is a linear combination of the functions

$$p^{-1-\gamma_0^f}(1-p)^{\gamma_0^f-2} \quad \text{and} \quad p^{2-\gamma_0^f}(1-p)^{-1-\gamma_0^f},$$

for $[p^*, 1]$, $f(p)$ is a linear combination of the functions

$$p^{-1-\gamma_1^f}(1-p)^{\gamma_1^f-2} \quad \text{and} \quad p^{2-\gamma_1^f}(1-p)^{-1-\gamma_1^f}$$

Full details are given in the Appendix.

Firms enter at a rating p_0 , causing a kink at $f(p_0)$, and flow to either side depending on the reviews left by consumers. Crucially, if $p^* > p_0$, then users of the platform *do not visit new firms*, and thus entrants must hope that the background flow of information ϵ is enough to push their rating above p^* so that they might make profits. Once above p^* , a firm's rating moves at a faster rate, since the added flow of consumer reviews creates a more precise feedback process. Figure 1 summarizes these findings graphically.

This asymmetry in the flow speed of information between firms has a striking effect on f . Slow learning for the worst firms induces a steep left tail. Indeed, as background learning becomes negligible, firms below p^* face slim prospects for their rating to climb and thus exit rapidly. This

¹¹Typically, the state is a measure of productivity. See for instance Hopenhayn (1992), Jovanovic (1982), Luttmer (2007). See Atkeson et al. (2015), Board and Meyer-ter-Vehn (2014) for examples where the state is a belief.

outflow thins the left tail of the distribution, transferring weight to the right.

Corollary 1. *The ratings distribution exhibits a discontinuous drop at the consumer's threshold p^* , and has a fatter right tail than left.*

Cross-sectioning by Firm Age - While unnecessary for the determination of equilibrium, it is possible to analytically solve for the distribution of firms by age as well as ratings, $f(p, \tau)$, $\tau \in [0, \infty)$. Several intuitive predictions can be extracted from these expressions. Older firms exhibit a conditional distribution of output with higher mean - an immediate consequence of selection - but lower variance. Exiting firms tend to be middle-aged; new entrants inevitably begin their ratings process bounded away from the exit threshold, while established incumbents are necessarily of high quality simply through selection, and thus have high ratings. Furthermore, cross-correlating firm age with the quantity of reviews received generates a further prediction that is almost immediate in a model with reputation formation and exit decisions: firms with many reviews tend to have more good than bad reviews.

Welfare - To prepare for the subsequent normative analysis, it is instructive to briefly discuss how one might make welfare comparisons across stationary equilibria of the model. The welfare of consumers is simple and unambiguous, as they are short-lived. Accounting for long-lived firms subject to entry and exit across time is less straightforward. However, due to the free entry condition imposed on firms, standard arguments exert that a suitable measure of social welfare at a stationary equilibrium is consumer welfare¹²:

$$SW \equiv \int_p^1 pg(p)dp + \eta [V(p_0) - K] = CW \tag{13}$$

This equivalence is important for the normative theory of Section 3 — a platform attempting to maximize either social or consumer welfare will adopt the same policy.

¹²See Burdett and Menzio (2018) for a discussion.

3 RATINGS DESIGN

In reality, full transparency provides only partial guidance as to the value a platform can provide to its users. Typically, platforms can shape consumer choice through various means - here, I focus on the design of ratings. Before turning to this, as with any theory of second-best, it is instructive to consider what the platform can achieve when subject to neither firms' nor consumers' participation constraints. With consumers performing directed search, it is almost immediate that they benefit from more firms of any type in the market. That is, the platform wants to *maximize entry* and *minimize exit*. Recall the key insight from the proof of Theorem 1 - a lower exit threshold \underline{p} increases $f(p)$ everywhere, i.e. at low- as well as high-rated firms. Since consumers can effectively ignore bad firms, this is unambiguously good for them.¹³ However, a primary role of such platforms is to help avoid such search frictions, and thus directed search seems a sensible benchmark.

Recent studies show that, while many platforms use simple aggregation methods to construct product ratings, gains could be enjoyed from a more carefully designed ratings system (Dai et al. (forthcoming)). Furthermore, some sites, such as Zomato, engage in so-called *ratings normalization*, whereby they transform a restaurant's rating as a function of the entire distribution $f(p)$ with the explicit intention of guiding consumer choice.¹⁴ I interpret ratings design as the mapping from a firm's history of reviews into those reviews included when updating the rating according to Bayes' rule - under full transparency, this mapping is the identity map. My analysis allows direct comparison of the ratings distribution that obtain under optimal and transparent design, and can thus speak to the type of policy Zomato adopts.

Definition 2 (Ratings). Given a cumulative output process X , a **rating** p is a real-valued process that is progressively measurable w.r.t. \mathbb{F}^x . A rating p is **simple** if there exists a process Y and a function $\lambda : [0, 1] \rightarrow \mathbb{R}_{++}$ such that:

1. $\lambda(p) \geq \iota$, for some $\iota > 0$

¹³Such a relationship would not hold if for example consumers performed McCall-style search and thus could not avoid bad firms.

¹⁴See <https://www.zomato.com/blog/simplifying-ratings-for-a-better-dining-experience>.

2.

$$dY_t = \lambda(p)dX_t, \quad \text{where } p = \mathbb{E}_t^y(\theta)$$

In this case, Y is referred to as a **cumulative review process**, whereas the function λ is referred to as the **rating policy**. I denote simple policies by $\{\lambda\}$.

I endow the platform with the ability to commit to a ratings policy. Practically speaking, it would be challenging to update a rating policy in light of firm-specific or even distributional evidence as it evolves.¹⁵ In the Appendix, I demonstrate that the optimal rating system within a broad class is simple. Henceforth, I focus solely on simple ratings. Such ratings arise from the following simple procedure. The platform commits to randomly excluding a fraction of incoming reviews from the cumulative review process Y . This class of ratings is tractable to study and simple to implement.

What is important to highlight is the lower bound ι on the review inflow rate. This is imposed foremost for expository reasons; market equilibria with respect to such policies are defined with precisely the same conditions as under full transparency. Were the policy to involve p for which $\lambda(p) = 0$, the conditions associated with the Fokker-Planck equation as well as firms' value functions would need to be amended for technical reasons.¹⁶

Definition 3 (Implementability). A policy $\{\lambda\}$ is **implementable** if there exists a collection $E = \{V, J, g, f, p, \eta\}$ such that equations (4), (6), (7), (9), (10) and the relevant conditions listed in Lemma 2 and Proposition 1 all hold, and for all $p \in [0, 1]$,

$$0 \leq \lambda(p) \leq \pi(p) + \epsilon \tag{14}$$

In this case, we say that $\{\lambda\}$ **implements** E and that E is a (stationary) equilibrium w.r.t $\{\lambda\}$.

¹⁵Yelp! 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?l=en_US.

¹⁶Specifically, the state p^* would become a *sticky boundary*, and hence the ratings distribution would be described a $f(p)$ over $[p, p^*)$ and a mass point at p^* .

Let $\mathcal{E}(\lambda) = \{E : \{\lambda\} \text{ implements } E\}$. Finally, a simple policy $\{\lambda\}$ is **non-empty** if it implements an equilibrium E with $\eta > 0$.

Thus, a policy is implementable if it supports a stationary equilibrium as previously defined, in which consumer behavior induces service rates that are consistent with the prescribed inflow function; the rate of inflow cannot exceed the maximum feasible flow rate ($\lambda(p) \leq \pi(p) + \epsilon$), and reviews cannot be destroyed ($\lambda(p) \geq 0$). Empty policies are of course always dominated, since full transparency supports non-zero entry rates and thus strictly positive consumer welfare.

The central question in this section is then, what can the platform achieve by designing a simple rating system? To begin with, since consumers are effectively solving the same problem, the allocation $g(p)$ retains its p^* -structure *in any equilibrium for any policy*.

Lemma 5. 1. *Let $\{\lambda\}$ be non-empty with $E \in \mathcal{E}(\lambda)$. Then there exists $p_E^*(\lambda) \in (0, 1)$ such that*

$$\lim_{p \nearrow p_E^*(\lambda)} G(p) = 0 \quad \text{and} \quad \Theta(p_E^*(\lambda)) = 1, \Theta(p) < 1 \quad \forall p > p_E^*(\lambda)$$

$$\text{Let } p^*(\lambda) = \max_{E \in \mathcal{E}(\lambda)} p_E^*(\lambda)$$

2. *Consumer welfare in equilibrium outcome E is given by $p_E^*(\lambda)$.*

To see that $p_E^*(\lambda)$ equals consumer welfare, note that at this rating consumers are guaranteed service, and thus their expected value is simply $p_E^*(\lambda)$. The indifference condition in equation (11) completes the proof. This argument holds irrespective of the ratings distribution, and thus for any policy. As such, the platform's optimization problem is remarkably simple to state:¹⁷

$$\underset{\{\lambda\}}{\text{maximize}} \quad p^*(\lambda)$$

Crucially, an increase in p^* directly depresses firms' profits. Here then, we see the stark tension faced by the platform; the objective is to ensure as great a split of the surplus for consumers as

¹⁷Embedded in the definition is the assumption that, were a policy $\{\lambda\}$ to admit multiple equilibria ($|\mathcal{E}(\lambda)| > 1$), the platform can select the most favorable outcome, i.e. *ex-post implementation*. This assumption is irrelevant, as the optimal policy will be shown to induce a unique equilibrium.

possible, but, by doing so, this shrinks the surplus itself by depressing firms' incentives and thus reducing the quantity of highly-rated participating firms. The platform must then compensate firms through low-powered incentives; *choosing the policy that delivers firms' optimal feedback rates ensures that consumer welfare is maximized.*

This transforms the computationally intractable problem of searching over policies while solving for their respect equilibrium conditions into a single-firm, dynamic information control problem, much in the spirit of Moscarini and Smith (2001). The following result constitutes the main result of this section.

Theorem 2. *The optimal policy $\{\lambda\}$ involves:*

- $\lambda(p) = \epsilon$ for all $p < p^*(\lambda)$.
- $\lambda(p) = \iota$ for all $p \geq p^*(\lambda)$.


The structure of the solution is immediate, as discussed in Section 2.2; a firm's value function is maximized pointwise by having maximal feedback below p^* and minimal feedback above p^* . Recall the motivating example; the entrant would surely benefit from the platform throwing away half the incumbents reviews, levelling the playing field. What is less obvious is that this extreme feedback profile also constitutes the platform's optimal policy. Again, this reveals the powerful role general equilibrium interactions play in shaping incentives: entry/exit concerns are of first-order concern when it comes to providing consumers with high-quality options.

4 EXTENSIONS

4.1 INTENSIVE MARGINS - PRICES

A key simplification in the model is the absence of a control variable for firms that allows them to calibrate their terms of trade more scrupulously. For instance, once a firm's rating improves, it might be able to set higher prices or exert lower effort in an attempt to extract surplus from

Figure 2: **Timing over interval** $[\tau, \tau + d\tau]$



increased demand. Similarly, firms at lower ratings or indeed new entrants might set low prices in order to attract customers and generate reviews.

In this section, I explore these possibilities by extending to model to allow firms to flexible set prices, employing the approach of the competitive search literature. Specifically, the timing of events of a infinitesimal time interval is as described by Figure 4.1.

The firms problem is now a hybrid control-stopping-time problem; in addition to choosing a stopping time τ - or equivalently, an exit threshold p - the firm also solves a dynamic pricing problem. Of course, the subtlety of this control problem the current setting is that the pricing decision affects the firms present value not only through the usual demand schedule, but also through the informational rents that follow - that is, the SDE for ratings under full transparency, is:

$$dp_t = \frac{\sqrt{\pi(q_t, p_t) + \epsilon}}{\sigma} p_t(1 - p_t) d\bar{Z}_t, \quad (15)$$

where the service rate $\pi(\cdot)$ now depends explicitly on the chosen price via consumer choice. The firm's present value is now given by:

$$v(p) = \sup_{\tau^x, q \in \mathcal{Q}(p)} \mathbb{E}^x \left[\int_0^{\tau^x} e^{-(\rho+\delta)t} (q_t \pi(q_t, p_t) - c) dt | p_0 = p \right], \quad (16)$$

subject to (15), and where $\mathcal{Q}(p)$ denotes the set of processes progressively measurable with respect to \mathbb{F}^x and such that equation (15) has a strong solution, given that the process $(p_t)_{t \geq 0}$ is initially at p . Standards results in the theory of mixed stopping-control problems allow us to transform this expression into a tractable variational problem, itself a slightly extended version of the free boundary problem in Lemma 2.

Lemma 6. Let the pair $\{u(\cdot), \underline{p}\}$ for $u \in \mathcal{C}^1([0, 1])$ and $\underline{p} \in [0, 1]$ denote the variational inequality:

$$\begin{aligned}
\mathcal{A}_q\{u\} &= 0 \\
u(\underline{p}) &= 0 \\
u'(\underline{p}) &= 0 \\
u(p) &\geq 0 \quad \forall p \in [\underline{p}, 1] \\
u(p) &= 0 \quad \forall p \in [0, \underline{p}],
\end{aligned} \tag{17}$$

where

$$\mathcal{A}_q\{u\} = \sup_{q \in \mathbb{R}} \left\{ q\pi(p, q) + (\pi(p, q) + \epsilon) \frac{p^2(1-p)^2}{2\sigma^2} u''(p) \right\} - c - (\rho + \delta)u(p) \tag{18}$$

Suppose $\{V(\cdot), \underline{p}\}$ solves the problem (17), and that with respect to the process $(p_t)_{t \geq 0}$, the function $q(p)$ attains the supremum in (18), and that

$$\tau = \inf_s \{s \geq 0 : p_s \leq \underline{p}\} \tag{19}$$

Then the function $V(\cdot)$ coincides with the value $v(\cdot)$ in (16) and the pair $((q_t)_{t \geq 0}, \tau)$ attain the supremum in (16).

Thus, prices can be viewed as a function of firms' ratings. and thus the consumer's problem can yet again be written succinctly:

$$J = \max_{p \in [\underline{p}, 1], q \leq p} [(p - q) \Theta(p, q)], \tag{20}$$

Note that the indifference condition holds for all incentive-compatible prices, i.e. even those that are not charged by any firm on the equilibrium path. This embodies the *market utility (MU) assumption* commonly used in models of competitive search - were a firm to deviate in its posted price, the resulting queue-length it expects to find is as if the price were on-path. (see Galenianos and Kircher (2012) for details) A stationary equilibrium is thus defined as before, but with the

addition of the policy $\{q\}$ as per (18).

The MU assumption, combined with the piecewise-linear matching, tells us that if firms have such price flexibility, congestion will cease to exist in equilibrium. Were a firm charging a price q and facing a lengthy queue, they would deviate to charge a higher price q' that partially alleviates congestion in a manner that keeps consumers happy, while clearly benefiting the firm. An immediate corollary is that equilibrium prices are affine in quality:

Lemma 7. *There exists a $p^* \geq 0$ such that firms' equilibrium pricing policy $\{q\}$ satisfies $q(p) = p - p^*$. Furthermore, $\Theta(p) = 1$ for all $p \in [p, 1]$. Equilibrium consumer welfare equals p^* .*

Price floors and caps - At this stage, several observations regarding the above mechanism for price setting are worth making. First, the absence of congestion at all firms is inconsistent with a wealth of empirical findings.¹⁸ Indeed, studies show that less than 50% of the increase in demand that arises from higher ratings is passed into higher prices (Lewis and Zervas (2016)). In many of these cases, congestion is likely the result of firms' inability to price discriminate, e.g. healthcare practitioners. In the restaurant industry, the reason could either be menu costs, or other pro-social forces that exist outside of the model.¹⁹ In any case, a pattern shared by all these findings is that the highest-rated sellers tend to exhibit congestion. Beyond these empirical results, a theoretical literature also attempts to rationalize the presence of congestion externalities in two-sided markets and matching platforms. For instance, Eeckhout and Kircher (2010) and the competitive search literature at large employ frictional matching technologies - unlike the frictionless, linear rule I employ - that allow congestion and pricing to co-exist, but as such, embed an additional inefficiency.²⁰ More recently, Romanyuk and Smolin (forthcoming) outline a model with random matching that exhibits congestion at highly-rated sellers.

The outcome of such deliberations is that a model that exhibits some amount of price rigidity,

¹⁸See Andersen and Magruder (2012) for restaurants, Horton (2018) for evidence from AirB'nB, and the following link for a discussion on wait times in US healthcare: <https://www.hhnmag.com/articles/6417-the-push-is-on-to-eliminate-hospital-wait-times>.

¹⁹See Becker (1991) for an early theoretical model, and Hörner and Lambert (2018) for further empirical discussion.

²⁰Such models tend to exhibit concave matching rates, and as such would resemble my baseline setting even with pricing.

as well as some congestion at highly-rated firms in particular would best describe outcomes in such marketplaces. To this end, I introduce both lower and upper bounds on the prices available to firms to charge, i.e. the process $(q_t)_{t \geq 0}$ takes values in $[q, \bar{q}]$ for some $-\infty < q \leq \bar{q} \leq 1$. If $q = 1$ then the cap is inoperative, as any higher price will always be rejected by consumers.

Turning to the incumbent's problem, how does the additional tool of pricing help shape incentives and present values? Having pinned down prices as linear in rating, firms have one remaining choice - how many consumers to serve. That is, the service rate $\pi(p, q)$ is also a control variable for the firms. To understand this choice further, we can effectively re-write the HJB equation governing incumbents' value functions:

$$\begin{aligned} V(p) &= \max_{q \in [q, \bar{q}]} \left\{ \frac{1}{\rho + \delta} \left[q\pi(p, q) - c + (\pi(p, q) + \epsilon) \frac{p^2(1-p)^2}{2\sigma^2} V''(p) \right] \right\} \\ &= \max_{\pi \in [0, \bar{\lambda}]} \left\{ \frac{1}{\rho + \delta} \left[\left(p - p^* + \frac{p^2(1-p)^2}{2\sigma^2} V''(p) \right) \pi - c + \epsilon \frac{p^2(1-p)^2}{2\sigma^2} V''(p) \right] \right\} \end{aligned}$$

From this expression, it is clear that optimal service rates are bang-bang, i.e. there exists a \tilde{p} such that $\pi(p) = \mathbb{I}_{p \geq \tilde{p}}$, and so firms' flow profit function is piece-wise linear, while of course satisfying the market clearing condition:

$$\int_{\tilde{p}}^1 g(p) dp = 1 \quad (21)$$

In summary, equilibrium strategies are then characterized by the scalars $\{p, p^*, \tilde{p}, \eta\}$. This tractable structure allows me to derive the following results in closed-form:

Proposition 3. *There exists a $\tilde{q} \in (0, 1)$ such that:*

1. *If $\bar{q} < 1$ and $p_0 < \frac{K(\rho+\delta)}{\lambda\bar{q}-c} < \min\{\frac{p_0}{\bar{q}}, 1\}$, then $V(p)$ is S-shaped on $[p, 1]$.*
2. *If $\bar{q} = 1$ and $q \geq \tilde{q}$, then $V(p)$ is S-shaped on $[p, 1]$.*
3. *If $\bar{q} = 1$ and $q \leq \tilde{q}$, then $V(p)$ is convex on $[p, 1]$.*

Thus, if entry costs are sufficiently high, and a price cap is in place, the optimal rating design will again involve upper censorship. In this case, however, there is an intermediate phase of a

firm's life-cycle during which it operates at capacity, and the platform fully incorporates these reviews into the ratings process. If no price cap is in place, then owing to the convexity of the firm's problem - it is now information-loving throughout the ratings space - the optimal simple rating is fully transparent. However, in ongoing numerical work, I aim to show that more complex ratings systems dominate fully transparency while involving suppression of reviews for highly-rated firms. By pooling firms over an interval $[x, 1]$, firms with a rating of x get to charge a higher price than under full transparency, since consumers rationally believe that the firm has an average rating in this interval. This effectively transfers value from firms close to 1 to firms close to x , thus front-loading incentives: for high enough x , firms at x are necessarily closer to p_0 . Finally, it is worth noting that prices cannot depend on a firm's age, since consumers care only about the rating of a firm.²¹ Thus, phenomena such as introductory or predatory pricing are ruled out.

4.2 ADVERSE SELECTION

My analysis assumes that firms do not know their own quality and learn alongside the market, as in Jovanovic (1982), Board and Meyer-ter-Vehn (2014), Ericson and Pakes (1995). This is meant to capture the notion that the suitability of its product within a market is not fully understood by the firm *ex ante*. With asymmetric information, the observability of entry/exit decisions could allow firms to signal their quality through these actions, as in Bar-Isaac (2003), Atkeson et al. (2015). In particular, if each firm perfectly knew their type *ex ante*, free entry would dictate that only good firms enter. To see this, note that the continuation values for each type of firm now satisfy:

$$V_\theta(p) = \frac{1}{\rho + \delta} [\pi(p) - c + (\theta - p)V'(p) + \Sigma(p)V''(p)]$$

Since $V(p)$ is increasing, it is immediate that $V_1(p) > V_0(p)$ for all $p \in [0, 1]$. This would imply no *ex-post* heterogeneity, and in particular a degenerate ratings distribution. Such a profile is clearly

²¹Suppose a firm at state (p, t) trades at price q , whereas a firm at state (p, t') trades at price q' , for $t \neq t', q \neq q'$ (take $q > q'$ wlog). Consumer indifference then implies that $\Theta(p, t) \neq \Theta(p, t')$, and thus $\pi(p, t) = \pi(p, t')$. But then $q' \mapsto q$ would constitute a profitable deviation for the firm (p, t') .

violated empirically - see Section 5 and Appendix E for details - thus motivating the assumption that, at least to some degree, firms do not fully know their quality *ex ante*. I further assume that types are fixed, whereas one could imagine types can evolve according to investment (Board and Meyer-ter-Vehn (2014)).

5 CALIBRATION

The positive theory in section 2 combined with the normative theory of section 3 yields a clear counterfactual hypothesis; were platforms that currently use a policy of full transparency to adopt the optimal policy described by Theorem 2, consumer welfare would increase. In this section, I quantify this claim by calibrating the baseline model to data provided by Yelp! It should be stressed that the model was kept purposefully simple in order to maintain tractability and provide clear insights, and thus the quantitative findings reported here should be interpreted as back-of-the-envelope calculations.

Full transparency seems an obvious benchmark, not only from a theoretical standpoint, but also in light of the numerous predictions it generates that echo existing empirical results. However, for my approach to be valid, it is worth discussing the assumption with regards the data used here. Yelp! report that approximately 25% of reviews are excluded from the ratings process, suggesting a significant departure from full transparency. That said, their stated reason is quite different; they only include reviews that "...contribute reliable and useful content".²² Insofar as maintaining maximal accuracy and full transparency are qualitatively identical, my approach seems well justified.

The parameters to be determined are: the discount rate ρ , the attrition rate δ , the rate of background learning ϵ , the standard deviation of consumer reviews, σ , the capacity constraint $\bar{\lambda}$, firms' flow and entry costs c, K , and the fraction of high-quality entrants, p_0 . Whereas the measure of consumers was normalized to 1 in the theoretical analysis, I will parametrize it by B here, in order to make meaningful statements regarding relative populations. See Appendix E.1 for

²²See https://www.yelp-support.com/article/Does-Yelp-recommend-every-review?l=en_US.

Table 1: **Calibrated Model**

Parameter	Value	Quantity	FT	50% sup	90% sup
ρ	0.05	η	0.363	0.374	0.389
δ	0.18	p^*	0.606	0.620	0.647
σ	2	\underline{p}	0.474	0.474	0.474
ϵ	0.2	Av. lifespan	4.13	4.19	4.35
$\bar{\lambda}$	0.7	Av. quality	0.616	0.611	0.607
c	0.1	Exit rate	0.170	0.167	0.166
K	0.2	Exit ratio	46.8%	44.7%	42.7%
p_0	0.55				
B	0.28				

Average lifespan is calculated as below. Average quality is $\int_{\underline{p}}^1 pf(p)dp / \int_{\underline{p}}^1 f(p)dp$. The exit rate is $\Sigma(\underline{p})f'(\underline{p})$. The exit ratio is defined as $\frac{\Sigma(\underline{p})f'(\underline{p})}{\eta}$.

methodological details.

The left panel of Table 1 reports the calibrated parameters, while the right panel reports key equilibrium quantities under full transparency and two counterfactuals in which the stated fraction of reviews posted for firms with ratings above p^* are suppressed. Of particular importance is the expected lifetime of an incumbent firm with current rating p , $\mathcal{T}(p)$. In this setting, this is given by the mean first passage time to \underline{p} discounted by δ which satisfies:²³

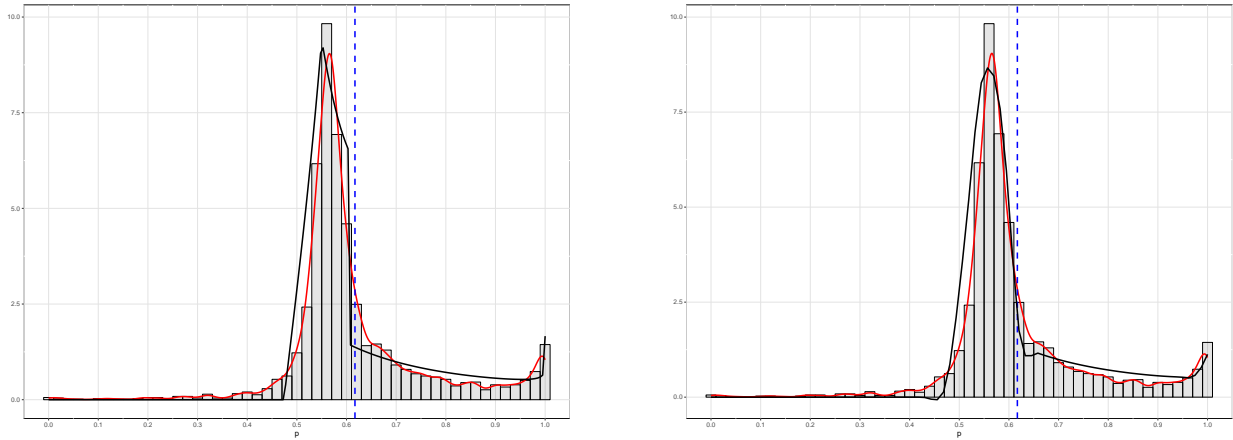
$$\delta\mathcal{T}(p) = 1 + \Sigma(p)\mathcal{T}''(p), \quad \mathcal{T}(\underline{p}) = 0, \quad \mathcal{T}_-(p^*) = \mathcal{T}_+(p^*)$$

I define the *exit ratio* as $\frac{\Sigma(\underline{p})f'(\underline{p})}{\eta}$. Interpreting exogenous attrition as a reduced form for any firm exit unrelated to ratings, the exit ratio is the proportion of exit caused by low ratings - the model suggests that roughly 47% of firms that shut down do so due to low ratings, a figure that would drop by roughly 4% under the optimal ratings policy.

The key statistics relate to improvements in consumer welfare - the platform's objective - under various regimes of suppression. Were the platform to throw away half the reviews generated by its users (i.e. consumers searching at a rating p^* and above), consumer welfare would increase by 2.3%.

²³See Moscarini (2005), Jovanovic (1982) for similar expressions, and Appendix A for a detailed derivation for the current setting.

Figure 3: Calibrated Model



Dotted blue line: empirical mean. Bars: empirical histogram density. Red lines: empirical kernel density. Black lines: theoretical ratings distribution. Left panel: pure calibrated model. Right panel: smoothed fit using a Loess line.

Were this to increase to 90%, consumer welfare would increase by 6.8%. Beyond this findings, the model estimates that the population of platform users is three times that of non-users. Perhaps this seems surprisingly large, in light of the assumption that the background term ϵ should be thought of as relatively small. One explanation is that the sample in question here is that of Las Vegas, well-known for supporting large numbers of tourists who perhaps do not use such platforms over short stays. As expected, stronger upper censorship boosts entry rates and reduces the exit ratio. As a result, firms live longer on average. That said, the large increase in consumer standards offsets these gains and leads to an mild increase in the exit threshold. Most interestingly, the average quality of incumbent firms *decreases* under these market conditions.²⁴ If one interprets background learning as driven by consumers that do not use the platform to direct their search as consequently ignored from the platform’s objective, then the optimal policy increases the welfare of users while decreasing the welfare of non-users.

²⁴The prediction is ambiguous here - for one, suppression has ambiguous equilibrium effects on the exit threshold.

6 CONCLUSION

This paper began with the basic observation that consumer reviews can form a *barrier to entry* for new firms. I built a tractable, equilibrium model of dynamic platform formation in which firms with heterogeneous, hidden types make entry and exit decisions, and whose type is gradually revealed via consumer reviews. A *rating platform* controls the inclusion of these reviews into a firm’s rating. Ratings guide consumer search and thus provide firms with incentives to remain active. Under a regime of *full transparency*, in which the platform includes all consumer reviews received into a firm’s rating, consumers flock to the highest-rated firms, while ignoring struggling firms altogether. This demand cliff induces a powerful non-convexity in the firm’s problem: its continuation value is S-shaped as a function of its rating. Thus, while struggling firms covet feedback and established firms dislike feedback, *precisely the reverse* profile obtains. Turning to design, I modelled the platform as maximizing consumer welfare via its control of the rating process. The central result was that optimal rating design involves *upper censorship* — the exclusion of reviews from established firm’s ratings — as a means of making the task of surmounting the ratings hill less daunting, thus stimulating participation. Finally, I performed an exploratory calibration using data from Yelp!, showing how consumer welfare would increase by roughly 7% under the optimal policy.

REFERENCES

- Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar**, “Fast and Slow Learning From Reviews,” *working paper*, 2018.
- Andersen, M. and J. Magruder**, “Learning from the Crowd: Regression Discontinuity Estimates of the Effects of an Online Review Database,” *Economic Journal*, 2012, 122 (563), 957–989.
- Atkeson, A., C. Hellwig, and G. Ordoñez**, “Optimal Regulation in the Presence of Reputation Concerns,” *Quarterly Journal of Economics*, 2015, 130 (1), 415–464.

- Bagwell, K.**, “Informational Product Differentiation as a Barrier to Entry,” *International Journal of Industrial Organization*, 1990, 8, 207–223.
- Banerjee, A. V.**, “A simple model of herd behavior,” *Quarterly Journal of Economics*, 1992, 107 (3), 797–817.
- Bar-Isaac, H.**, “Reputation and Survival: Learning in a Dynamic Signalling Model,” *Review of Economic Studies*, 2003, 70 (2), 231–251.
- Becker, G.**, “A Note on Restaurant Pricing and Other Examples of Social Influences on Prices,” *Journal of Political Economy*, 1991, 99, 1109–1116.
- Bergemann, D. and J. Välimäki**, “Market Diffusion with Two-Sided Learning,” *RAND Journal of Economics*, 1997, 28 (4), 773–795.
- Biglaiser, G.**, “Middlemen as Experts,” *RAND Journal of Economics*, 1993, 24 (2), 212–223.
- 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.
- Bloedel, A. and I. Segal**, “Persuasion with Rational Inattention,” *working paper*, 2018.
- Board, S. and M. Meyer-ter-Vehn**, “A Reputational Model of Firm Dynamics,” *Working Paper*, 2014.
- Bolton, P. and C. Harris**, “Strategic Experimentation,” *Econometrica*, 1999, 67 (2), 349–374.
- Burdett, K. and G. Menzio**, “The (Q,S,s) Pricing Rule,” *Quarterly Journal of Economics*, 2018, 85 (2), 892–928.
- Campbell, A., C.M.Leister, and Y. Zenoux**, “Word-of-Mouth Communication and Search,” *Working paper*, 2018.
- Che, Y-K. and J. Hörner**, “Recommender Systems as Mechanisms for Social Learning,” *Quarterly Journal of Economics*, forthcoming.

- Chevalier, J.A. and D. Mayzlin**, “The Effect of Word of Mouth on Sales: Online Book Reviews,” *Journal of Marketing Research*, 2006, 43 (3), 345–354.
- Dai, W., G. Jin, J. Lee, and M. Luca**, “Aggregation of Consumer Ratings: An Application to Yelp.com,” *Quantitative Marketing and Economics*, forthcoming.
- Eeckhout, J. and P. Kircher**, “Sorting and Decentralized Price Competition,” *Econometrica*, 2010, 78 (2), 539–574.
- Engelbert, H. J. and W. Schmidt**, “Strong Markov continuous local martingales and solutions of one-dimensional stochastic differential equations,” *Mathematische Nachrichten*, 1991, 151 (1), 149–197.
- Ericson, R. and A. Pakes**, “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *Review of Economic Studies*, 1995, 62 (1), 53–82.
- Farronato, C. and A. Fradkin**, “The Welfare Effects of Peer Entry in the Accommodations Market: The Case of Airbnb,” *working paper*, 2018.
- Feller, W.**, “Diffusion Processes in One Dimension,” *Transactions of the American Mathematical Society*, 1954, pp. 988–1012.
- Fradkin, A.**, “Search, Matching, and the Role of Digital Marketplace Design in Enabling Trade,” *working paper*, 2018.
- Gabaix, X., J-M. Lasry, P-L. Lions, and B. Moll**, “The Dynamics of Inequality,” *Econometrica*, 2016, 84 (6), 2071–2111.
- Galenianos, M. and P. Kircher**, “On the Game-theoretic Foundations of Competitive Search Equilibrium,” *International Economics Review*, 2012, 53, 1–21.
- Gardiner, C.**, *Stochastic methods: A handbook for the natural sciences* 2009.
- Goel, A.M. and A.V. Thakor**, “Information reliability and welfare: A theory of coarse credit ratings,” *Journal of Financial Economics*, 2015.

- Grossman, G.M. and H. Horn**, “Infant-Industry Protection Reconsidered: The Case of Informational Barriers to Entry,” *Quarterly Journal of Economics*, 1988, *103* (4), 767–787.
- 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.
- Hagiu, A. and B. Jullien**, “Why do intermediaries divert search?,” *RAND Journal of Economics*, 2011, *42* (2).
- 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,” *Working Paper*, 2018.
- Horton, J.**, “Buyer Uncertainty about Seller Capacity: Causes, Consequences, and a Partial Solution,” *Management Science*, *forthcoming*, 2018.
- Hui, X., M. Saeedi, G. Spagnalo, and S. Tadelis**, “Certification, Reputation and Entry: An Empirical Analysis,” *NBER Working Paper*, 2018, (24916).
- Jovanovic, B.**, “Selection and the Evolution of Industry,” *Econometrica*, 1982, *50* (3), 649–670.
- Karatzas, I. and H. Wang**, “Utility maximization with discretionary stopping,” *SIAM Journal on Control and Optimization*, 2001, *39*, 306–329.
- and **S.E. Shreve**, *Methods of Mathematical Finance* 1998.
- Keilson, J. and J.A. Wellner**, “Oscillating Brownian motion,” *Journal of Applied Probability*, 1978, *15* (2), 300–310.
- Kolotilin, A., T. Mylovanov, A. Zapelchelnjuk, and M. Li**, “Persuasion of a Privately Informed Receiver,” *Econometrica*, 2017, *85* (6), 1949–1964.
- Kovbasyuk, S. and G. Spagnolo**, “Memory and Markets,” *working paper*, 2018.

- 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,” *working paper*, 2018.
- le Gall, J-F.**, “One-dimensional stochastic differential equations involving the local times of the unknown process,” *Stochastic Analysis. Lecture Notes Math.*, 1985, 1095, 5182.
- Lester, B.**, “Information and Prices with Capacity Constraints,” *American Economic Review*, 2011, 101 (4), 1591–1600.
- Lewis, G. and G. Zervas**, “The Welfare Impact of Consumer Reviews: A Case Study of the Hotel Industry,” *Working Paper*, 2016.
- Li, L., S. Tadelis, and X. Zhou**, “Buying Reputation as a Signal of Quality: Evidence from an Online Marketplace,” *working paper*, 2018.
- Li, X. and L. Hitt**, “Self-Selection and Information Role of Online Product Reviews,” *Information Systems Research*, 2008, 19 (4).
- Lipster, R. and A. Shryaev**, *Statistics of Random Processes* 1977.
- Lizzeri, A.**, “Information Revelation and Certification Intermediaries,” *RAND Journal of Economics*, 1999, 30 (2), 214–231.
- Luca, D.L and M. Luca**, “Survival of the Fittest: The Impact of the Minimum Wage on Firm Exit,” *working paper*, 2018.
- Luca, M.**, “Reviews, Reputation, and Revenue: The Case of Yelp.com,” *working paper*, 2016.
- Luo, T. and P.B. Stark**, “Only the Bad Die Young: Restaurant Mortality in the Western US,” *working paper*, 2014.
- Luttmer, E.G.**, “Selection, Growth, and the Size Distribution of Firms,” *The Quarterly Journal of Economics*, 2007, 122 (3), 1103–1144.

- Marinovic, I., A. Skrzypacz, and F. Varas**, “Dynamic Certification and Reputation for Quality,” *AEJ: Microeconomics*, 2018, *10* (2), 58–82.
- 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.
- Powell1, D., J. Yu, M. DeWolf, and K.J. Holyoak**, “The Love of Large Numbers: A Popularity Bias in Consumer Choice,” *Psychological Science*, 2017, *28* (10), 14321442.
- Romanyuk, G. and A. Smolin**, “Cream Skimming and Information Design in Matching Markets,” *AEJ: Microeconomics*, forthcoming.
- 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.
- Schmalensee, R.**, “Product Differentiation Advantages of Pioneering Brands,” *American Economic Review*, 1982, *72* (3), 349–365.
- Stokey, N.**, *The Economics of Inaction* 2008.
- Tadelis, S.**, “Reputation and Feedback Systems in Online Platform Markets,” *Annual Review of Economics*, 2016, *8*, 321–340.
- **and F. Zettelmeyer**, “Information Disclosure as a Matching Mechanism: Theory and Evidence from a Field Experiment,” *American Economic Review*, 2015, *105* (2), 886–905.
- Yang, K.H.**, “Non-linear Pricing of Information Structure: An Equivalence Result,” *working paper*, 2018.

Appendices

A DERIVATION OF LOCAL TIME FOR RATINGS PROCESS

In this section, I derive the general expression of the so-called *local time* of a firm's rating process. See Stokey (2008) for an excellent treatment of this concept.

Let the ratings process solve the SDE given by (2) defined over $[0, 1]$. This process is formally an *oscillating diffusion process*, for which the existence of a strong solution was shown in le Gall (1985) and explicit formulae for the density and occupation times were derived in Keilson and Wellner (1978). In particular, p_t converges to either 0 or 1 when unstopped, but achieves neither state in finite time. Now, let τ be the \mathbf{F}^x -measurable stopping time defined as $\tau = T(\underline{p})$, where $T(p) = \min\{t : p_t = p\}$ is the first hitting time of the ratings process for a given level. Since $T(1) = \infty$, it follows that $\tau = T(\underline{p}) \wedge T(1)$. Taking expectations of Theorem 3.7 in Stokey (2008), the firm's value function can be re-written:

$$\begin{aligned} v(\hat{p}) &= \mathbb{E}^x \left[\int_0^\tau e^{-(\rho+\delta)t} (\lambda(p_t) - c) dt | p_0 = \hat{p} \right] \\ &= \mathbb{E}^x \left[\int_0^1 l(p; \hat{p}, \tau, \rho + \delta) (\lambda(p) - c) dp \right] \\ &= \int_{\underline{p}}^1 L(p; \hat{p}, \underline{p}, 1, \rho + \delta) (\lambda(p) - c) dp \end{aligned}$$

where $l(p; \hat{p}, \tau, \rho + \delta)$ is the discounted local time for the ratings process defined as in Theorem 3.4 in Stokey (2008),

$$L(p; \hat{p}, a, b, \rho + \delta) = \mathbb{E}^x [l(p; \hat{p}, T(a) \wedge T(b), \rho + \delta)] \quad \text{for } a, b \in [0, 1]$$

The purpose of this appendix is to derive closed-form solutions for the function $L(p; \hat{p}, a, b, \rho + \delta)$ where a, b are such that $\tilde{\lambda}(p) = \lambda \in \mathbb{R}_+$. To this end, let $T = T(a) \wedge T(b)$, and define

$$w(\hat{p}) = \mathbb{E}^x \left[\int_0^\tau e^{-rt} g(p_t) dt | p_0 = \hat{p} \right] \tag{A.1}$$

for some $g \in \mathcal{C}^2([0, 1])$. Standard arguments yield that w solves the second-order ODE

$$rw(p) = g(p) + \Sigma(p)w''(p), \tag{A.2}$$

where $\Sigma(p) = \frac{\lambda}{2\sigma^2}p^2(1-p)^2$ and subject to the boundary conditions $w(a) = w(b) = 0$. The two independent general solutions to equation (A.2) are

$$w_1(p) = p^\gamma(1-p)^{1-\gamma}$$

$$w_2(p) = p^{1-\gamma}(1-p)^\gamma$$

where $\gamma = \frac{1}{2} + \sqrt{\frac{1}{2} + \frac{2r\sigma^2}{\lambda}}$. A particular solution for (A.2) is

$$\bar{w}(p) = \sum_{i=1}^2 \xi_i(p)w_i(p)$$

where $\xi_i(p), i = 1, 2$ are functions to be determined. Conjecture that $\sum_{i=1}^2 \xi'_i(p)w_i(p) = 0$, in which case $\Sigma(p) \sum_{i=1}^2 \xi_i(p)w_i(p) = -g(p)$ and thus for the conjecture to be true, it must be that

$$\xi'_1(p) = \frac{w_2(p)g(p)}{\mathbb{W}(w_1, w_2)(p)\Sigma(p)}, \quad \xi'_2(p) = \frac{w_1(p)g(p)}{\mathbb{W}(w_1, w_2)(p)\Sigma(p)},$$

where $\mathbb{W}(w_1, w_2) = w_1w'_2 - w_2w'_1$ is the Wronskian. Direct calculation (omitted algebra) results in

$$\xi'_1(p) = \frac{2\sigma^2}{(1-2\gamma)\lambda}p^{-1-\gamma}(1-p)^{\gamma-2}g(p), \quad \xi'_2(p) = \frac{2\sigma^2}{(1-2\gamma)\lambda}p^{\gamma-2}(1-p)^{-1-\gamma}g(p),$$

and thus

$$\xi_1(p) = \frac{2\sigma^2}{(1-2\gamma)\lambda} \int_a^p z^{-1-\gamma}(1-z)^{\gamma-2}g(z)dz, \quad \xi_2(p) = \frac{2\sigma^2}{(1-2\gamma)\lambda} \int_p^b z^{\gamma-2}(1-z)^{-1-\gamma}g(z)dz$$

Combining the expressions for the general and particular solutions for $w(p)$, and imposing the boundary

conditions $w(a) = w(b) = 0$ gives rise to the full expression:

$$\begin{aligned}
w(p) &= \xi_1(p)w_1(p) + \xi_2(p)w_2(p) - \psi(p)\xi_2(a)w_2(a) - \Psi(p)\xi_1(b)w_1(b) \\
&= \frac{2\sigma^2}{(1-2\gamma)\lambda} \left[\int_a^p \left(\frac{p}{z}\right)^\gamma \left(\frac{1-p}{1-z}\right)^{1-\gamma} \frac{1}{z(1-z)} g(z) dz \right. \\
&\quad + \int_p^b \left(\frac{p}{z}\right)^{1-\gamma} \left(\frac{1-p}{1-z}\right)^\gamma \frac{1}{z(1-z)} g(z) dz \\
&\quad - \psi(p)a^{1-\gamma}(1-a)^\gamma \int_a^b z^{-1-\gamma}(1-z)^{\gamma-2} g(z) dz \\
&\quad \left. - \Psi(p)b^\gamma(1-b)^{1-\gamma} \int_a^b z^{\gamma-2}(1-z)^{-1-\gamma} g(z) dz \right]
\end{aligned} \tag{A.3}$$

where

$$\psi(p) = \frac{w_1(p)w_2(b) - w_1(b)w_2(p)}{w_1(a)w_2(b) - w_1(b)w_2(a)}, \quad \Psi(p) = \frac{w_1(p)w_2(a) - w_1(a)w_2(p)}{w_1(b)w_2(a) - w_1(a)w_2(b)}$$

To compute $L(p; \hat{p}, a, b, r)$, set $g(p) = 1$ and take derivatives of equation (A.3):

Proposition A.1. *Suppose $a, b \in [0, 1]$ such that $\tilde{\lambda}(p) = \lambda \in \mathbb{R}_+$ for all $p \in [a, b]$. Let $L(p; \hat{p}, a, b, r)$ denote the expected discounted local time for the ratings process p between levels a, b , with initial rating \hat{p} and discount rate r . Then*

$$\begin{aligned}
L(p; \hat{p}, a, b, r) &= \frac{2\sigma^2}{(1-2\gamma)\lambda} \left[\Xi(p; \hat{p}) - \psi(\hat{p})a^{1-\gamma}(1-a)^\gamma p^{-1-\gamma}(1-p)^{\gamma-2} \right. \\
&\quad \left. - \Psi(\hat{p})b^\gamma(1-b)^{1-\gamma} p^{\gamma-2}(1-p)^{-1-\gamma} \right]
\end{aligned} \tag{A.4}$$

where

$$\Xi(p; \hat{p}) = \begin{cases} \left(\frac{\hat{p}}{p}\right)^\gamma \left(\frac{1-\hat{p}}{1-p}\right)^{1-\gamma} & \text{if } p \leq \hat{p} \\ \left(\frac{\hat{p}}{p}\right)^{1-\gamma} \left(\frac{1-\hat{p}}{1-p}\right)^\gamma & \text{if } p \geq \hat{p} \end{cases}$$

B MISCELLANEOUS PROOFS

B.1 LEMMA 2

Note that the conditions of Rüschemdorf and Urusov (2008) hold, since the SDE for the state variable p has no drift coefficient and a diffusion coefficient given by $\lambda(p)p(1-p)$, which is locally bounded since

$\lambda(p) \in [0, \bar{\lambda} + \epsilon]$ for all $p \in [0, 1]$. Their Theorem 2.1 implies that the unique value function satisfying (3) belongs to $\mathcal{C}^1([0, 1])$, and solves (4), and that \underline{p} is unique. Finally, boundedness of the value function follows from the Extreme Value Theorem.

B.2 PROPOSITION 1

Here, I derive the boundary conditions for the Fokker-Planck equation (7), 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) 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 Brownian motion 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 (7) directly and use the Fundamental Theorem of Calculus).

Thus, for the ratings process defined by the SDE (2),

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

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

where $\Sigma(p)$ is as defined in Lemma 2. Standard results in statistical mechanics then imply that the transition measures form a root of the adjoint operator to (B.1).

Remark 1. For the process Y defined above, the stationary distribution satisfies $\mathcal{A}_f f = 0$, where $\mathcal{A}_f = \mathcal{A}_b^*$.

This verifies that f must solve (7). 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 B.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*

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

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

Proof.

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

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

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

1. $f(\underline{p}) = 0, \Sigma(1)f(1) = 0$

We have that $f(\underline{p})\Sigma(\underline{p})u' = 0$ for all u . Since $\underline{p} > 0$, it follows that $\Sigma(\underline{p}) \neq 0$, and hence $f(\underline{p}) = 0$. This is the standard ‘‘attainable boundary’’ condition (see Feller (1954) for a comprehensive classification of boundary conditions for diffusion processes over bounded intervals). The second condition follows similarly. This is a ‘‘natural boundary’’ condition, as $p = 1$ cannot be attained in finite time.

2. $\Sigma(p_0) [f'_-(p_0) - f'_+(p_0)] = \eta, \frac{\partial}{\partial p} \Sigma(p) f(p) \in \mathcal{C}^1([\underline{p}, p_0) \cup (p_0, 1]), \Sigma(p) f(p) \in \mathcal{C}^1([\underline{p}, 1])$

An inflow rate of η yields the boundary condition $[\Sigma(p_0)u(p_0)]_{-}^{\pm} = \eta$ on any u that solves $\mathcal{A}_f\{u\} = 0$. Furthermore, since u and u' are continuous, this implies that:

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

Intuitively, the condition states that the total outflow of mass from p_0 , given by $[\mathcal{J}\{f\}(p_0)]_{-}^{\pm}$, 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 that 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.

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

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

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

where all other terms in $[(\Sigma f)']_{\mathcal{D}}$ vanish due to continuity of $(\Sigma f)'$ and $f(\underline{p}) = 0$.

B.3 LEMMA 3

First, note that in the presence of background learning at rate $\epsilon > 0$, the stationary distribution must admit a density with full support on $[\underline{p}, 1]$. To obtain indifference in (9), it must be that for any rating visited in equilibrium, congestion must (weakly) occur, i.e. $\frac{g(p)}{\lambda f(p)} \geq 1$ 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 $\frac{g(p_1)}{\lambda f(p_1)} < 1, \frac{g(p_2)}{\lambda f(p_2)} < 1$. 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 $\pi(p) \in \{0, \bar{\lambda}\}$. 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 $\pi(p) = \bar{\lambda}$ for $p \geq p^*$. This also establishes the conjectured form of $g(p)$.

Finally, that $p^* > 0$ is a simple consequence of maintaining indifference.

Finally, I argue that $p^* > \underline{p}$. First, since $p^* > 0$, the qualifier in Lemma 2 obtains, and thus $\underline{p} > 0$ in equilibrium. Suppose instead that $p^* \leq \underline{p}$. Then $\pi(p) = \bar{\lambda}$ for all $p \in [\underline{p}, 1]$, and it is then readily shown that the firm's value satisfies $V(p) = \frac{\bar{\lambda} - c}{\rho + \delta}$ for all $p \in [\underline{p}, 1]$, violating the boundary condition $V(\underline{p}) = 0$.

B.4 LEMMA 7

Fix an equilibrium price path $(q_t)_{t \geq 0} = \{q(p)\}$ and associated value function $V(p)$, and suppose there exists a p such that $\Theta(p, q(p)) < 1$. Let $\Theta(p, q(p)) = 1 - \psi$ for some $\psi > 0$. I claim there exists $q' > q(p)$ such that $(p - q') > (p - q(p)) \Theta(p, q(p))$, i.e. such a price offers consumers a strictly higher expected value. To see this, let $q' = q + \chi$. Then:

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

which holds for sufficiently small χ . (note that $p - q(p) > 0$ in equilibrium) Thus, under the market utility assumption, $\Theta(p, q(p)) < \Theta(p, q') \leq 1$ must hold. Finally, note that such a price offer would constitute a profitable deviation for a firm with rating p , since $\Theta(p, q') \leq 1$ implies that $\pi(p, q') = \bar{\lambda}$, and hence:

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

B.5 PROPOSITION 3

1. I claim that there exists a $\bar{p} < 1$ such that the equilibrium price path $\{q(p)\}$ is such that $q(p) = \bar{q}$ for all $p \in [\bar{p}, 1]$.

Suppose not, i.e. all firms charge less than the price cap, and thus $q(p) = p - a$ for some a such that $1 - a \leq \bar{q}$. Then I claim that $V(p)$ is strictly convex on $[\underline{p}, 1]$. To see this, note that:

$$\pi(p) = \begin{cases} \bar{\lambda}(p - a) & \text{if } p \geq \tilde{p} \\ 0 & \text{if } p < \tilde{p} \end{cases} \quad (\text{B.5})$$

and hence

$$V(p) = \frac{1}{\rho + \delta} \left[\bar{\lambda}(p - a) - c + (\bar{\lambda} + \epsilon) \frac{p^2(1-p)^2}{2\sigma^2} V''(p) \right] \quad (\text{B.6})$$

The general solution to equation (B.6) is as for (5), while the particular solution is now affine. Thus, the full solution is:

$$V(p) = \begin{cases} c_1^0 p^{1-\gamma_1^v} (1-p)^{\gamma_1^v} + \frac{\bar{\lambda}(p-a)-c}{\rho+\delta} & \text{if } p \geq p^* \\ c_0^0 p^{1-\gamma_0^v} (1-p)^{\gamma_0^v} + c_1^1 p^{\gamma_0^v} (1-p)^{1-\gamma_0^v} - \frac{c}{\rho+\delta} & \text{if } p < p^* \end{cases} \quad (\text{B.7})$$

where the exponents γ_0^v, γ_1^v and coefficients c_0^0, c_1^0, c_1^1 solve the matrix as detailed in Proposition C.1. Strict convexity is now easily verified. Thus, $V(p) < \frac{p[\bar{\lambda}\bar{q}-c]}{\rho+\delta}$, i.e. a straight line connecting 0 with $\frac{[\bar{\lambda}\bar{q}-c]}{\rho+\delta}$ and hence the free entry condition is violated. This proves the claim.

Thus, for $p \in [\underline{p}, 1]$, $pi(p) = \bar{\lambda}\bar{q}$ and hence $V(p) \leq \frac{[\bar{\lambda}\bar{q}-c]}{\rho+\delta}$ which in turn implies that $V''(p) \leq 0$. Parts 2) and 3) are proven in a similar fashion.

C PROOF OF THEOREM 1

The outline of the proof is as follows:

1. Fix the entry rate η
 - (a) Solve explicitly for f for arbitrary \underline{p}, p^* .
 - (b) Solve explicitly for (V, \underline{p}) for arbitrary p^* .
 - (c) Recast \underline{p}, p^* as a pair of reaction correspondences, $R_-(p^*), R_*(\underline{p})$.
 - (d) Prove that these correspondences are functions, with precisely one intersection.

2. Prove that the value of entry $V(p_0)$ is decreasing in the entry rate, with $V(p_0) = K$ always obtaining for some unique entry rate $\eta \in (0, \infty)$.

C.1 COMPUTING THE RATINGS DISTRIBUTION

The entry rate η will be assumed fixed until later in the proof. As such, 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:

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

Case 1: $p^* > p_0$

The general solution to the Fokker-Planck equation (7) is obtained from the derivations outlined in Appendix A:

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

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

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

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

$$c_2^2 = 0$$

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

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

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

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

Case 2: $p^* < p_0$

The general solution to the Fokker-Planck equation (7) now becomes

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

As before, $c_2^2 = 0$, and so coefficients $\mathbf{c} = [c_0^1 \ c_0^2 \ c_1^1 \ c_1^2 \ c_2^1]^T$ solve $\mathbf{A}\mathbf{c} = \mathbf{b}$, where

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

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

C.2 COMPUTING THE FIRM'S VALUE FUNCTION

By Lemma 2, we know that the firm's value function belongs to $\mathcal{C}^1([0, 1])$. Thus, the boundary conditions for the associated variational problem reduce to:

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

The general and particular solutions to equation (5) can be found by applying the derivation obtained in Appendix A:

Proposition C.1. *Let*

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

Then firms' value function given by

$$V(p) = \begin{cases} c_1^0 p^{1-\gamma_1^v} (1-p)^{\gamma_1^v} + \frac{\bar{\lambda}-c}{\rho+\delta} & \text{if } p \geq p^* \\ c_0^0 p^{1-\gamma_0^v} (1-p)^{\gamma_0^v} + c_0^1 p^{\gamma_0^v} (1-p)^{1-\gamma_0^v} - \frac{c}{\rho+\delta} & \text{if } p < p^* \end{cases} \quad (\text{C.1})$$

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

$$\Theta_v = \begin{bmatrix} \left[\frac{1-\gamma_0}{p} - \frac{\gamma_0}{1-p} \right] p^{1-\gamma_0} (1-p)^\gamma & \left[\frac{\gamma_0}{p} - \frac{1-\gamma_0}{1-p} \right] p^{\gamma_0} (1-p)^{1-\gamma_0} & 0 \\ p^{*1-\gamma_0} (1-p^*)^{\gamma_0} & p^{*\gamma_0} (1-p^*)^{1-\gamma_0} & p^{*1-\gamma_1} (1-p^*)^{\gamma_1} \\ \left[\frac{1-\gamma_0}{p^*} - \frac{\gamma_0}{1-p^*} \right] p^{*1-\gamma_0} (1-p^*)^{\gamma_0} & \left[\frac{\gamma_0}{p^*} - \frac{1-\gamma_0}{1-p^*} \right] p^{*\gamma_0} (1-p^*)^{1-\gamma_0} & - \left[\frac{1-\gamma_1}{p^*} - \frac{\gamma_1}{1-p^*} \right] p^{*1-\gamma_1} (1-p^*)^{\gamma_1} \end{bmatrix}$$

and

$$\mathbf{b}_v = \begin{bmatrix} 0 \\ \frac{\bar{\lambda}}{\rho+\delta} \\ 0 \end{bmatrix}$$

and the threshold value p solves

$$c_0^0 p^{1-\gamma_0} (1-p)^{\gamma_0} + c_0^1 p^{\gamma_0} (1-p)^{1-\gamma_0} = \frac{c}{\rho + \delta} \quad (\text{C.2})$$

noting that the coefficients c_0^0, c_0^1 implicitly depend on p .

C.3 REACTION FUNCTIONS

In light of Lemmas 2 and 3, equilibrium strategies are fully summarized by the pair of scalars $\{p, p^*\}$.

Let $R_-(p^*)$ solve (C.2) with respect to p and $R_*(p)$ solve (10) subject to (11). That is, $R_-(p^*)$ computes optimal exit thresholds, fixing p^* , while $R_*(p)$ computes optimal demand thresholds p^* fixing p .

Lemma C.2. 1. Both $R_-(p^*)$ and $R_*(\underline{p})$ are functions on $[0, 1]$. $R_-(p^*) \in \mathcal{C}^1([0, 1])$, while $R_*(\underline{p}) \in \mathcal{C}([0, 1])$.

2. For all $p^* > 0$, $R_-(p^*) < p^*$. $R_-(0) = 0$.

3. $R_*(0) > 0$.

Proof. 1. That $R_-(p^*)$ is single-valued is a consequence of the uniqueness of \underline{p} in Lemma 2. To show that $R_*(\underline{p})$ is single-valued, note that it is easily verified that for profile $g(p)$ to solve (10), p^* must satisfy the equation:

$$p^* = \frac{\bar{\lambda}}{B} \int_{p^*}^1 p f(p; p^*) dp \quad (\text{C.3})$$

where $f(p; p^*)$ is used to denote the explicit dependence of f on p^* . It remains to prove that a solution to (C.3) exists, and is unique.


Suppose that $p^* > p_0$ (a similar argument will apply in the case $p^* < p_0$). Then (C.3) reduces to

$$\begin{aligned} p^* &= \frac{\bar{\lambda}}{B} \int_{p^*}^1 p f(p; p^*) dp \\ &= \frac{\bar{\lambda}}{B} \int_{p^*}^1 c_1^2(p^*) p^{-\gamma_1^f} (1-p)^{\gamma_1^f-2} dp \\ &= \frac{\bar{\lambda}}{B} c_1^2(p^*) \left[\frac{(1-p^*)^{\gamma_1^f-1} p^{*1-\gamma_1^f}}{\gamma_1^f-1} \right] \end{aligned}$$

Consider the function $h(x) = c_1^2(x)(1-x)^{\gamma_1^f-1} x^{1-\gamma_1^f} \equiv c_1^2(x)j(x)$. By the final part of the proof of Lemma 3, $h(0) > 0$, and $h(1) = 0$. Since $\gamma_1^f > 1$, $j(x)$ is decreasing in x . I now prove that $c_1^2(x)$ is decreasing in x .

To do so, I re-write the system of equations governing f , with the goal of splitting the system in two parts, defined by p_0 . This allow direct computation of derivatives with respect to p^* . To this end, re-write the boundary conditions $\Sigma(p_0)[f'_-(p_0) - f'_+(p_0)] = \eta$ as $\Sigma(p_0)f'_-(p_0) = \alpha$ and $-\Sigma(p_0)f'_+(p_0) = \eta - \alpha$ for some $\alpha \in \mathbb{R}$. Thus, f solves the forward equation (7) on $[\underline{p}, p_0)$ subject to the conditions $f(\underline{p}) = 0$ and $\Sigma(p_0)f'_-(p_0) = \alpha$, and on $(p_0, 1]$ subject to $\Sigma(1)f(1) = 0$ and $-\Sigma(p_0)f'_+(p_0) = \eta - \alpha$. The constant α is then pinned down by the aggregate condition 6 in Proposition 1.

Figure 4: Ratings Density $f(p)$ for fixed \underline{p} as p^* varies, $p_1^* > p_2^* > p_0 > p_3^*$



On $[\underline{p}, p_0)$, f is determined by the two coefficients $[c_0^1 c_0^2]^T$ that solve the system:

$$\begin{bmatrix} \alpha_0^1(\underline{p}) & \alpha_0^2(\underline{p}) \\ \frac{\partial \alpha_0^1}{\partial p}(p_0) & \frac{\partial \alpha_0^2}{\partial p}(p_0) \end{bmatrix} \begin{bmatrix} c_0^1 \\ c_0^2 \end{bmatrix} = \begin{bmatrix} 0 \\ \alpha \end{bmatrix} \quad (\text{C.4})$$

while on $(p_0, 1)$, f is determined by the three coefficients $[c_1^1 c_1^2 c_2^2]^T$ that solve the system:

$$\begin{bmatrix} \frac{\partial \alpha_0^1}{\partial p}(p_0) & \frac{\partial \alpha_0^2}{\partial p}(p_0) & 0 \\ \frac{\partial}{\partial p} \Sigma_0(p^*) \alpha_0^1(p^*) & \frac{\partial}{\partial p} \Sigma_0(p^*) \alpha_0^2(p^*) & \frac{\partial}{\partial p} \Sigma_1(p^*) \alpha_1^1(p^*) \\ \Sigma_0(p^*) \alpha_0^1(p^*) & \Sigma_0(p^*) \alpha_0^2(p^*) & \Sigma_1(p^*) \alpha_1^1(p^*) \end{bmatrix} \begin{bmatrix} c_1^1 \\ c_1^2 \\ c_2^2 \end{bmatrix} = \begin{bmatrix} \eta - \alpha \\ 0 \\ 0 \end{bmatrix} \quad (\text{C.5})$$

Direct computation then yields that c_2^2 is strictly decreasing in p^* , regardless of the value of α . In summary, we have shown that $h(0) > 0, h(1) = 0$ and $h(x)$ is strictly decreasing. Hence, by the Intermediate Value Theorem, there exists a solution to (C.3) (uniqueness obtains since h is injective), denoted p^* . If $p^* \geq \underline{p}$. Then set $R_*(\underline{p}) = p^*$. If $p^* < \underline{p}$, the unique profile that solves (9) subject to (10) is $g(p) = \frac{Bp f(p)}{\int_{\underline{p}}^1 p f(p) dp}$ and thus $R_*(\underline{p}) = \underline{p}$.

Finally, continuous differentiability of $R_-(p^*)$ follows from applying the Implicit Function Theorem to (C.2).

2. For the first part, see the final part of the proof of Lemma 3. $R_-(0) = 0$ then follows by applying the Sandwich Theorem.

3. Let $f_0(;, p^*)$ denote the ratings distribution under no exit, for a given p^* ²⁵. It is readily shown that $f_0(;, p^*)$ is strictly positive almost everywhere for all $p^* \in [0, 1]$, and hence there exists $\epsilon > 0$ such that the RHS of equation (C.3) evaluated with respect to f_0 must be greater than ϵ .

□

Lemma C.3. 1. $\frac{\partial R_-(p^*)}{\partial p^*} > 0$ for all $p^* \in [0, 1]$

2. There exists $\tilde{p} \in (0, p_0)$ such that $\frac{\partial R_*(\underline{p})}{\partial \underline{p}} < 0$ for all $\underline{p} \leq \tilde{p}$ and $R_-(p^*) = \underline{p}$ for all $\underline{p} \geq \tilde{p}$

Proof. 1. Take $p_1^* < p_2^*$. Then $\lambda(p; p_1^*) > \lambda(p; p_2^*)$ for all $p \in (p_1^*, p_2^*]$ and $\lambda(p; p_1^*) = \lambda(p; p_2^*)$ everywhere else.

Hence, by equation (3), a firm's value is strictly decreasing in p^* . The result then follows immediately from the conditions (4).

2. I first show that as the exit threshold decreases, more firms exist *at all ratings*, i.e. if $p_1 < p_2$, then $f(p, p_1) > f(p, p_2)$ for all $p > p_1$. The result then follows immediately by expression (C.3) - a higher \underline{p} makes the RHS smaller, and hence p^* must decline to maintain equality.

I utilize the decomposition of f introduced in the proof of Lemma C.2 part 1. Take the system (C.4).

Direct computation yields that both c_0^1, c_0^2 are decreasing in \underline{p} . Thus, $f_-(p_0)$ is decreasing in \underline{p} . Now

consider the system (C.5). Recall that f is continuous at p_0 , as obtained in the proof of Proposition

1. As such, the boundary condition $-\Sigma(p_0)f'_+(p_0) = \eta - \alpha$ can be replaced by the simpler condition

$f_+(p_0) = f_-(p_0)$, transforming the system into:


$$\begin{bmatrix} \alpha_0^1(p_0) & \alpha_0^2(p_0) & 0 \\ \frac{\partial}{\partial p} \Sigma_0(p^*) \alpha_0^1(p^*) & \frac{\partial}{\partial p} \Sigma_0(p^*) \alpha_0^2(p^*) & \frac{\partial}{\partial p} \Sigma_1(p^*) \alpha_1^1(p^*) \\ \Sigma_0(p^*) \alpha_0^1(p^*) & \Sigma_0(p^*) \alpha_0^2(p^*) & \Sigma_1(p^*) \alpha_1^1(p^*) \end{bmatrix} \begin{bmatrix} c_1^1 \\ c_1^2 \\ c_2^2 \end{bmatrix} = \begin{bmatrix} f_-(p_0) \\ 0 \\ 0 \end{bmatrix} \quad (\text{C.6})$$

From this expression, it is then immediate that c_1^1, c_1^2, c_2^2 are all increasing in $f_-(p_0)$ and hence decreasing in \underline{p} .

□

²⁵Technically, I must adjust the boundary conditions appropriately, replacing condition 1 with the condition $\Sigma(0)f(0) = 0$ and amending condition 6 to read $\delta \int_0^1 f(p) dp = \eta$.

Figure 5: Ratings Density $f(p)$ for fixed p^* as \underline{p} varies, $\underline{p}_1 > \underline{p}_2 > \underline{p}_3$



C.4 FIXED POINT

Combining the various results of the previous section yields a simple fixed point problem for the pair of functions R_- and R_* .


Lemma C.4. *For fixed $\eta > 0$, there exists a unique tuple $\{V, f, J, \underline{p}, p^*\}$ that solve (4), (7), (9), (10), with g defined as in Lemma 3.*

Proof. $R_-^{-1}(p)$ is strictly increasing on $[0, \tilde{p}]$ with $R_-^{-1}(0) = 0$, $R_-^{-1}(\tilde{p}) > \tilde{p}$, while $R_*(p)$ is strictly decreasing on $[0, \tilde{p}]$ with $R_*(0) > 0$ and $R_*(\tilde{p}) = \tilde{p}$. The uniqueness of the pair \underline{p}, p^* follows from the Intermediate Value Theorem.. The uniqueness of V, J, f then follows from Lemma 2, Lemma C.2 and the derivation in Section C.1. □

C.5 ENTRY

To summarize, I have shown that for a fixed entry rate η , firms' exit and consumers' demand thresholds are uniquely determined, in turn uniquely determining the ergodic ratings distribution. I now complete the proof of Theorem 1 by endogenizing η and proving that a unique η closes the system. Note that

Figure 6: Reaction Functions



η effects an incumbent's value only indirectly through its effect on p^* . Let $p^*(\eta), V(p_0; \eta)$ denote these relationships.

Lemma C.5. $f(p; \eta)$ is homogeneous of degree 1 in η , and hence $p^*(\eta)$ is strictly increasing.

Proof. By Lemma 2, the coefficients determining f solve $\mathbf{Ac} = \eta \mathbf{1}$. Thus, using equation (C.3),

$$p^*(\eta) = \frac{\bar{\eta}\bar{\lambda}}{B} \int_{p^*}^1 pf(p; p^*)dp \quad (\text{C.7})$$

Implicit differentiation yields

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

□

as the integral is positive and strictly decreasing in p^* - see Lemma C.2.

Lemma C.6. 1. $V(p_0; \eta)$ is strictly decreasing in η

2. $\lim_{\eta \rightarrow 0} V(p_0; \eta) = \frac{\bar{\lambda} - c}{\rho + \delta}$

3. $\lim_{\eta \rightarrow \infty} V(p_0; \eta) = 0$

Proof. 1. Follows from Lemma C.5 and Lemma C.3.

2. First, note that as $\eta \rightarrow 0$, $p^* \rightarrow 0$. This follows from the fact that the integrand in (C.7) is bounded above for an entry rate of 1. The remainder of the proof is as follows. Take p^* small, and consider a slightly higher rating. This rating is also small, and hence the local evolution of the process is arbitrarily slow. As such, the firm's value is approximately equal to the discounted flow payoff of selling out. The result then follows from monotonicity of the firm's value function.

Formally, set $p^* = \epsilon > 0$ and $\hat{p} = 2\epsilon$. I will show that the local time converges to a Dirac mass at \hat{p} with measure bounded above by $\frac{1}{r}$ as $\epsilon \rightarrow 0$. By Proposition A.1, we have that

$$\begin{aligned}
L(\hat{p}; \hat{p}, \hat{p} - \epsilon, \hat{p} + \epsilon, r) &= \frac{2\sigma^2}{(1 - 2\gamma)\lambda\hat{p}(1 - \hat{p})} - \psi(\hat{p})(\hat{p} - \epsilon)^{1-\gamma}(1 - \hat{p} + \epsilon)^\gamma \hat{p}^{\gamma-2}(1 - \hat{p})^{-1-\gamma} \\
&\quad - \Psi(\hat{p})(\hat{p} + \epsilon)^\gamma(1 - \hat{p} - \epsilon)^{1-\gamma} \hat{p}^{-1-\gamma}(1 - \hat{p})^{\gamma-2} \\
&= \frac{2\sigma^2}{(1 - 2\gamma)\lambda\epsilon(1 - 2\epsilon)} - \psi(2\epsilon)\epsilon^{1-\gamma}(1 - \epsilon)^\gamma(2\epsilon)^{\gamma-2}(1 - 2\epsilon)^{-1-\gamma} \\
&\quad - \Psi(2\epsilon)(3\epsilon)^\gamma(1 - 3\epsilon)^{1-\gamma}(2\epsilon)^{-1-\gamma}(1 - 2\epsilon)^{\gamma-2} \\
&\geq \frac{2\sigma^2}{(1 - 2\gamma)\lambda\epsilon(1 - 2\epsilon)} - \epsilon^{1-\gamma}(1 - \epsilon)^\gamma(2\epsilon)^{\gamma-2}(1 - 2\epsilon)^{-1-\gamma} \\
&\quad - (3\epsilon)^\gamma(1 - 3\epsilon)^{1-\gamma}(2\epsilon)^{-1-\gamma}(1 - 2\epsilon)^{\gamma-2} \\
&= \frac{2\sigma^2}{(1 - 2\gamma)\lambda\epsilon(1 - 2\epsilon)} + A\epsilon^{-1}(1 - 2\epsilon)^{-1}, \quad \text{for some } A \in (0, \infty) \\
&\geq \frac{2\sigma^2}{(1 - 2\gamma)\lambda\epsilon(1 - 2\epsilon)} \\
&\rightarrow \infty \quad \text{as } \epsilon \rightarrow 0
\end{aligned}$$

where the inequality on the third line is valid since $|\psi(p)|, |\Psi(p)| \leq 1$. On the other hand, for any $a, b \in (0, 1)$, $\int_a^b L(p; \hat{p}, a, b, r) dp$ is clearly bounded above by $\frac{1}{r}$ - set $g(p) = 1$ in equation (A.1) and note

that $T = T(a) \wedge T(b) < \infty$. Hence

$$\begin{aligned}
V(\hat{p}) &= \int_{\underline{p}}^1 L(p; \hat{p}, \underline{p}, 1, \rho + \delta)(\lambda(p) - c) dp \\
&\geq \int_{\hat{p}-\epsilon}^{\hat{p}+\epsilon} L(p; \hat{p}, \underline{p}, 1, \rho + \delta)(\lambda(p) - c) dp \\
&\rightarrow \int_{\hat{p}-\epsilon}^{\hat{p}+\epsilon} \frac{\delta(\hat{p} - p)}{\rho + \delta} (\lambda(p) - c) dp \quad \text{as } \epsilon \rightarrow 0 \\
&= \frac{\bar{\lambda} - c}{\rho + \delta}
\end{aligned}$$

Since $p_0 > \hat{p}$ for sufficiently small ϵ , the result follows from the monotonicity of $V(p)$.

3. Follows by a similar argument to 2. □

In light of Definition 3, the uniqueness of η follows by the Intermediate Value Theorem:

Lemma C.7. *There exists a unique η such that $V(p_0; \eta) = K$*

D PROOF OF THEOREM 2

The following lemma will be used throughout the proof. Its proof is a direct application of the derivation in, for instance, Appendix A of Karatzas and Wang (2001) or more generally Section 2.6 in Karatzas and Shreve (1998).

Lemma D.1. *Let the stochastic process $(z_t)_{t \geq 0}$ solve the SDE:*

$$dz_t = \sigma(z_t) \sqrt{\lambda_t} dW_t, \tag{D.1}$$

on the interval $[0, 1]$, $\sigma(\cdot)$ satisfies conditions (1) and where $(W_t)_{t \geq 0}$ is a standard Wiener process. Let $\Lambda(z_0)$ denote the set of processes progressively measurable with respect to \mathbb{F}^z that take values in some bounded interval $[\underline{\lambda}, \bar{\lambda}] \subset \mathbb{R}_+$ such that equation (D.1) has a strong solution, and where the process $(z_t)_{t \geq 0}$

takes an initial value z_0 . Let the function f satisfy conditions [1](#), and

$$v(z_0) = \sup_{\tau, \lambda \in \Lambda(z_0)} \mathbb{E}^z \left[\int_0^{\tau} e^{-(\rho+\delta)t} f(z_t) dt \mid z(0) = z_0 \right] \quad (\text{D.2})$$

Finally, let the tuple $\{V(\cdot), \underline{z}, \lambda(\cdot)\}$ denote the variational problem:

$$\begin{aligned} \mathcal{A}_z\{V\} &= 0 \\ V(\underline{z}) &= 0 \\ V'(\underline{z}) &= 0 \\ V(z) &\geq 0 \quad \forall z \in [\underline{z}, 1] \\ V(z) &= 0 \quad \forall z \in [0, \underline{z}], \end{aligned} \quad (\text{D.3})$$

where

$$\mathcal{A}_z\{u\} = \max_{\lambda \in [\underline{\lambda}, \bar{\lambda}]} \left\{ \frac{1}{2} \sigma^2(z) \lambda u''(z) \right\} + f(z) - (\rho + \delta)u(z)$$

Suppose $\{V(\cdot), \underline{z}, \lambda(\cdot)\}$ solves the problem [\(D.3\)](#), and that with respect the process $(z_t)_{t \geq 0}$, the function $\lambda(z)$ is given by:

$$\lambda(z) = \begin{cases} \bar{\lambda} & \text{if } u''(z) \geq 0 \\ \underline{\lambda} & \text{if } u''(z) \leq 0 \end{cases} \quad (\text{D.4})$$

and that

$$\tau = \inf\{t \geq 0 : z_t \leq \underline{z}\} \quad (\text{D.5})$$

Then the function $V(\cdot)$ coincides with the value [\(D.2\)](#) and the pair $((\lambda_t)_{t \geq 0}, \tau)$ attain the supremum in [\(D.2\)](#).

Define the family of policies:

$$\mathcal{P} = \{\{\lambda\} \text{ non-empty} : \lambda(p) = \iota, \ p < p^*(\lambda), \ \lambda(p) = \epsilon, \ p \geq p^*(\lambda)\}$$

I claim that \mathcal{P} is non-empty and furthermore that $|\mathcal{E}(\lambda)| = 1$ for each $\lambda \in \mathcal{P}$. To see this, fix an $x \in [0, 1]$, define the policy $\{\lambda_x\}$ by $\lambda(p) = \iota, \ p < x, \ \lambda(p) = \epsilon, \ p \geq x$, and let $r(x) = p^*(\lambda_x)$. The mapping

$r \in \mathcal{C}^1([0, 1])$ with $r(0) > 0$ and $r(1) < 1$, all of which follow from an identical argument to that in Lemma C.2 part i). The claim follows by the Intermediate Value Theorem.

Now take an arbitrary non-empty policy $\{\lambda\}$. I claim that there exists a policy $\{\lambda_\iota\} \in \mathcal{P}$ such that $p^*(\lambda_\iota) \geq p^*(\lambda)$. Suppose not, that is $p^*(\lambda_\iota) < p^*(\lambda)$ for all $\{\lambda_\iota\} \in \mathcal{P}$. Under $\{\lambda\}$, the value function for a firm is given by:

$$v(p) = \sup_{\tau^x} \mathbb{E}^y \left[\int_0^{\tau^x} e^{-(\rho+\delta)t} (\bar{\lambda} \mathbb{I}_{p_t \geq p^*(\lambda)} - c) dt | p(0) = p \right] \quad (\text{D.6})$$

subject to the SDE:

$$dp_t = \frac{\sqrt{\lambda(p_t)}}{\sigma} p_t (1 - p_t) d\tilde{Z}_t,$$

and the initial condition $v(p_0) = K$. where Y_t is the cumulative review process for the policy $\{\lambda\}$. By Lemma 2, this value coincides with the function $V(p)$, where:

$$V_\lambda(p) = \frac{1}{\rho + \delta} \left[\bar{\lambda} \mathbb{I}_{p \geq p^*(\lambda)} - c + \lambda(p) \frac{p^2(1-p)^2}{2\sigma^2} V_\lambda''(p) \right], \quad V_\lambda(p(\lambda)) = V_\lambda'(p(\lambda)) = 0,$$

where $p(\lambda) \in (0, 1)$ such that $\tau_\lambda = \inf\{\tau \geq 0 | p_t = p(\lambda)\}$ solves (D.6). Let the pair $(\tilde{V}(p), \tilde{p})$ solve:

$$\tilde{V}(p) = \frac{1}{\rho + \delta} \left[\bar{\lambda} \mathbb{I}_{p \geq p^*(\lambda)} - c + \lambda_\iota(p) \frac{p^2(1-p)^2}{2\sigma^2} \tilde{V}''(p) \right], \quad \tilde{V}(\tilde{p}) = \tilde{V}'(\tilde{p}) = 0,$$

I claim that $\tilde{V}(p) \geq V_\lambda(p)$ for all $p \in [0, 1]$. In particular, I will show that $(\tilde{V}(\cdot), \tilde{p}, \lambda_\iota(\cdot))$ solves the variational problem (4). First, note that $\tilde{V}(p) \leq \bar{V} = \frac{\bar{\lambda}-c}{\rho+\delta}$, and hence for $p \geq p^*(\lambda)$,

$$\tilde{V}(p) - \bar{V} = \frac{\lambda_\iota(p)}{\rho + \delta} \frac{p^2(1-p)^2}{2\sigma^2} \tilde{V}''(p) \leq 0,$$

i.e. $\tilde{V}''(p) \leq 0$. Similarly, for $p < p^*(\lambda)$, $\tilde{V}''(p) > 0$. By definition of $\lambda_\iota(\cdot)$, it satisfies (D.4), and thus by Lemma D.1 the proof of the claim is complete. But then if $p^*(\lambda_\iota) \geq p^*(\lambda)$, it follows from equation (3) that $V_\iota(p) > \tilde{V} \geq V_\lambda(p)$, where $V_\iota(p)$ is the firm's equilibrium value under the policy λ_ι . Evaluating these inequalities at p_0 violates the condition $V_\iota(p_0) = K$.

E CALIBRATION

E.1 DATA


The data used for the figures in this section is taken from restaurants in Las Vegas. This city is the largest sample in the public data set provided by Yelp!. The price range is fixed at the “\$\$” range. Higher ranges have too few restaurants to provide reliability estimates, whereas lower ranges comprise largely of franchises and chains, the survival prospects of whom are typically uncorrelated with ratings. (see Luca (2016)) All restaurants are open as of June 2018. Overall, this sub-sample consists of 2,602 restaurants and 216,571 reviews over the sample period.

E.2 METHODOLOGY

The calibration proceeds as follows. First, I set the discount rate at an annual rate of 5%, taking dt as 1 year, in line with various macroeconomic studies. Data provided by the Bureau of Labor Statistics (BLS) on restaurant survival rates in the Western US allows me to externally calibrate the attrition rate δ . Using this data, Luo and Stark (2014) estimate a median lifespan of 4.5 years, dropping to 3.75 years when focusing on small start-ups with a labor force of 5 people or less. The remaining parameters are determined by matching the empirical ratings distribution to the theoretical distribution $f(p)$. The benefit of this approach is that it relies mainly on cross-sectional rather than time-series data. The dataset has been running for only 8 years. Furthermore, in the absence of precise markers for firm entry and exit times, inference based on dynamic hazard-rates is imprecise. Contrastingly, given the large number of cross-sectional observations (roughly 6 million reviews for 200,000 businesses), the empirical ratings distribution is robust. Two important transforms to the raw data must be applied before performing the calibration. First, one must take a stand on how reviews on Yelp! that have support $\{1, 2, 3, 4, 5\}$ map into the model’s normally distributed signal space. That is, consumer outcomes are not observed, and are captured by coarse feedback. I proceed with the assumption of symmetry - a review of 3 stars is equally likely to be generated by all firms, while a 1(2) star review is as informative as a 5(4) star review. Second, the raw data pertains to each firm’s cumulative review process X , not its rating p . I thus take the following approach: fix a guess for p_0 and σ , use the previous symmetry argument on reviews and


combine with the observed sample X to compute an estimate for a firm's rating p via Bayes' Rule. Along with the chosen value $\rho = 0.05$, the remaining parameters are then chosen in order that the model $f(p)$ fit the empirical distribution. The results of this procedure applied to the sample described in section E are reported in Table 1.

Figure 7: Calibration



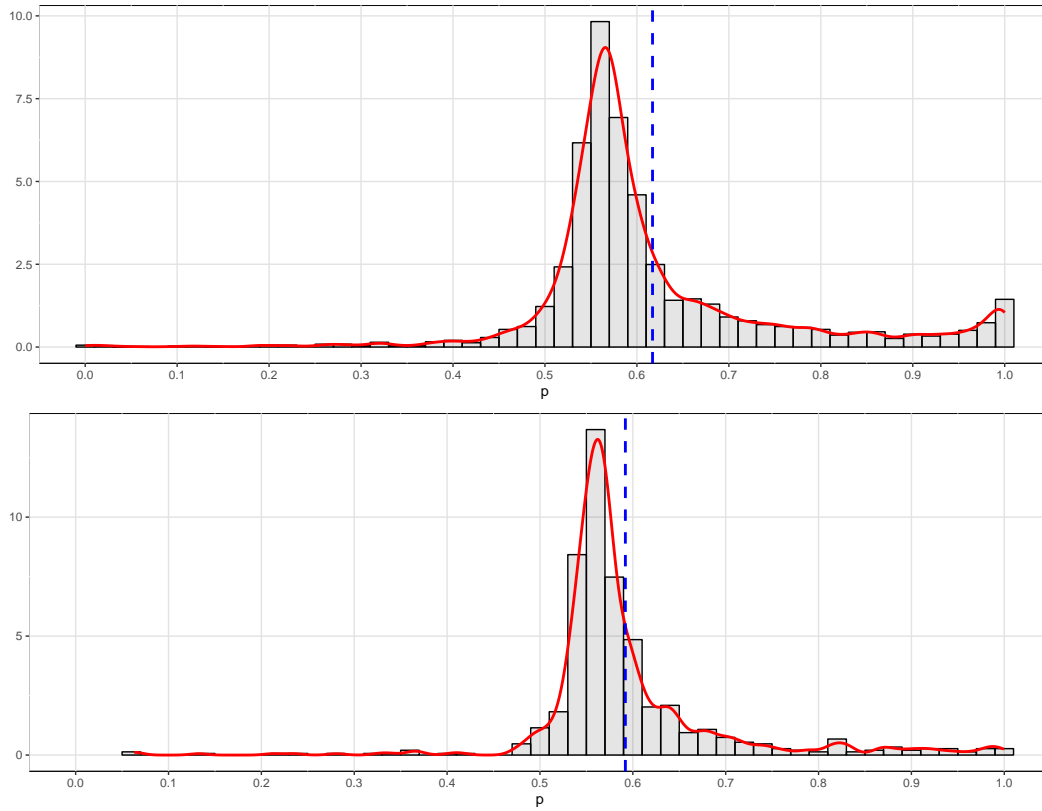
The shaded grey bars represent the empirical density histogram, with kernel density plot in red. The dotted blue line denotes the empirical mean. The black line plots the density from the calibrated model. The areas of worst fit are around the discontinuity at the theoretical value of p^* , as well as the left extreme of the support. Evidently, the notion that the increase in feedback rates for higher-rated firms is discontinuous is empirically invalid, and represents a key simplification of the model. The model is able to reproduce the fat right tail, as well as the right-skewness of the distribution.

Figure 8: Age distributions



These figures cross-section the ratings distribution by number of reviews, as a proxy for the age of the firms. The top panel uses firms with between 5 and 20 reviews, the middle panel between 20 and 200 reviews, and the bottom panel between 200 and 500 reviews. Together, they show that average quality is increasing with the age of the firm (selection), as well as that exit rates are hump-shaped in firm age: note that, for this data, the estimated exit rating \underline{p} is approximately 0.42, and that at this rating, the middle panel exhibits the largest density.

Figure 9: Selection effects



The top panel is the ratings distribution for firms that were active at the end of the sample. The bottom panel is firms that have closed down at some stage during the sample window. Inactive firms have on average a lower rating than active firms. That the bottom panel has non-trivial support indicates the empirical importance of the attrition parameter δ .