Insider Imitation[*]

Erik Madsen[†] Nikhil Vellodi[‡]

[Current Draft]: April 23, 2021

Abstract

Many platforms operating online marketplaces sell their own products alongside those of third-party sellers, and may use marketplace data on the demand for third-party products to inform the launch of their own products. To evaluate the anti-competitive implications of this practice, we model how a platform who can commit to a product introduction policy optimally uses marketplace data. An optimal policy trades off the ex post profitability of imitating successful third-party products against the ex ante reduction in innovation caused by imitation. We find that a regulation which bans the sharing of marketplace data stimulates innovation for "experimental" products with significant upside demand potential, but stifles it for "incremental" products with little upside potential.

Keywords: Innovation, online platforms, privacy, long-tail products.

1 Introduction

The internet has become an important channel for the sale of consumer products. In the United States, 14% of all retail sales took place online in 2020.[1] A large fraction of these online sales are transacted through *marketplaces*, websites that allow consumers to find and

[*]We would like to thank Panle Barwick, Olivier Compte, Phillipe Jehiel, Alessandro Lizzeri, Leslie Marx and seminar participants at PSE and Cornell for insightful discussions. Vellodi acknowledges support from the EUR grant ANR-17-EURE-0001.

[†]Department of Economics, New York University.

[‡]Paris School of Economics.

[1]Current US Census Bureau retail sales figures are available at https://tinyurl.com/cfzuaynb. In several key product categories, online purchases are even more important, accounting for over 60% of sales of books, music, and video; half of all sales of consumer electronics and toys; and over a third of apparel sales (Coppola 2020).

purchase products from third-party sellers.[2] While some marketplaces exclusively sell third-party products, others additionally offer house-brand products marketed by the owner of the marketplace. For instance, Amazon markets a variety of products under its "AmazonBasics and "Amazon Essentials" brands, among others, while Walmart sells house brands under labels such as "Sam's Choice" and "Great Value".

The sale of house brands alongside third-party products on online marketplaces has attracted the scrutiny of lawmakers and antitrust regulators, who have expressed concern that marketplace owners may be using house brands to compete unfairly with third-party sellers. One practice of interest, and the focus of this paper, is the use of marketplace sales data to inform decisions about the introduction of house-brand products. News reports indicate that Amazon in particular has leveraged its marketplace data extensively in this manner.[3] The European Commission recently issued a preliminary view that Amazon's practices breach EU antitrust law by allowing the company to unfairly "avoid the normal risks of retail competition and to leverage its dominance in the market for the provision of marketplace services" (European Commission 2020). A United States Congressional investigation of competition in digital marketplaces similarly concluded that "Amazon's dual role... incentivizes Amazon to exploit its access to competing sellers' data and information", and expressed the view that this behavior constitutes anticompetitive conduct (U.S. House Committee on the Judiciary 2020, p.16).

In this paper, we model a marketplace's usage of seller data to inform product introduction, and examine the implications of the practice for competition, innovation, and regulation. Our analysis proceeds from two main assumptions about competition on online marketplaces. First, we suppose that online marketplaces provide outlets for the sale of products which would not otherwise exist. This assumption is in line with evidence on the existence of a "long tail" of products which are sold online but not in any physical store (Anderson 2006; Brynjolfsson, Hu, and Smith 2003, 2006). It is also consistent with the

[2]In 2020, an estimated 39% of all online retail sales in the United States took place on Amazon, the dominant online marketplace, while the second- and third-largest online retailers Walmart and eBay – both of which also run online marketplaces – each accounted for approximately 5% of online sales (Lipsman 2020).

[3]This practice has been independently documented by Bloomberg (Soper 2016), Buzzfeed (Miranda 2018), and the Wall Street Journal (Mattioli 2020). It is also discussed in detail in U.S. House Committee on the Judiciary (2020, pp.274-282).

much lower barriers to entry needed to sell in a large online marketplace as compared to a physical retailer of similar scale. In particular, online marketplaces often allow sellers to list products for free, require little or no vetting, and handle payments and even order fulfillment. By contrast, large physical retailers extensively vet suppliers, require proof of ability to reliably fulfill minimum volumes, and prefer suppliers with proven products who maintain a diverse set of sales channels.[4] The reduced scale requirements of online marketplaces are reflected in the demographics of marketplace sellers, many of whom are small companies comprised of a handful of employees who make most or all of their revenue by selling on a single marketplace, and who often make only a few thousand dollars in monthly revenue.[5]

Second, we assume that marketplace owners internalize the impact of aggressive competition by house brands on product diversity in the marketplace. Due to the small scale and lack of brand recognition of many marketplace sellers, such competition has the potential to significantly impact innovation, particularly when targeted at successful products using marketplace data. We take seriously the likelihood that vertically integrated marketplaces recognize this danger and make decisions which maximize joint long-run profits across divisions. We therefore focus on the role of regulation in curbing behavior which advantages vertically integrated marketplaces at the expense of social or consumer welfare. We abstract from the question of whether marketplace owners might fail to coordinate their policies across divisions, and whether regulation might be helpful for enforcing such coordination.

In our model, a vertically integrated online platform hosts a marketplace on which a third-party entrepreneur can sell her products, and on which the platform can additionally sell its own products. The entrepreneur is uniquely able to *innovate* by developing a new product not already sold on the marketplace. The platform can only *imitate* by developing a derivative product to compete with one already sold. Innovation is an inherently risky

[4]Small business consultants advise that to be approved as a Walmart supplier, a seller should rely on Walmart for no more than 30% of their sales and should first establish a track record by selling their product in smaller retailers or online marketplaces (Bose 2018; "How to Sell to Walmart" 2020).

[5]A recent survey of Amazon sellers found that over 50% of sellers had five or fewer employees, and roughly 50% of sellers made over 80% of their revenue selling on Amazon. 13% of sellers reported selling *nowhere* other than Amazon ("The State of the Amazon Marketplace 2019" 2020). Another survey found that only half of Amazon sellers exceed $5,000 in monthly sales ("The State of the Amazon Seller 2021" 2021).

activity, as demand for a new product on the marketplace is uncertain and can be discovered only by incurring a fix cost to develop and market the product. By contrast, the profitability of an imitation product can be predicted from the demand for the baseline product. The platform earns profits from a combination of sales of its own product and transaction fees for sales of the entrepreneur's product.

Prior to the entrepreneur's innovation decision, the platform commits to a *competition policy*, which specifies how aggressively the platform imitates existing products. Absent regulation, the platform can observe the sales of existing products and condition its imitation decisions on that data. A competition policy therefore dictates both *which* products are imitated, as a function of the demand for the product, as well as *how quickly* a competing product is introduced.

When designing a competition policy, the platform balances its ex ante desire to attract innovative products, which generate revenue for the platform via fees, against the ex post profit gains from imitating successful products. In this sense, the platform solves a problem similar to that of a planner designing a patent policy providing incentives for innovation, as in, for instance, Hopenhayn and Squintani (2016) and Henry and Ponce (2011). We show that under an optimal competition policy, sufficiently successful products are imitated immediately, while unsuccessful products are never imitated. Such a policy generally does not maximize social welfare, but does not inevitably lead to too little innovation. We show that from a planner's point of view the platform may imitate either too much or too little, depending on how much the planner values competition.[6]

We then study the impact of a regulation which blocks *insider imitation*, the ability of the platform to use marketplace data to tailor their competition policy. Our main result connects the impact of such a regulation to uncertainty about the demand for a new product. We define a novel notion of *upside demand potential*, measuring the fraction of total demand which is concentrated in "right-tail" states of the world with very high demand. We show that for "experimental" products with significant upside demand potential, banning insider imitation stimulates innovation. Conversely, for "incremental" products, which have little

[6]The platform may imitate too little relative to the first-best because of an uninternalized consumer surplus gain from competition. This externality is reminiscent of the spillovers effects highlighted in the dynamic entry models of Rob (1991) and Jovanovic and Lach (1989), where firm entry generates positive spillover effects for existing firms and entry rates may be inefficiently low.

upside potential and tend to exhibit demand concentrated near the average, banning insider imitation stifles innovation. Whether a ban on marketplace data usage enhances social surplus therefore hinges on two features of the product market: The shape of the distribution of demand for new products, and the welfare gains from competition.

Our work is complementary to a pair of recent papers by Hagiu, Teh, and Wright (2020) and Etro (2020) which examine regulation of vertically integrated online platforms. Hagiu, Teh, and Wright (2020) suppose that the platform cannot commit to avoid imitating third-party sellers. They find that a regulation which bans imitation of third-party products may increase not only social surplus but also the platform's equilibrium profits, as the regulation makes up for a lack of commitment power by the platform. Etro (2020) allows the platform to commit to a competition policy, and finds that a ban on imitation decreases long-run consumer surplus. He also considers the impact of a data regulation which prevents the platform from targeting a house brand product at the customers of competing products, and finds that such a regulation also decreases consumer surplus. Neither paper models demand uncertainty or considers regulations which restrict the ability of the platform to use demand data to target product introduction. To the best of our knowledge, our paper is the first to provide a formal economic analysis of a setting with those features.[7]

Our work also connects to a rapidly growing literature studying the implications of increased consumer data collection for market outcomes (Acemoglu et al. 2021; Bergemann, Bonatti, and Gan 2021; Liang and Madsen 2021), as well as to a closely related literature examining the role of consumer privacy in online markets (Eilat, Eliaz, and Mu 2020; Fainmesser, Galeotti, and Momot 2021; Gomes and Pavan 2021; Hidir and Vellodi 2020; Ichihashi 2019, 2020). These papers highlight the potential value of consumer privacy regulations which restrict the use of consumer data by sellers. The insider imitation ban we evaluate can be viewed as an analogous data privacy regulation aimed at protecting particular sellers. Our work highlights that increased data usage has important implications for producer as well as consumer privacy, expanding the scope of applications in which regulation may be warranted.

The remainder of the paper is structured as follows. In Section 2, we introduce the

[7] The antitrust implications of demand data usage by platforms have also been discussed by legal scholars. See, e.g., Khan (2017), §IV.D.

baseline model, which we analyze in a benchmark setting without uncertainty in Section 3, and with demand uncertainty under a linear demand specification in Section 4. In Section 5, we characterize the impact of an insider imitation ban on innovation and study the welfare implications of the regulation. Finally, in Section 6, we extend our model to include transfers and nonlinear demand specifications.

2 Model

2.1 The environment

An online platform and an entrepreneur interact on a marketplace run by the platform. The interaction unfolds in three stages. In the first stage, the platform publicly commits to a competition policy, which we discuss further in Section 2.2. In the second stage, the entrepreneur chooses whether to *innovate* by developing a new product. If she chooses not to innovate, the game ends and both players make zero profits. Otherwise, the entrepreneur pays a fixed cost $k > 0$ to develop the product. In the final stage, the entrepreneur sells her product on the marketplace, generating a stream of profits over an infinite, continuous time horizon. At any time, the platform can choose to *imitate* the entrepreneur's profit, paying a fixed cost $k_P > 0$ to instantaneously develop and market a competing product.

We model the profits generated by selling products on the marketplace in a flexible reduced-form way. So long as the entrepreneur's product is the only one present in the market, she earns a flow of profits $\Pi_E^M(\alpha) > 0$, where α is a *demand state*, to be described in further detail shortly. Meanwhile the platform makes a flow of profits $\Pi_P^M(\alpha) > 0$ which reflect fees it collects from the entrepreneur's sales. We will refer to $\Pi_E^M(\alpha)$ and $\Pi_P^M(\alpha)$ as each player's *monopoly profits*, though they may embed competition from other products on the marketplace.

Once the platform has introduced an imitating product, the entrepreneur's profits fall to $\Pi_E^D(\alpha) \in (0, \Pi_E^M(\alpha))$, while the platform's profits rise to $\Pi_P^D(\alpha) > \Pi_M^D(\alpha)$, with $\Pi_P^D(\alpha)$ capturing a combination of the platform's profits from sales of its own product as well as fees from sales of the entrepreneur's product. These profit specifications may be interpreted as a stand-in for the outcome of a static competition between the platform and entrepreneur, for instance via quantity- or price-setting. We will refer to $\Pi_E^D(\alpha)$ and $\Pi_P^D(\alpha)$ as the *duopoly profits* of each player.

Both players are expected profit maximizers who share a common discount rate r. As a result, if the entrepreneur chooses to innovate and the platform chooses to imitate at time T, the entrepreneur's lifetime profits are

$$U_E = -k + (1 - e^{-rT})r^{-1}\Pi_E^M(\alpha) + e^{-rT}r^{-1}\Pi_E^D(\alpha),$$

while the platform's lifetime profits are

$$U_P = (1 - e^{-rT})r^{-1}\Pi_P^M(\alpha) + e^{-rT}(r^{-1}\Pi_E^D(\alpha) - k_P).$$

Two important forms of uncertainty arise in the model. First, while the entrepreneur observes the innovation cost k prior to deciding whether to innovate, from the platform's perspective k is unobserved and uncertain. At the time it formulates a competition policy, the platform believes that $k \sim G$, where G is continuous, has support on the bounded interval $[\underline{k}, \overline{k}]$ with $0 < \underline{k} < \overline{k}$, and is continuously differentiable on its support. We additionally impose a standard monotone hazard rate condition on the distribution of costs:

Assumption 1. *The cost hazard rate $G'(k)/G(k)$ is weakly decreasing on $[\underline{k}, \overline{k}]$, and is zero at $k = \overline{k}$.*

Second, both players view the demand for a new product as uncertain prior to the product's introduction to the market. They share a common prior that $\alpha \sim F$, where $\alpha \in \mathbb{R}_+$ is normalized so that F has mean 1. We do not require that F be continuous, allowing for the possibility of atoms in the demand distribution. To streamline exposition, we will maintain the assumption that F has support on a (possibly unbounded) interval.[8] As soon as a product is introduced, all demand uncertainty is resolved and α is observed by both players.

We interpret higher demand states as corresponding to a more profitable market for

[8] Our main results continue to hold without the interval support assumption. Assuming interval support ensures continuous differentiability of the platform's objective function and continuity of the threshold demand function $\alpha^*(\kappa)$, simplifying the statement of results.

each player regardless of the amount of competition, so that each $\Pi_i^s(\cdot)$ is strictly increasing in α for each player $i = P, E$ and market structure $s = M, D$. We will also assume that consumer surplus under each market structure, which we will denote $CS^s(\alpha)$, is increasing in the demand state. To streamline exposition, we will additionally assume that each $\Pi_i^s(\cdot)$ and $CS^s(\cdot)$ is continuous in α. An immediate consequence of these assumptions is that total social welfare, which we will denote $W^s(\alpha) \equiv \Pi_P^s(\alpha) + \Pi_E^s(\alpha) + CS^s(\alpha)$, is continuous and increasing in demand.

Our main results will focus on a linear demand model in which $\Pi_i^s(\alpha) = \alpha\mu_i^s$, $CS_i^s(\alpha) = \alpha\nu^s$, and $W^s(\alpha) = \alpha w^s$ for each s and i. Given that $\mathbb{E}[\alpha] = 1$, the μ_i^s represent average profits for each player under a given market structure, while ν^s represents average consumer surplus and w^s represents average social welfare. This linear specification can be microfounded by an environment in which α represents uncertainty about the size of the market, while the distribution of values among consumers is known, and in which both players have linear product costs.

2.2 Competition policies

Prior to the entrepreneur's innovation decision, the platform publicly commits to an *competition policy* dictating the time at which it develops an imitation product, supposing that the entrepreneur has innovated. This policy cannot condition on the entrepreneur's innovation cost, which is the private information of the entrepreneur.[9] However, it can condition on the market demand state α, which is observed by the platform as soon as the entrepreneur's product is introduced.

Formally, a competition policy is a function $T : \mathbb{R}_+ \to \mathbb{R}_+ \cup \{\infty\}$, where $T(\alpha)$ denotes the time at which the platform imitates the entrepreneur's product when the demand state is α. We allow the imitation time to be infinite to accommodate the possibility that the platform never imitates some products.

[9] Because the cost of innovation does not impact the entrepreneur's willingness to be imitated, the platform would not benefit from offering a menu of competition policies. This remains true even if the platform can additionally offer transfers, as in Section 6.2.

2.3 Notation and assumptions

We will write $\overline{\Pi}_i^s \equiv \mathbb{E}[\Pi_i^s(\alpha)]$ to denote the expected flow profits for a given player under each market structure, averaging over all possible demand states. We will also write $\Delta\Pi_i^s(\alpha) \equiv \Pi_i^D(\alpha) - \Pi_i^M(\alpha)$ to denote the change in each player's flow profits when moving from the monopoly to duopoly market structure in demand state α, and $\Delta\overline{\Pi}_i \equiv \mathbb{E}[\Delta\Pi_i^s(\alpha)]$ to denote a player's average change in flow profits. $\overline{CS}^s$, $\Delta CS(\alpha)$, and $\Delta\overline{CS}$ are defined similarly for consumer surplus, as are $\overline{W}^s$, $\Delta W(\alpha)$, and $\Delta\overline{W}$ for social welfare. We will assume that each $\overline{\Pi}_i^s$ and $\overline{CS}^s$ is finite. In the linear demand model, we will use the alternative notation $\Delta\mu_i \equiv \mu_i^D - \mu_i^M$, $\Delta\nu \equiv \nu^D - \mu^M$, and $\Delta w \equiv w^D - w^M$ to denote changes in average profits, consumer surplus, and social welfare.

We make several assumptions on the size of profit changes between the monopoly and duopoly states to focus our analysis on settings with nontrivial regulatory concerns surrounding marketplace data usage. We will assume that the entrepreneur always enters if assured a perpetual monopoly, no matter the cost of innovation, but that an entrepreneur with sufficiently high innovation costs would not enter if imitated immediately:

Assumption 2. $\overline{\Pi}_E^M > r\bar{k} > \overline{\Pi}_E^D$.

The first part of the assumption is without loss, as entrepreneurs who would not enter even when assured a monopoly are irrelevant to the platform and regulator. The second part of the assumption ensures that imitation has the potential to reduce innovation. If it did not hold, there would be no regulatory concerns around marketplace data usage.

We will also assume that imitation is more profitable for the platform in high demand states than low ones, and that the platform finds imitation profitable on average, but unprofitable in very low demand states:

Assumption 3. *$\Delta\Pi_P(\cdot) - rk_P$ satisfies single-crossing in α, and $\Delta\overline{\Pi}_P > rk_P > \Delta\Pi_P(0)$.*

Single-crossing is a regularity condition which is satisfied, for instance, in the linear demand model, since in that case $\Delta\overline{\Pi}_P = \Delta\mu_P > 0$ implies that $\Delta\Pi_P(\alpha) - rk_P = \alpha\Delta\mu_P rk_P$ is increasing in α. The condition $\Delta\overline{\Pi}_P > rk_P$ ensures that, under a regulation banning usage of marketplace data, the platform still finds imitation potentially profitable. Meanwhile the condition $rk_P > \Delta\Pi_P(0)$ rules out the possibility that the platform optimally imitates immediately in all demand states, in which case marketplace data would not be useful to the platform and regulating it would be ineffective.[10]

An immediate corollary of Assumption 3 is $\Delta\Pi_P(\alpha) > rk_P$ for α sufficiently large. It will be useful to develop notation for the lowest demand state at which imitation yields non-negative profits. We will write α_0 for this state, which given our continuity assumptions satisfies $\Delta\Pi_P(\alpha_0) = rk_P$ and $\alpha_0 > 0$. We will let k_0 be the corresponding innovation cost above which an entrepreneur would not innovate, supposing its product is imitated immediately if $\alpha \geq \alpha_0$, but is never imitated otherwise. That is,

$$k_0 \equiv \max\left\{\underline{k}, r^{-1}\left(\overline{\Pi}_E^M + \int_{\{\alpha \geq \alpha_0\}} dF(\alpha)\, \Delta\Pi_E(\alpha)\right)\right\}.$$

3 The No-Uncertainty Benchmark

We begin by deriving the platform's optimal competition policy in a benchmark setting without demand uncertainty, where profits are equal to their average levels with probability 1. This analysis serves two purposes. First, it illuminates important aspects of the platform's policy design problem, which will carry over into the full problem with uncertainty. Second, an optimal policy absent uncertainty corresponds precisely to the policy the platform would choose under a regulation which prevents it from conditioning imitation on the market demand state.

Absent demand uncertainty, the platform's competition policy is a single imitation time $T \in \mathbb{R}_+ \cup \{\infty\}$. Let $\widetilde{k}_0$ be the highest cost an entrepreneur would pay to innovate supposing she were imitated immediately. That is, $\widetilde{k}_0 \equiv \max\{\underline{k}, r^{-1}\overline{\Pi}_E^D\}$. To streamline our analysis, we impose a lower bound on the cost hazard rate at $\widetilde{k}_0$ which ensures an interior solution to the platform's policy design problem:

Assumption 4. $\frac{G'(\widetilde{k}_0)}{G(\widetilde{k}_0)} \geq r\frac{\Delta\overline{\Pi}_P - rk_P}{|\Delta\overline{\Pi}_E|\left(\overline{\Pi}_P^D - rk_P\right)}$.

This assumption is satisfied if $\widetilde{k}_0 = \underline{k}$, or if $G(\widetilde{k}_0)$ is sufficiently small.

[10] Nothing in our analysis would change if $\Delta\Pi_P(0) \geq rk_P$, so long as the mass of entrepreneurs who would enter supposing they were imitated immediately in all demand states is sufficiently low that such a policy is not optimal for the platform.

Given a competition policy T, the entrepreneur's total profits from innovating are

$$U_E(T,k) \equiv \left(1-e^{-rT}\right) r^{-1}\overline{\Pi}_E^M + e^{-rT} r^{-1}\overline{\Pi}_E^D - k.$$

Since $\overline{\Pi}_E^M > \overline{\Pi}_E^D$, $U_E(T,k)$ is strictly increasing in T and strictly decreasing in k. As a result, for each competition policy there exists a corresponding cost threshold above which innovation is suppressed, and this threshold is decreasing in T. By inverting this relationship, the platform's competition policy design problem can be equivalently cast as targeting a threshold innovation cost $\kappa \in [\underline{k}, \overline{k}]$ by choosing the unique imitation time $\widetilde{T}(\kappa)$ satisfying $U_E(\widetilde{T}(\kappa), \kappa) = 0$.

Not all cost thresholds in $[\underline{k}, \overline{k}]$ are guaranteed to be implementable by a competition policy. In particular, cost thresholds below $\widetilde{k}_0$ are not implementable if $r^{-1}\overline{\Pi}_E^D > \underline{k}$, as entrepreneurs with costs below this level are undeterred even by very aggressive imitation. On the other hand, Assumption 2 ensures that all thresholds up to $\overline{k}$ are implementable. Therefore $\widetilde{T}(\kappa)$ is well-defined, finite, and strictly increasing on $[\widetilde{k}_0, \overline{k}]$.

Given a threshold innovation cost $\kappa \in [\widetilde{k}_0, \overline{k}]$, the platform's total profits are

$$\begin{aligned}
\widetilde{U}_P(\kappa) &= G(\kappa) r^{-1}\left(\overline{\Pi}_P^M + (\Delta\overline{\Pi}_P - rk_P) e^{-r\widetilde{T}(\kappa)}\right) \\
&= G(\kappa) r^{-1}\left(\overline{\Pi}_P^M + (\Delta\overline{\Pi}_P - rk_P)\frac{\overline{\Pi}_E^M - r\kappa}{|\Delta\overline{\Pi}_E|}\right).
\end{aligned}$$

The following lemma establishes that the platform's optimal innovation cost threshold is interior and is the unique critical point of $\widetilde{U}_P(\kappa)$:

Lemma 1. *In the absence of demand uncertainty, the optimal innovation threshold κ^{**} is unique, satisfies $\kappa^{**} \in (\widetilde{k}_0, \overline{k})$, and is characterized by the first-order condition*

$$\frac{G'(\kappa^{**})}{G(\kappa^{**})}\left(\overline{\Pi}_P^M + (\overline{\Pi}_E^M - r\kappa^{**})\bar{\Lambda}\right) = r\bar{\Lambda},$$

where

$$\bar{\Lambda} \equiv \frac{\Delta\overline{\Pi}_P - rk_P}{|\Delta\overline{\Pi}_E|}.$$

Several of our assumptions play a key role in this result. Assumption 1, which imposes a monotone cost hazard rate, ensures that $\widetilde{U}_P$ has at most one critical point, which if it exists is a global optimum. It also imposes a zero hazard rate at $\overline{k}$, implying that $\widetilde{U}_P(\kappa)$ is downward-sloping there. Meanwhile Assumption 4 ensures that $\widetilde{U}_P(\kappa)$ is upward-sloping at $\widetilde{k}_0$, so that an interior critical point must exist.

3.1 The Value of Imitation

As a prelude to our goal of comparing optimal rates of innovation across the regulated and unregulated problems, we extract some important general insights from the analysis without uncertainty.

We begin by slightly rearranging the first-order condition derived in Lemma 1 to read

$$G'(\kappa)\left(\overline{\Pi}_P^M + (\overline{\Pi}_E^M - r\kappa)\bar{\Lambda}\right) = G(\kappa)r\bar{\Lambda}.$$

The platform's first-order condition written in this form represents a balancing of marginal gains from innovation (the left-hand side) against inframarginal losses (the right-hand side) when the platform raises κ. Such an increase induces the entry of a mass of entrepreneurs of size $G'(\kappa)\,d\kappa$, from whom the platform collects the baseline monopoly profit $\overline{\Pi}_P^M$ plus an additional surplus from imitation. This additional surplus is equal to the marginal entrepreneur's entire profits converted into platform profits at rate $\bar{\Lambda}$. This quantity $\bar{\Lambda}$ therefore represents the platform's *average conversion rate*, and measures the marginal rate of transformation of the entrepreneur's profits into the platform's profits achievable by increased imitation, assuming imitation is uniform across all demand states.

To achieve this additional entry, the platform must give up profits from all existing innovators, constituting a mass of size $G(\kappa)$, by delaying imitation. A given rise in κ directly translates into an equivalent rise in the lifetime profits of every innovator, reducing platform flow profits by an amount $r\bar{\Lambda}\,d\kappa$. This trade-off between marginal gains and infra-marginal losses is analogous to the classical monopoly pricing problem, wherein the monopolist balances the inframarginal revenue gain from existing customers against the marginal loss of customers when deciding how far to raise prices.

The average conversion rate $\bar{\Lambda}$ measures how efficiently the platform can convert the entrepreneur's profits into its own. As such, it is likely a crucial determinant of the platform's incentive to imitate. The following lemma establishes that as the conversion rate improves, the platform imitates more aggressively.

Lemma 2. *Suppose* $\Lambda > 0$. *Let*

$$\widehat{U}_P(\kappa, \Lambda) \equiv G(\kappa) r^{-1} \left(\overline{\Pi}_P^M + \Lambda \left(\overline{\Pi}_E^M - r\kappa \right) \right)$$

and define $\widetilde{\kappa}(\Lambda)$ *to be the unique maximizer of* $\widehat{U}_P(\cdot, \Lambda)$ *on* $[\widetilde{k}_0, \overline{k}]$. *Then* $\widetilde{\kappa}(\Lambda)$ *is decreasing in* Λ, *and strictly so whenever* $\widetilde{\kappa}(\Lambda) > \widetilde{k}_0$.

An appealing but incomplete explanation of this result is that an improved conversion rate makes imitation more attractive due to the increased platform profits per unit of profits extracted from inframarginal entrepreneurs. But this argument neglects the fact that an increased conversion rate also increases the losses from excluding marginal entrepreneurs, as those innovators were more profitable for the platform under a higher conversion rate. As a result, the net effect on innovation of a change in Λ is not immediately obvious.

The crucial insight needed to establish Lemma 2 is that the marginal gains from innovation are determined not only by the profits from imitation, but also by the platform's baseline profits under monopoly. These gains are therefore less sensitive to the conversion rate than are the inframarginal losses of innovation, which are solely determined by changes in profits from imitation. As a result, if marginal gains balanced inframarginal losses at a given Λ, raising the conversion rate raises marginal gains by less than it raises inframarginal losses. The amount of innovation must then fall to optimally re-balance these levels. A corollary of this reasoning is that if $\overline{\Pi}_M^P = 0$, i.e. the platform makes no profits merely by hosting the entrepreneur's product, then the conversion rate would have no impact on the optimal amount of innovation.

4 The Laissez-Faire Outcome

We now derive the form of an optimal competition policy absent regulation under the linear demand model. (We generalize these results to nonlinear specifications in Section 6.1.)

Recall that a competition policy is a function $T : \mathbb{R}_+ \to \mathbb{R}_+ \cup \{\infty\}$, specifying an imitation time for each possible demand state α. As in the no-uncertainty benchmark of Section 3, the problem of designing an optimal competition policy can be reduced to optimally targeting a threshold innovation cost $\kappa \in [\underline{k}, \overline{k}]$ to innovate. Unlike in that benchmark, however, there are now many ways the platform could induce a particular threshold, as the competition policy has many degrees of freedom.

A first step in the analysis is therefore to characterize the optimal competition policy given a desired level of innovation. That is, the platform must design a policy T to maximize

$$G(\kappa)r^{-1}\left(\mu_P^M + \int dF(\alpha)e^{-rT(\alpha)}(\alpha\Delta\mu_P - rk_P)\right),$$

subject to the marginal entrepreneur's zero-profit constraint

$$\mu_E^M + \Delta\mu_E \int dF(\alpha)\, e^{-rT(\alpha)}\alpha = r\kappa.$$

As usual, only cost thresholds satisfying $\kappa \geq \widetilde{k}_0$ can be implemented by some competition policy. However, no policy which imitates in states $\alpha < \alpha_0$, where α_0 is as defined in Section 2.3 and satisfies $\alpha_0\Delta\mu_P = rk_P$, can ever be optimal for the platform in the global problem in which the platform optimizes over κ. As a result, we may assume that $T(\alpha) = \infty$ for $\alpha < \alpha_0$, in which case only cost thresholds satisfying the stricter lower bound $\kappa \geq k_0$ are achievable, where $k_0 \geq \widetilde{k}_0$ is as defined in Section 2.3. Going forward, we will restrict attention to cost thresholds satisfying this latter lower bound.

The following lemma establishes that an optimal competition policy inducing a threshold $\kappa \in [k_0, \overline{k}]$ takes a simple threshold form: The platform imitates immediately in demand states above a marginal state $\alpha^*(\kappa^*)$, and avoids imitating in all other state.

Lemma 3. *The optimal competition policy inducing an innovation cost threshold $\kappa \in [k_0, \overline{k}]$*

is uniquely defined up to a measure-zero set of demand states. One optimal policy satisfies

$$T^*(\alpha,\kappa) = \begin{cases} 0, & \alpha > \alpha^*(\kappa), \\ \overline{T}(\kappa), & \alpha = \alpha^*(\kappa) \\ \infty, & \alpha < \alpha^*(\kappa), \end{cases}$$

where

$$\alpha^*(\kappa) \equiv \sup\left\{\alpha \ : \ \mu_E^M + \Delta\mu_E \int_{\{\alpha' \geq \alpha\}} dF(\alpha')\,\alpha' \leq r\kappa\right\}$$

and $\overline{T}(\kappa) \in \mathbb{R}_+ \cup \{\infty\}$ *is the smallest solution to*

$$\mu_E^M + \Delta\mu_E \left(\int_{\{\alpha > \alpha^*(\kappa)\}} dF(\alpha)\alpha + \alpha^*(\kappa)\Delta F(\alpha^*(\kappa))e^{-rT}\right) = r\kappa.$$

The threshold $\alpha^*(\kappa)$ *satisfies* $\alpha^*(\kappa) \in [\alpha_0, \infty)$ *and is continuous and increasing in* κ*.*

Optimal imitation exhibits a threshold structure with respect to the demand state because, under a linear demand specification, the *marginal conversion rate*

$$\Lambda(\alpha) \equiv \frac{\Delta\Pi_P(\alpha) - rk_P}{|\Delta\Pi_E(\alpha)|}$$

is larger in better demand states. The marginal conversion rate is closely analogous to the average conversion rate $\bar{\Lambda}$ defined in Section 3.1, but conditions on imitation in a particular demand state rather than assuming uniform imitation across all states. Under linear demand $\Lambda(\alpha)$ is strictly increasing in α, and so the most profitable way for the platform to extract a given amount of surplus from entrepreneurs is to imitate in high demand states.

Whenever the demand distribution is continuous at the threshold demand state $\alpha^*(\kappa)$, the threshold state arises with zero probability and the choice of imitation time in this state is immaterial. In this case, we follow a convention which selects an optimal policy with $\overline{T}(\kappa) = 0$. By contrast, whenever the threshold demand state arises with positive probability, due to the presence of an atom in the demand distribution, the choice of $\overline{T}(\kappa)$ affects the average profits of both players. In that case, $\overline{T}(\kappa)$ is uniquely pinned down by the entrepreneur's zero-profit condition, and may be both nonzero and finite.[11]

Given a threshold innovation cost $\kappa \in [k_0, \overline{k}]$, the platform's total profits are

$$U_P(\kappa) = G(\kappa) r^{-1} \left(\mu_P^M + \int dF(\alpha)\, e^{-rT^*(\alpha,\kappa)} \left(\alpha \Delta\mu_P - r k_P \right) \right).$$

The following lemma establishes that the platform's optimal innovation cost threshold is interior and is the unique critical point of $U_P(\kappa)$. It also writes the corresponding FOC in a form which will facilitate analysis of the impact of a data regulation.

Lemma 4. *The optimal innovation threshold κ^* is unique, satisfies $\kappa^* \in (k_0, \overline{k})$, and is characterized by the first-order condition*

$$\frac{G'(\kappa^*)}{G(\kappa^*)} \left(\mu_P^M + (\mu_E^M - r\kappa^*) \Lambda(\alpha^*(\kappa^*)) + R(\alpha^*(\kappa^*)) \right) = r \Lambda(\alpha^*(\kappa^*)),$$

where

$$R(\alpha) \equiv r k_P \int_{\{\alpha' > \alpha\}} dF(\alpha') \left(\frac{\alpha'}{\alpha} - 1 \right).$$

This first-order condition is very similar to the one derived in Lemma 1 for the no-uncertainty case, with two important distinctions. First, the average conversion rate $\bar{\Lambda}$ is replaced by the marginal conversion rate $\Lambda(\alpha^*(\kappa))$, reflecting the fact that under demand uncertainty, a marginal increase in imitation occurs solely in states near $\alpha^*(\kappa)$ and thus converts the entrepreneur's profits into platform profits at the rate $\Lambda(\alpha^*(\kappa))$. Second, the left-hand size of the FOC, capturing the profit to the platform from admitting an additional entrepreneur, contains an additional non-negative remainder term $R(\alpha^*(\kappa))$. This term accounts for the fact that the marginal conversion rate is not constant as κ changes, and is larger for higher κ. Thus the conditional average conversion rate, governing the overall profits from an entering entrepreneur, is larger than the marginal conversion rate. The term $R(\alpha^*(\kappa))$ corrects for this discrepancy. In general this remainder term pushes the platform toward a larger threshold innovation cost, i.e. more innovation and less imitation, than would be suggested by the marginal conversion rate alone.

[11] If we allowed the platform to commit to a randomized competition policy, then delayed imitation could be replicated by randomization over extremal imitation.

5 Regulating Data Usage

We now evaluate the impact of a regulation which prohibits the platform from using marketplace data to inform product introduction. We will refer to such a regulation as an *insider imitation ban*. In the context of our model, this restriction amounts to requiring that the platform choose a competition policy which is independent of the demand state α. Under such a regulation, uncertainty in market demand is immaterial to both the platform and the entrepreneur, as each player's payoff under any permitted policy is a function only of average profits under each market structure. As a result, the optimal competition policy under the regulation is exactly the optimal policy of the no-uncertainty benchmark characterized in Section 3.

In Section 5.1, we characterize the impact of an insider imitation ban on the quantity of innovation. In Section 5.2, we use this result to assess when the regulation increases or decreases social welfare. Throughout this section, we will maintain a linear demand specification, as in Section 4.

5.1 Does Regulation Stimulate Innovation?

We first determine how an insider imitation ban impacts the total quantity of innovation. Since innovation is increasing in the threshold innovation cost, this comparison is equivalent to characterizing the sign of $\kappa^* - \kappa^{**}$, recalling that κ^* is the platform's optimal innovation cost threshold absent regulation, while κ^{**} is the optimal threshold under regulation.

Our results establish that regulation stimulates innovation if the demand distribution F has a sufficiently large *upside demand potential*, with a significant fraction of total demand distributed far above the mean; and that conversely regulation stifles innovation if the demand distribution has little upside potential, with most demand distributed below the mean. These findings are formally captured in the following pair of propositions:

Proposition 1. *There exists an $\varepsilon > 0$ such that for any demand distribution F, if*

$$\int_{\{\alpha<1\}} dF(\alpha)\,\alpha > 1 - \varepsilon$$

then $\kappa^* > \kappa^{**}$.

Proposition 2. *There exists a demand state $\overline{\alpha} > 1$ and an $\varepsilon > 0$ such that for any demand distribution F, if*

$$\int_{\{\alpha \geq \overline{\alpha}\}} dF(\alpha)\, \alpha > 1 - \varepsilon$$

then $\kappa^* < \kappa^{**}$.

These two propositions connect the impact of regulation to the behavior of the *tail weight function*

$$\Phi(\alpha) \equiv \int_{\{\alpha' \geq \alpha\}} dF(\alpha')\, \alpha'.$$

This function tracks the fraction of total demand lying above a given threshold demand state. By definition $\Phi(0) = 1$, as all demand must be accounted for at some demand state and total demand is normalized to 1. Additionally, Φ is decreasing in α and satisfies $\lim_{\alpha \to \infty} \Phi(\alpha) = 0$. We will refer to $\Phi(\alpha)$ as the α*-tail weight*, and $\Phi(1)$ in particular as the *mean tail weight*. Proposition 1 says that if the mean tail weight is sufficiently close to 0, then the regulation must reduce innovation. Meanwhile Proposition 2 says that if, for a large enough $\overline{\alpha}$, the $\overline{\alpha}$-tail weight is sufficiently close to 1, then the regulation boosts innovation.

Distributions with a large $\overline{\alpha}$-tail weight correspond to *experimental* products: They exhibit a high incidence of demand at very high demand states, counterbalanced by a high probability of very low demand. For instance, consider a 2-point demand distribution with mass on $\alpha \in \{\alpha_*, \alpha^*\}$, where $\alpha^* \geq \overline{\alpha}$ and $\alpha_* < 1$.[12] The mass $\rho = \Delta F(\alpha^*)$ placed on state α^* must ensure that the mean demand is 1, i.e. $\rho = (1 - \alpha_*)/(\alpha^* - \alpha_*) \in (0, 1)$. This distribution has an $\overline{\alpha}$-tail weight

$$\Phi(\overline{\alpha}) = \rho\alpha^* = \frac{1 - \alpha_*}{1 - \alpha_*/\alpha^*},$$

which may be taken arbitrarily close to 1 by taking α_* close to zero. In this limit ρ converges to $1/\alpha^*$, meaning that the complementary probability $1 - 1/\alpha^*$ is placed on a state with very little demand. This probability may be very close to 1 if α^* is large, indicating a significant

[12] We have imposed the regularity condition that F have support on an interval, and so formally such a distribution must also have a positive density on the interval (α_*, α^*). But this density may be taken arbitrarily small, approximating a 2-point distribution as closely as desired.

chance that the product fails to find a market. For such experimental products, a ban on insider imitation will lead to increased innovation.

Distributions with a small mean tail weight, on the other hand, correspond to *incremental* products: They exhibit very little demand above the mean, with a corresponding concentration of overall demand near the mean. Consider the same two-point distribution as in the previous paragraph, with the requirement that $\alpha^* \geq \overline{\alpha}$ weakened to $\alpha^* > 1$. The mean tail weight of this distribution is $\Phi(1) = \rho\alpha^*$, which may be made arbitrarily small by taking α_* to 1. Note that in this limit, ρ converges to zero. Thus not only does the lower mass move toward 1, but all of the weight of the distribution shifts toward this mass. This limit therefore corresponds to a demand distribution which is concentrated around the mean. For such incremental products, a ban on insider imitation leads to a reduction in innovation.

The demand tail weights govern the impact of regulation because they control the marginal demand state implementing a target innovation cost threshold. When the right tail of the distribution has significant weight, most cost thresholds correspond to threshold demand states below the mean. At such demand thresholds, the marginal conversion rate is lower than the average rate. In particular, the demand threshold $\alpha^*(\kappa^{**})$ which would implement the optimal regulated cost threshold falls below the mean. As a result, the platform finds imitation less attractive absent the regulation versus under it, in line with Lemma 2 and the discussion which follows. This force leads the regulation to suppress innovation. The opposite logic holds when the right tail has significant weight and most cost thresholds, in particular the optimal regulated threshold κ^{**}, are implemented with demand thresholds significantly above the mean, leading to reduced imitation and a rise in innovation under regulation.

The one factor complicating this reasoning is that, in general, the unregulated platform's incentive to imitate is lower than indicated by the marginal conversion rate, as highlighted in the discussion following Lemma 4 and reflected in the remainder term $R(\alpha^*(\kappa))$ in the unregulated FOC. As a result, the marginal demand state must be sufficiently far above the mean state for the larger conversion rate to imply an increase in imitation absent the regulation. This force drives the threshold state $\overline{\alpha}$ appearing in Proposition 2 above 1.

The logic underpinning Propositions 1 and 2 connects a product's upside potential, an exogenous feature of the demand distribution, to whether an endogenous marginal demand state lies above or below the mean. This finding is reminiscent of the results of Johnson and Myatt (2006), who identify a demand rotation transformation which controls whether a monopolist's profit-maximizing marginal customer has a value above or below the value of the average customer. While their concept of a demand rotation differs from our notion of right-tail weight, both results similarly emphasize how the shape of demand may impact market outcomes via a comparison between marginal and average demand states.

5.2 Welfare

We now evaluate the implications of an insider imitation ban for social welfare. We have seen that, depending on the shape of the distribution of demand, an insider imitation ban may either stimulate or suppress innovation. Whether this outcome is good for welfare depends on whether a regulator, concerned with maximizing social welfare, prefers more or less innovation than would prevail in the absence of regulation. We will show that, holding fixed an arbitrary demand distribution, the direction of the inefficiency may go either direction, depending on how much consumer welfare increases following imitation. As a result, an insider imitation ban improves welfare only in product markets satisfying particular combinations of demand uncertainty and consumer surplus gains from imitation.

We perform our welfare analysis under the following regularity assumption:

Assumption 5. $\Delta w > \Delta \mu_P$.

This assumption ensures that the platform does not fully internalize the welfare gains from imitation. It is likely to hold in any market in which introduction of an imitation product sufficiently lowers prices, so that consumer surplus gains from imitation are high. In particular, Assumption 5 may be equivalently written $\nu^D > \underline{\nu}^D$, where ν^D is average consumer surplus under duopoly and $\underline{\nu}^D \equiv \nu^M + |\Delta \mu_E|$.

We first compare the platform's optimal policy to the regulator's preferred outcome if it could freely design the platform's competition policy. The regulator's optimal policy can be derived in the same way as the platform's, substituting total welfare for platform profits throughout the derivation. An analog of Lemma 3 can be shown to hold, so that the regulator would optimally direct the platform to imitate immediately above some threshold demand state, and to never imitate otherwise, with the threshold state chosen to induce a target innovation cost threshold. Any disagreement between the regulator and platform's optimal policies is therefore confined to the choice innovation cost threshold. The regulator's optimal innovation cost threshold, which we will denote κ^{FB}, can then be shown to satisfy a first-order condition similar to the one derived in Lemma 4.

The following lemma characterizes how κ^{FB} compares to κ^* as ν^D, the average consumer welfare under duopoly, varies. Note that varying ν^D does not change κ^* or κ^{**}, as consumer surplus does not enter into the platform's profit function.

Lemma 5. *Fix all model parameters except for ν^D. Then κ^{FB} is strictly decreasing in ν^D. Further, there exists a threshold $\overline{\nu}^D > \underline{\nu}^D$ such that $\kappa^{FB} = \kappa^*$ if $\nu^D = \overline{\nu}^D$.*

This lemma establishes that the regulator prefers more innovation than the platform when ν^D is small, but less innovation when ν^D is large. This result reflects the fact that ν^D mediates the regulator's tradeoff between collecting the baseline welfare w^M from additional products under a monopolized market, versus the enhanced welfare w^D from fewer products under a duopoly. As ν^D rises, the regulator increasingly prefers imitation at the cost of innovation, and its optimal cost threshold falls relative to the platform's.

The following result establishes that for sufficiently extreme values of ν^D, an insider imitation ban impacts welfare in a manner reflecting the comparison between κ^{FB} and κ^*.

Proposition 3. *Fix all model parameters except for ν^D.*

1. *Fix a demand distribution F such that $\kappa^* < \kappa^{**}$. Then if ν^D is sufficiently large, an insider imitation ban strictly reduces social welfare.*

2. *Fix a demand distribution F such that $\kappa^* > \kappa^{**}$. Then:*

   - *If ν^D is sufficiently close to $\underline{\nu}^D$, an insider imitation ban strictly reduces social welfare.*

   - *If ν^D is sufficiently large, an insider imitation ban strictly increases social welfare.*

There are two reasons why this proposition characterizes the impact of the regulation on welfare only for extreme values of ν^D. The first reason is that for intermediate values of ν^D, the regulation may overshoot, moving the cost threshold further than the regulator would prefer to do if it could directly control the threshold. In particular, if ν^D is close to the value at which $\kappa^{FB} = \kappa^*$, then the regulator would prefer only a small change in the threshold, but the difference between κ^* and κ^{**} is independent of ν^D and so will in general be larger than optimal.

The second reason is that imposing the regulation always imposes an inefficiency, due to the elimination of the platform's ability to target imitation toward high demand states. This targeting is always beneficial for welfare, as it transforms entrepreneur profits into social welfare more efficiently. As a result, if the platform's optimal cost threshold remained fixed, then the regulation would necessarily degrade welfare. If the regulation is to improve social welfare, the efficiency gains achieved by moving from κ^* to κ^{**} must be sufficiently large so as to outweigh the inefficiency from eliminating targeting. The proof of the proposition establishes that this outcome arises only for sufficiently large ν^D.

The inefficiency arising from eliminating targeting makes it impossible to guarantee that, for sufficiently small values of ν^D, an insider imitation ban necessarily increases welfare when the ban leads to increased innovation. This is because for such values of ν^D, the inefficiency from lost targeting looms large. Depending on model parameters, a range of ν^D for which regulation is good for welfare may or may not exist. Perhaps surprisingly, an insider imitation ban can be ensured to increase welfare for some consumer surplus specifications only when the ban would lead to *less* innovation.

6 Extensions

6.1 Nonlinear demand states

We now extend our analysis to general nonlinear demand specifications. A basic complication arising when demand is nonlinear is that the conversion rate $\Lambda(\alpha)$ is not necessarily monotone. As a result, an optimal competition policy is not guaranteed to be a threshold rule in α. To avoid this complication and generalize our results from the linear case, we impose the following regularity assumption:

Assumption 6. $\Lambda(\alpha)$ *is strictly increasing in* α*, and* $\lim_{\alpha\to\infty}\Lambda(\alpha) < \infty$*.*

This assumption guarantees that as α increases, imitation transfers profits more efficiently from the entrepreneur to the platform. As a result, the most profitable way to induce a given level of post-entry profits for the entrepreneur is via a threshold rule. It also bounds the limiting conversion rate at high demand states, preserving a useful technical property satisfied by $\Lambda(\alpha)$ in the linear case.[13]

Under Assumption 6, the platform's optimal imitation threshold κ^* satisfies a first-order condition characterized in Lemma 4:

$$\frac{G'(\kappa^*)}{G(\kappa^*)}\left(\overline{\Pi}_P^M + \Lambda(\alpha^*(\kappa))(\overline{\Pi}_E^M - r\kappa^*) + R(\alpha^*(\kappa))\right) = r\Lambda(\alpha^*(\kappa)),$$

where now

$$R(\alpha) \equiv \int_{\{\alpha'>\alpha\}} dF(\alpha')\,\left(\Lambda(\alpha') - \Lambda(\alpha)\right)|\Delta\Pi_E(\alpha')|.$$

This remainder term reflects the difference between the marginal conversion rate $\Lambda(\alpha)$ and each of the inframarginal states $\Lambda(\alpha')$, weighted by the amount of profit $|\Delta\Pi_E(\alpha')|$ extracted from the entrepreneur in each demand state. Note that when demand is linear,

$$\left(\Lambda(\alpha') - \Lambda(\alpha)\right)|\Delta\Pi_E(\alpha')| = rk_P\left(\frac{\alpha'}{\alpha} - 1\right),$$

recovering the form of the remainder term obtained in the baseline analysis.

6.2 Transfers

We now study the platform's imitation problem when an additional tool is available: The platform may make lump-sum transfers to or from an entrepreneur who wishes to list their product on the platform. As in the baseline problem, there is no value to conditioning the transfer or the competition policy on a reported cost type by the entrepreneur. Additionally, without loss the platform does not condition the price on the realized demand α, as this

[13] Note that whenever F has support on a bounded interval, the limiting behavior of Λ is irrelevant and no assumptions need be placed on it.

simply introduces randomization into the transfer, and both the platform and entrepreneur care only about the total expected transfer. It is therefore sufficient to consider schemes of the form $(T(\cdot), p)$, where $T(\cdot)$ is a competition policy and p is a price charged to the entrepreneur to list on the platform (which could be negative). We will refer to such a scheme as a *listing policy*. Real-world practices vary on this front. For example, Walmart does not charge sign-up or fixed fees, whereas Amazon offers a plan with a fixed monthly fee.[14]

6.2.1 The unregulated problem

As in the baseline problem, any listing policy induces entrepreneurs below some threshold cost type to innovate. The platform's policy design problem can therefore be solved by first characterizing an optimal policy inducing a given entry threshold, and then optimizing over the threshold. Given a cost type κ to be made indifferent over entry, the platform's profits are maximized by first choosing a competition policy which maximizes industry surplus, and then setting the price of entry to extract all surplus from the marginal entrepreneur. Surplus is maximized by immediate imitation for demand states above the industry break-even state α^I satisfying

$$\alpha^I(\Delta\mu_P + \Delta\mu_E) - rk_P = 0.$$

The associated optimal listing policy is

$$T^*(\alpha, \kappa) = \begin{cases} 0, & \alpha \geq \alpha^I \\ \infty, & \alpha < \alpha^I \end{cases}, \quad p^*(\kappa) = r^{-1}\left(\mu_E^M + \Delta\mu_E \int_{\{\alpha \geq \alpha^I\}} dF(\alpha)\, \alpha\right) - \kappa.$$

Unlike in the baseline problem, the platform faces no restrictions on the threshold types it may induce—any choice of κ in $[\underline{k}, \overline{k}]$ is implementable.

Whether p^* is positive or negative depends on how α^I compares to $\alpha^*(\kappa)$, which is the break-even demand state absent payments. If $\alpha^*(\kappa) < \alpha^I$, then $p^*(\kappa)$ is positive, while if $\alpha^*(\kappa) > \alpha^I$, then $p^*(\kappa)$ is negative and entrepreneurs are subsidized to enter. Further, the

[14] For Walmart marketplace fees, see https://tinyurl.com/3dn22hnx. For Amazon fees, see https://tinyurl.com/46u8tmwe.

price charged to the entrepreneur declines as κ rises.

Define Π_I^* to be total industry flow profits under an optimal listing policy. That is,

$$\Pi_I^* = \mu_E^M + \mu_P^M + \int_{\{\alpha \geq \alpha^I\}} dF(\alpha)\,(\alpha(\Delta\mu_E + \Delta\mu_P) - rk_P).$$

Then the platform's profits as a function of κ are

$$U_P(\kappa) = G(\kappa)(r^{-1}\Pi_I^* - \kappa).$$

Since $\Pi_I^* > \Pi_E^M > r\overline{k}$, this expression is strictly positive for all choices of κ. The platform's optimal innovation threshold κ^*, if it is interior, satisfies the first-order condition

$$\frac{G'(\kappa)}{G(\kappa)}(\Pi_I^* - r\kappa) = 1.$$

Our assumptions guarantee that κ^* is interior, because $G'(\overline{k}) = 0$ by Assumption 1, while $G'(\kappa)/G(\kappa)$ goes to ∞ at $\kappa = \underline{k}$ given our monotone hazard rate condition combined with $G(\underline{k}) = 0$.

6.2.2 The regulated problem

Now consider imposing the usual regulation that the platform may not condition T on α. It continues to be the case that, conditional on κ, the platform chooses a competition policy which maximizes industry profits (subject to the regulatory constraint) and then charges a price which extracts all surplus from the marginal entrepreneur. So an optimal listing policy is

$$T^{**}(\kappa) = \begin{cases} 0, & \alpha^I \leq 1 \\ \infty, & \alpha^I > 1 \end{cases}, \quad p^*(\kappa) = r^{-1}(\mu_E^M + \Delta\mu_E e^{-\rho T^{**}(\kappa)}) - \kappa.$$

Let Π_I^{**} be the industry's optimal flow profits under the regulation:

$$\Pi_I^{**} = \max\{\mu_E^M + \mu_P^M, \mu_E^D + \mu_P^D - rk_P\}.$$

Then the platform's regulated profits as a function of κ are

$$\widetilde{U}_P(\kappa) = G(\kappa)(r^{-1}\Pi_I^{**} - \kappa).$$

Because $\Pi_I^{**} \geq \mu_E^M > r\overline{k}$, the platform's optimal regulated innovation threshold κ^{**} is interior and satisfies the first-order condition

$$\frac{G'(\kappa)}{G(\kappa)}(\Pi_I^{**} - r\kappa) = 1.$$

In general $\Pi_I^{**} \leq \Pi_I^{*}$, and so $\kappa^{**} \leq \kappa^{*}$. Further, whenever α^I lies in the interior of the support of F both inequalities must be strict.

Note that unlike in the problem without transfers, where the impact of the regulation on innovation was ambiguous, with transfers the regulation unambiguously reduces innovation.

6.2.3 When is regulation optimal?

We now compare the regulator's welfare with and without regulation. In case α^I does not lie in the interior of the support of F, then the regulation does not change market outcomes and so has no impact on welfare. So assume going forward that $\Pr(\alpha > \alpha^I) > 0$ and $\Pr(\alpha < \alpha^I) > 0$.

When unrestricted transfers are available, the regulation always strictly reduces entry. If it additionally weakly reduces total welfare post-innovation, then the regulation is unambiguously bad for ex ante welfare.

Suppose first that $\alpha^I > 1$, so that the regulation suppresses all imitation. Then total welfare drops post-innovation if

$$w^M + \int_{\{\alpha \geq \alpha^I\}} dF(\alpha)\,(\alpha \Delta w - rk_P) \geq w^M.$$

Since $\Delta w = \Delta\mu_E + \Delta\mu_P + \Delta\nu > \Delta\mu_E + \Delta\mu_P$, we must have $\alpha^I \Delta w - rk_P > 0$. Therefore

$$\int_{\{\alpha \geq \alpha^I\}} dF(\alpha)\,(\alpha \Delta w - rk_P) > 0.$$

So welfare indeed drops post-imitation. Hence the regulation is unambiguously bad for welfare whenever its effect is to suppress imitation.

Suppose instead that $\alpha^I < 1$, so that the regulation leads to imitation in all demand states. Then total welfare drops post-innovation if

$$w^M + \int_{\{\alpha \geq \alpha^I\}} dF(\alpha)\,(\alpha \Delta w - rk_P) \geq w^D - rk_P,$$

or equivalently, if

$$\int_{\{\alpha \geq \alpha^I\}} dF(\alpha)\,(\alpha \Delta w - rk_P) < 0.$$

This inequality holds whenever ν^D is sufficiently close to ν^M. So the regulation harms welfare if it increases imitation but the consumer surplus gains from imitation are small.

References

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar (2021). "Too Much Data: Prices and Inefficiencies in Data Markets". *American Economic Journal: Microeconomics*, Forthcoming.

Anderson, Chris (2006). *The Long Tail: Why the Future of Business Is Selling Less of More*. New York: Hyperion.

Bergemann, Dirk, Alessandro Bonatti, and Tan Gan (2021). "The Economics of Social Data". Working Paper.

Bose, Shubhomita (July 29, 2018). "How to Get Your Product in Walmart". *Small Business Trends*. URL: https://smallbiztrends.com/2018/05/get-your-product-sold-at-walmart.html.

Brynjolfsson, Erik, Yu (Jeffrey) Hu, and Michael D. Smith (2003). "Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers". *Management Science* 49.11, pp. 1580–1596.

Brynjolfsson, Erik, Yu (Jeffrey) Hu, and Michael D. Smith (2006). "From niches to riches: The anatomy of the long tail". *Sloan Management Review* 47.4, pp. 67–71.

Coppola, Daniela (Nov. 27, 2020). "Online share of total U.S. retail revenue in 2020, by product category". *Statista*. URL: https://www.statista.com/statistics/203043/online-share-of-total-us-retail-revenue-projection/.

Eilat, Ran, Kfir Eliaz, and Xiaosheng Mu (2020). "Bayesian Privacy". *Theoretical Economics*, Forthcoming.

Etro, Federico (2020). "Product Selection in Online Marketplaces". Working paper.

European Commission (Nov. 10, 2020). "Antitrust: Commission sends Statement of Objections to Amazon for the use of non-public independent seller data and opens second investigation into its e-commerce business practices". URL: https://ec.europa.eu/commission/presscorner/detail/en/ip_20_2077.

Fainmesser, Itay P., Andrea Galeotti, and Ruslan Momot (2021). "Digital Privacy". Working Paper.

Gomes, Renato and Alessandro Pavan (2021). "Price Customization and Targeting in Matching Markets". Working Paper.

Hagiu, Andrei, Tat-How Teh, and Julian Wright (2020). "Should platforms be allowed to sell on their own marketplaces?" Working paper.

Henry, Emeric and Carlos J. Ponce (2011). "Waiting to imitate: on the dynamic pricing of knowledge". *Journal of Political Economy* 119.5, pp. 959–981.

Hidir, Sinem and Nikhil Vellodi (2020). "Personalization, Discrimination and Information Disclosure". *Journal of the European Economic Association*, forthcoming.

Hopenhayn, Hugo A. and Francesco Squintani (2016). "Patent Rights and Innovation Disclosure". *The Review of Economic Studies* 83.1, pp. 199–230.

"How to Sell to Walmart" (Nov. 1, 2020). *8th & Walton*. URL: https://www.8thandwalton.com/blog/how-to-sell-to-walmart/.

Ichihashi, Shota (2019). "Online Privacy and Information Disclosure by Consumers". *American Economic Review* 110.2, pp. 569–595.

Ichihashi, Shota (2020). "Dynamic Privacy Choices". Working Paper.

Johnson, Justin P. and David P. Myatt (2006). "On the Simple Economics of Advertising, Marketing and Product Design". *American Economic Review* 96.3.

Jovanovic, Boyan and Saul Lach (1989). "Entry, Exit, and Diffusion with Learning by Doing". *American Economic Review* 79.4, pp. 690–99.

Khan, Lina M. (2017). "Amazon's Antitrust Paradox". *Yale Law Journal* 126.3, pp. 710–805.

Liang, Annie and Erik Madsen (2021). "Data and Incentives". *working paper*.

Lipsman, Andrew (Mar. 10, 2020). "Top 10 US Ecommerce Companies 2020". *eMarketer*. URL: https://www.emarketer.com/content/top-10-us-ecommerce-companies-2020.

Mattioli, Dana (Apr. 23, 2020). "Amazon Scooped Up Data From Its Own Sellers to Launch Competing Products". *The Wall Street Journal*. URL: https://www.wsj.com/articles/amazon-scooped-up-data-from-its-own-sellers-to-launch-competing-products-11587650015.

Miranda, Leticia (June 13, 2018). "Amazon Sellers Say The Tech Giant Is Crushing Them With Competitive Pricing". *Buzzfeed News*. URL: https://www.buzzfeednews.com/

article/leticiamiranda/amazon-sellers-say-the-tech-giant-is-crushing-them-with.

Rob, Rafael (1991). “Learning and Capacity Expansion under Demand Uncertainty”. *The Review of Economic Studies* 58.4, pp. 655–675.

Soper, Spencer (Apr. 20, 2016). “Got a Hot Seller on Amazon? Prepare for E-Tailer to Make One too”. *Bloomberg*. URL: https://www.bloomberg.com/news/articles/2016-04-20/got-a-hot-seller-on-amazon-prepare-for-e-tailer-to-make-one-too.

“The State of the Amazon Marketplace 2019” (Nov. 27, 2020). *Feedvisor*. URL: https://feedvisor.com/resources/amazon-trends/the-state-of-the-amazon-marketplace-2019/.

“The State of the Amazon Seller 2021” (2021). *JungleScout*. URL: https://www.junglescout.com/wp-content/uploads/2021/02/State-of-the-Amazon-Seller-2021-Jungle-Scout.pdf.

U.S. House Committee on the Judiciary (2020). “Investigation of Competition in Digital Markets: Majority Staff Report and Recommendations”. URL: https://judiciary.house.gov/uploadedfiles/competition_in_digital_markets.pdf.

A Proofs

A.1 Lemma 1

The optimal threshold κ^{**} must satisfy the FOC

$$\frac{G'(\kappa^{**})}{G(\kappa^{**})}\left(\overline{\Pi}_P^M + (\Delta\overline{\Pi}_P - rk_P)\frac{\overline{\Pi}_E^M - r\kappa^{**}}{|\Delta\overline{\Pi}_E|}\right) \leq r\frac{\Delta\overline{\Pi}_P - rk_P}{|\Delta\overline{\Pi}_E|},$$

with equality if the lower bound constraint $\kappa \geq \widetilde{k}_0$ does not bind. Note that given our monotone hazard rate condition, the lhs of the FOC is strictly decreasing in κ^{**}, while the rhs is independent of κ^{**}. Thus either there exists a unique interior optimum, or else the lhs is smaller than the rhs for all $\kappa^{**} > \widetilde{k}_0$, and the lower bound constraint binds. Assumption 4 ensures the former case obtains.

A.2 Lemma 2

Fix $\Lambda > 0$. By arguments very similar to the proof of Lemma 1, $\widetilde{\kappa}(\Lambda)$ can be shown to satisfy $\widetilde{\kappa}(\Lambda) < \overline{k}$ and $\Omega(\widetilde{\kappa}(\Lambda), \Lambda) \leq 0$, with equality whenever $\widetilde{\kappa}(\Lambda) > \widetilde{k}_0$, where

$$\Omega(\kappa, \Lambda) \equiv \frac{G'(\kappa)}{G(\kappa)}\left(\overline{\Pi}_P^M + (\overline{\Pi}_E^M - r\kappa)\Lambda\right) - r\Lambda.$$

Now, since $\overline{\Pi}_P^M > 0$, it must be that

$$\frac{G'(\widetilde{\kappa}(\Lambda))}{G(\widetilde{\kappa}(\Lambda))}(\overline{\Pi}_E^M - r\widetilde{\kappa}(\Lambda)) < r.$$

Therefore $\Omega(\kappa, \Lambda)$ is decreasing in Λ when $\kappa = \widehat{\kappa}(\Lambda)$. Since Ω is additionally decreasing in κ everywhere, it must be that $\widehat{\kappa}(\Lambda') \leq \widehat{\kappa}(\Lambda)$ if $\Lambda' > \Lambda$, and this inequality must hold strictly if $\widehat{\kappa}(\Lambda) > \widetilde{k}_0$.

A.3 Lemma 3

The Lagrangian for the competition policy design problem is

$$\mathscr{L} = G(\kappa)r^{-1}\left(\mu_P^M + \int dF(\alpha)\, e^{-rT(\alpha)}(\alpha\Delta\mu_P - rk_P)\right) \tag{A.1}$$

$$+ \lambda r^{-1}\left(\mu_E^M + \Delta\mu_E \int dF(\alpha)\, e^{-rT(\alpha)}\alpha - r\kappa\right). \tag{A.2}$$

Re-arranging this expression yields

$$\mathscr{L} = r^{-1}\int dF(\alpha)\, e^{-rT(\alpha)}(\alpha(G(\kappa)\Delta\mu_P + \lambda\Delta\mu_E) - G(\kappa)rk_P) + \text{const},$$

where "const" are additive constants not depending on $T(\cdot)$.

We look for possible saddle points of this problem in two different regimes, depending on the sign of $G(\kappa)\Delta\mu_P + \lambda\Delta\mu_E$. If this expression is strictly is strictly positive, then letting $\alpha^\dagger(\lambda) \equiv rk_P/(\Delta\mu_P + \Delta\mu_E\lambda/G(\kappa)) > 0$, the Lagrangian is maximized by setting

$$T(\alpha) = \begin{cases} 0, & \alpha > \alpha^\dagger(\lambda), \\ \infty, & \alpha < \alpha^\dagger(\lambda), \end{cases}$$

with any choice of $T(\alpha^\dagger(\lambda))$ consistent with optimality. Such a policy satisfies the zero-profit condition iff

$$\mu_E^M + \Delta\mu_E\left(\int_{\{\alpha>\alpha^\dagger(\lambda)\}} dF(\alpha)\,\alpha + \Delta F(\alpha^\dagger(\lambda))e^{-rT(\alpha^\dagger(\lambda))}\right) = r\kappa.$$

The threshold demand state $\alpha^\dagger(\lambda)$ may be set anywhere in the interval $(0,\infty)$ by an appropriate choice of λ. By definition of $\alpha^*(\kappa)$, the zero-profit condition can be satisfied only if $\alpha^\dagger(\lambda) = \alpha^*(\kappa)$. This demand threshold corresponds to the Lagrange multiplier

$$\lambda^*(\kappa) \equiv G(\kappa)\Lambda(\alpha^*(\kappa)).$$

Further, the condition $\mu_E^M > r\overline{k}$ implies that $\alpha^*(\kappa) < \infty$ for every $k \le \overline{k}$, so that when $\alpha^\dagger(\lambda) = \alpha^*(\kappa)$, an appropriate value of $T(\alpha^*(\kappa))$ ensures that the zero profit condition will be satisfied. If F is continuous at $\alpha = \alpha^*(\kappa)$, then any associated choice of $T(\alpha^*(\kappa))$ satisfies the zero-profit condition; in particular, we may choose $T(\alpha^*(\kappa)) = 0$. Otherwise, there exists a unique $\overline{T}(\kappa) \in \mathbb{R}_+ \cup \{\infty\}$ such that the zero-profit condition is satisfied when $\alpha^\dagger(\lambda) = \alpha^*(\kappa)$. This characterizes the essentially unique saddle point satisfying $G(\kappa)\Delta\mu_P + \lambda\Delta\mu_E > 0$.

Suppose instead that $\Delta\mu_P + \lambda\Delta\mu_E$ is non-positive. In this case, the integrand in the Lagrangian is everywhere strictly negative, and the Lagrangian is maximized by setting $T(\alpha) = \infty$ everywhere. But then the marginal entrepreneur's profits from entry are $r^{-1}\mu_E^M - \kappa$, which does not satisfy the zero-profit condition. So no such saddle point exists.

As an essentially unique saddle point exists, it must represent the solution to the platform's problem. It remains only to prove that $\alpha^*(\kappa)$ is continuous, increasing, and satisfies $\alpha^*(\kappa) \geq \alpha_0$ for all $\kappa \geq k_0$. Let

$$V(\alpha) \equiv \mu_E^M + \Delta\mu_E \int_{\{\alpha' \geq \alpha\}} dF(\alpha')\,\alpha'.$$

Then $V(\cdot)$ is an increasing function, and by continuity of measure it is left-continuous. First observe that $V(\alpha_0) = rk_0$ by definition of k_0, and so $\alpha^*(\kappa) \geq \alpha_0$ for all $\kappa \geq k_0$. Next, left-continuity combined with the definition of $\alpha^*(\kappa)$ implies that $V(\alpha^*(\kappa)) \leq r\kappa$ for all κ. But then also $V(\alpha^*(\kappa)) \leq r\kappa'$ for every $\kappa' > \kappa$, implying that $\alpha^*(\kappa') \geq \alpha^*(\kappa)$, i.e. $\alpha^*(\cdot)$ is an increasing function.

Now, suppose that $\lim_{\kappa' \uparrow \kappa} \alpha^*(\kappa') = \alpha^*(\kappa-) < \alpha^*(\kappa)$. Then for every $\alpha' \in (\alpha^*(\kappa-), \alpha^*(\kappa))$, we have $V(\alpha') > r\kappa'$ for all $\kappa' < \kappa$, meaning that $V(\alpha') \geq r\kappa$. Since $V(\alpha^*(\kappa)) \leq r\kappa$, it must be that $V(\alpha') \geq V(\alpha^*(\kappa))$. Given that V is increasing and $\alpha^*(\kappa) > \alpha'$, this inequality can hold only if $V(\alpha') = V(\alpha^*(\kappa))$. In other words, V is constant on the interval $(\alpha^*(\kappa-), \alpha^*(\kappa)]$. But we must also have $V(\alpha) > V(\alpha^*(\kappa))$ for all $\alpha > \alpha^*(\kappa)$, and $V(\alpha) > V(\alpha^*(\kappa-))$ for every $\alpha > \alpha^*(\kappa-)$. In other words, both $\alpha^*(\kappa-)$ and $\alpha^*(\kappa)$ must lie in the support of F. But by assumption the support of F is an interval, meaning that V must be strictly increasing on $(\alpha^*(\kappa-), \alpha^*(\kappa)]$, a contradiction. So it must be that $\alpha^*(\kappa-) = \alpha^*(\kappa)$, i.e. α^* is left-continuous.

Suppose instead that $\lim_{\kappa' \downarrow \kappa} \alpha^*(\kappa') = \alpha^*(\kappa+) > \alpha^*(\kappa)$. Then for every $\alpha' \in (\alpha^*(\kappa), \alpha^*(\kappa+))$, we have $V(\alpha') \leq r\kappa'$ for all $\kappa' > \kappa$, meaning that $V(\alpha') \leq r\kappa$. But also $V(\alpha') > r\kappa$ given that $\alpha' > \alpha^*(\kappa)$, a contradiction. So it must be that $\alpha^*(\kappa+) = \alpha^*(\kappa)$, i.e. α^* is right-continuous.

A.4 Lemma 4

Recall that $U_P(\kappa)$ is the value function of the maximization problem

$$G(\kappa)r^{-1}\left(\mu_P^M + \int dF(\alpha)e^{-rT(\alpha)}(\alpha\Delta\mu_P - rk_P)\right)$$

subject to the constraint

$$r^{-1}\left(\mu_E^M + \Delta\mu_E \int dF(\alpha)e^{-rT(\alpha)}\alpha - r\kappa\right) = 0.$$

The proof of Lemma 3 establishes that $\lambda^*(\kappa)$, the Lagrange multiplier on the constraint at the optimal policy, satisfies $\lambda^*(\kappa) = G(\kappa)\Lambda(\alpha^*(\kappa))$. Hence by the envelope theorem

$$U_P'(\kappa) = G'(\kappa)r^{-1}\left(\mu_P^M + \int dF(\alpha)\, e^{-rT^*(\alpha,\kappa)}(\alpha\Delta\mu_P - rk_P)\right) - G(\kappa)\Lambda(\alpha^*(\kappa)).$$

We characterize the maximum of U_P by establishing that U_P' crosses zero exactly once. To this end, define

$$\Omega(\kappa) \equiv \frac{G'(\kappa)}{G(\kappa)}\left(\mu_P^M + \int dF(\alpha)\, e^{-rT^*(\alpha,\kappa)}(\alpha\Delta\mu_P - rk_P)\right) - r\Lambda(\alpha^*(\kappa)).$$

By construction, the signs of Ω and U_P' agree everywhere. We first establish that $\Omega(\kappa)$ is strictly decreasing in κ. The monotone hazard rate assumption on G ensures that $G'(\kappa)/G(\kappa)$ is weakly decreasing in κ. Meanwhile, Lemma 3 implies that $T^*(\alpha, \kappa)$ is increasing in κ for every α, and that additionally

$$\int dF(\alpha)\, e^{-rT^*(\alpha,\kappa)}\alpha = \frac{\mu_E^M - r\kappa}{|\Delta\mu_E|},$$

so that the lhs is strictly decreasing in κ. Thus $T^*(\alpha, \kappa)$ must strictly increase on a positive-measure set of α when κ increases. This result implies that

$$\int dF(\alpha)\, e^{-rT^*(\alpha,\kappa)}(\alpha\Delta\mu_P - rk_P)$$

is strictly decreasing in κ for all $\kappa \geq k_0$. Finally,

$$\Lambda(\alpha) = \frac{\alpha\Delta\mu_P - rk_P}{\alpha|\Delta\mu_E|} = \frac{\Delta\mu_P}{|\Delta\mu_E|} - \frac{rk_P}{\alpha|\Delta\mu_E|},$$

which is strictly increasing in α. Then since $\alpha^*(\kappa)$ is increasing in κ, it must be that $\Lambda(\alpha^*(\kappa))$ is increasing in κ. It follows that $\Omega(\kappa)$ is strictly decreasing in κ.

We next establish that Ω is continuous in κ. To do so, we first rewrite Ω in terms of $R(\kappa)$, as in the lemma statement. Combining the algebraic identity

$$\begin{aligned}\int dF(\alpha)\, e^{-rT^*(\alpha,\kappa)}\,(\alpha\Delta\mu_P - rk_P) = \Lambda(\alpha^*(\kappa)) \int dF(\alpha)\, e^{-rT^*(\alpha,\kappa)}\alpha& \\ + rk_P \int dF(\alpha)\, e^{-rT^*(\alpha,\kappa)} \left(\frac{\alpha}{\alpha^*(\kappa)} - 1\right)&,\end{aligned}$$

with the entrepreneur's zero-profit condition, which may be written

$$\int dF(\alpha)\, e^{-rT^*(\alpha,\kappa)}\alpha = \frac{\mu_E^M - r\kappa}{|\Delta\mu_E|},$$

yields the identity

$$\begin{aligned}&\int dF(\alpha)\, e^{-rT^*(\alpha,\kappa)}\,(\alpha\Delta\mu_P - rk_P) \\ &= rk_P \int dF(\alpha)\, e^{-rT^*(\alpha,\kappa)} \left(\frac{\alpha}{\alpha^*(\kappa)} - 1\right) + (\mu_E^M - r\kappa)\Lambda(\alpha^*(\kappa)).\end{aligned}$$

Given that $T^*(\alpha,\kappa) = \infty$ for $\alpha < \alpha^*(\kappa)$ while $T^*(\alpha,\kappa) = 0$ for $\alpha > \alpha^*(\kappa)$, we therefore have

$$\int dF(\alpha)\, e^{-rT^*(\alpha,\kappa)}\,(\alpha\Delta\mu_P - rk_P) = R(\alpha^*(\kappa)) + (\mu_E^M - r\kappa)\Lambda(\alpha^*(\kappa)).$$

(This identity holds whether or not F has an atom at $\alpha^*(\kappa)$, as the integrand of the integral appearing in $R(\alpha^*(\kappa))$ vanishes at $\alpha = \alpha^*(\kappa)$.) Hence Ω may be written

$$\Omega(\kappa) \equiv \frac{G'(\kappa)}{G(\kappa)} \left(\mu_P^M + R(\alpha^*(\kappa)) + (\mu_E^M - r\kappa)\Lambda(\alpha^*(\kappa))\right) - r\Lambda(\alpha^*(\kappa)).$$

Since $G'(\kappa)/G(\kappa)$ and $\alpha^*(\kappa)$ are continuous in κ while $\Lambda(\alpha)$ is continuous in α, $\Omega(\kappa)$ is continuous in κ so long as $R(\alpha)$ is continuous in α. Using integration by parts, $R(\alpha)$ may be written

$$R(\alpha) = rk_P \lim_{\alpha' \downarrow \alpha} \left[(1 - F(\alpha'))\left(\frac{\alpha'}{\alpha} - 1\right)\right]_{\alpha'}^{\infty} + \frac{rk_P}{\alpha} \int_{\{\alpha > \alpha'\}} d\alpha' \, (1 - F(\alpha')).$$

The lower surface term trivially vanishes, while $\lim_{\alpha' \to \infty} \alpha'(1 - F(\alpha')) = 0$ given that F has a well-defined first moment, so that the upper surface term vanishes as well. Thus

$$R(\alpha) = \frac{rk_P}{\alpha} \int_{\{\alpha > \alpha'\}} d\alpha' \, (1 - F(\alpha')).$$

Since definite integrals are continuous in their limits of integration, $R(\alpha)$ is continuous, as desired.

Now, define $\kappa^* \equiv \sup\{\kappa \in [k_0, \overline{k}] \ : \ \Omega(\kappa) \geq 0\}$. By definition $\Omega(\kappa) < 0$ for all $\kappa > \kappa^*$, and additionally there must exist an increasing sequence $\kappa_n \to \kappa^*$ such that $\Omega(\kappa_n) > 0$, so that $\Omega(\kappa) > 0$ for all $\kappa < \kappa^*$ given that Ω is strictly decreasing. Hence κ^* is the unique maximizer of U_P. Further, $G'(\overline{k}) = 0$ implies that $\Omega(\overline{k}) < 0$. Meanwhile if $k_0 > \underline{k}$, then $\alpha^*(k_0) = \alpha_0$ and $\Lambda(\alpha^*(k_0)) = 0$, while if $k_0 = \underline{k}$, then $G(k_0) = 0$, so that in either case $\Omega(k_0) > 0$. Then by continuity of Ω it must be that $\kappa^* \in (k_0, \overline{k})$ and $\Omega(\kappa^*) = 0$, which is the desired first-order condition.

A.5 Proposition 1

We begin by observing that under an insider imitation ban, when only the average demand state enters each player's profit function due to the linear demand specification. As a result, κ^{**} is independent of the shape of the demand distribution. Throughout this proof, as well as the proof of Proposition 2, we will take κ^{**} to be fixed as we consider alternative shapes of F.

Note that $\widetilde{k}_0 \leq k_0$, so that it is possible for $\kappa^{**} < k_0$. In this case trivially $\kappa^* > \kappa^{**}$ regardless of the shape of F, and so in particular the result holds under the condition of the lemma statement. For the remainder of this proof, assume that $\kappa^{**} \geq k_0$.

Let

$$\Omega(\kappa) \equiv \frac{G'(\kappa)}{G(\kappa)} \left(\mu_P^M + (\mu_E^M - r\kappa)\Lambda(\alpha^*(\kappa)) + R(\alpha^*(\kappa))\right) - r\Lambda(\alpha^*(\kappa))$$

be the difference between the left- and right-hand sides of the unregulated FOC, evaluated at an arbitrary $\kappa \in [k_0, \overline{k}]$, with

$$\widetilde{\Omega}(\kappa) \equiv \frac{G'(\kappa)}{G(\kappa)} \left(\mu_P^M + (\mu_E^M - r\kappa)\bar{\Lambda}\right) - r\bar{\Lambda}$$

defined similarly with respect to the regulated FOC. Then

$$\Omega(\kappa) - \widetilde{\Omega}(\kappa) = \left(r - \frac{G'(\kappa)}{G(\kappa)}(\mu_E^M - r\kappa)\right)(\bar{\Lambda} - \Lambda(\alpha^*(\kappa))) + \frac{G'(\kappa)}{G(\kappa)} R(\alpha^*(\kappa)).$$

By assumption, $\kappa^{**} \geq k_0$, so it lies in the domain of the unregulated platform's problem. Since $\widetilde{\Omega}(\kappa^{**}) = 0$ by Lemma 1, we have

$$\Omega(\kappa^{**}) = \left(r - \frac{G'(\kappa^{**})}{G(\kappa^{**})}(\mu_E^M - r\kappa^{**})\right)(\bar{\Lambda} - \Lambda(\alpha^*(\kappa^{**}))) + \frac{G'(\kappa^{**})}{G(\kappa^{**})} R(\alpha^*(\kappa^{**})).$$

By rearranging the first-order condition appearing in Lemma 1, we obtain

$$r - \frac{G'(\kappa^{**})}{G(\kappa^{**})}(\mu_E^M - r\kappa^{**}) = \frac{G'(\kappa^{**})}{G(\kappa^{**})} \frac{\mu_P^M}{\bar{\Lambda}},$$

which when combined with the previous expression for $\Omega(\kappa^{**})$ yields

$$\Omega(\kappa^{**}) = \frac{G'(\kappa^{**})}{G(\kappa^{**})} \left(\mu_P^M \left(1 - \frac{\Lambda(\alpha^*(\kappa^{**}))}{\bar{\Lambda}}\right) + R(\alpha^*(\kappa^{**}))\right).$$

Lemma 4 establishes that if $\Omega(\kappa^{**}) > 0$, then $\kappa^* > \kappa^{**}$, and conversely if $\Omega(\kappa^{**}) < 0$ then $\kappa^* < \kappa^{**}$. Then since $R(\alpha) \geq 0$ for all α, it must be that $\Omega(\kappa^{**}) > 0$ if $\bar{\Lambda} > \Lambda(\alpha^*(\kappa))$. Note that

$$\Lambda(\alpha) = \frac{\alpha \Delta\mu_P - rk_P}{\alpha|\Delta\mu_E|} = \frac{\Delta\mu_P}{|\Delta\mu_E|} - \frac{rk_P}{\alpha|\Delta\mu_E|}$$

is strictly increasing in α, and is equal to $\bar{\Lambda}$ when $\alpha = 1$. Hence if $\alpha^*(\kappa^{**}) < 1$, we have $\Omega(\kappa^{**}) > 0$ and so $\kappa^* > \kappa^{**}$.

Now, by definition $\alpha^*(\kappa^{**})$ satisfies the zero-profit constraint

$$\mu_E^M - |\Delta\mu_E| \left(\int_{\{\alpha > \alpha^*(\kappa^{**})\}} dF(\alpha)\,\alpha + \alpha^*(\kappa^{**})\Delta F(\alpha^*(\kappa^{**}))e^{-r\overline{T}(\kappa^{**})} \right) = r\kappa^{**}.$$

But also, κ^{**} must satisfy the regulated zero-profit constraint

$$\mu_E^M - |\Delta\mu_E| e^{-r\widetilde{T}(\kappa^{**})} = r\kappa^{**}.$$

Hence

$$\int_{\{\alpha > \alpha^*(\kappa^{**})\}} dF(\alpha)\,\alpha + \alpha^*(\kappa^{**})\Delta F(\alpha^*(\kappa^{**}))e^{-r\overline{T}(\kappa^{**})} = e^{-r\widetilde{T}(\kappa^{**})},$$

where $\kappa^{**} > \widetilde{k}_0$, as established by Lemma 1, implies that $\widetilde{T}(\kappa^{**}) \in (0, \infty)$. It follows that if

$$\int_{\{\alpha \geq 1\}} dF(\alpha)\,\alpha < e^{-r\widetilde{T}(\kappa^{**})},$$

then we must have $\alpha^*(\kappa^{**}) < 1$. Letting $\varepsilon = e^{-r\widetilde{T}(\kappa^{**})} > 0$ yields the condition in the lemma statement.

A.6 Proposition 2

We first establish that if $\Phi(\alpha_0+) = \lim_{\alpha \downarrow \alpha_0} \Phi(\alpha)$ is sufficiently close to 1, then $\kappa^{**} > k_0$. Recall that by definition

$$k_0 = \max\{\underline{k}, r^{-1}(\mu_E^M + \Delta\mu_E \Phi(\alpha_0+))\}.$$

Therefore as $\Phi(\alpha_0+) \to 1$, we have

$$k_0 \to \max\{\underline{k} r^{-1}, \mu_E^D\} = \widetilde{k}_0.$$

Then since $\kappa^{**} > \widetilde{k}_0$, we have also that $\kappa^{**} > k_0$ for sufficiently large $\Phi(\alpha_0+)$.

Now, suppose that $\overline{\alpha}$ is chosen large enough that $\overline{\alpha} > \alpha_0$. Then given that Φ is a decreasing function, we have $\Phi(\alpha_0+) \geq \Phi(\overline{\alpha})$, and in the limit as $\Phi(\overline{\alpha})$ goes to 1, we must have $\Phi(\alpha_0+)$ go to 1 as well, in which case $\kappa^{**} > k_0$. Going forward, we will assume that $\overline{\alpha}$ is sufficiently large and $\Phi(\overline{\alpha})$ is sufficiently close to 1 so that $\kappa^{**} > k_0$, allowing us to evaluate Ω at κ^{**}.

By the proof of Proposition 1, we know that $\kappa^* < \kappa^{**}$ if

$$\Omega(\kappa^{**}) = \frac{G'(\kappa^{**})}{G(\kappa^{**})}\left(\mu_P^M\left(1 - \frac{\Lambda(\alpha^*(\kappa^{**}))}{\bar{\Lambda}}\right) + R(\alpha^*(\kappa^{**}))\right) < 0,$$

and that $\Lambda(\alpha)$ is strictly increasing in α and greater than $\bar{\Lambda}$ if $\alpha > 1$. Note that $R(\alpha)$ may be bounded above as

$$R(\alpha) \leq \frac{rk_P}{\alpha}\int_{\{\alpha'>\alpha\}} dF(\alpha')\,\alpha' \leq \frac{rk_P}{\alpha}\int dF(\alpha')\,\alpha' = \frac{rk_P}{\alpha}.$$

Hence

$$\Omega(\kappa^{**}) \leq \frac{G'(\kappa^{**})}{G(\kappa^{**})}\left(\mu_P^M\left(1 - \frac{\Lambda(\alpha^*(\kappa^{**}))}{\bar{\Lambda}}\right) + \frac{rk_P}{\alpha^*(\kappa^{**})}\right).$$

This bound depends on the demand distribution F only through $\alpha^*(\kappa^{**})$, and it is decreasing in $\alpha^*(\kappa^{**})$ and strictly negative for $\alpha^*(\kappa^{**})$ sufficiently large. Thus there must exist an $\overline{\alpha}$ such that if $\alpha^*(\kappa^{**}) \geq \overline{\alpha}$, then no matter the details of F, the inequality $\Omega(\kappa^{**}) < 0$ holds. Any such threshold $\overline{\alpha}$ must satisfy $\overline{\alpha} > 1$, since for $\alpha^*(\kappa^{**}) \leq 1$, we must have $\Omega(\kappa^{**}) \geq 0$.

Now, following logic similar to the proof of Proposition 1, if

$$\int_{\{\alpha \geq \overline{\alpha}\}} dF(\alpha)\,\alpha > e^{-r\widetilde{T}(\kappa^{**})},$$

then $\alpha^*(\kappa^{**}) \geq \overline{\alpha}$. Letting $\varepsilon = 1 - e^{-r\widetilde{T}(\kappa^{**})} > 0$ yields the condition in the proposition statement.

A.7 Lemma 5

Define

$$\Omega_R(\kappa) = \frac{G'(\kappa)}{G(\kappa)}\left[w^M + \Lambda_R(\alpha^*(\kappa))(\mu_E^M - r\kappa) + R(\alpha^*(\kappa))\right] - r\Lambda_R(\alpha^*(\kappa)),$$

where

$$\Lambda_R(\alpha) \equiv \frac{\alpha \Delta w - r k_P}{\alpha |\Delta \mu_E|}.$$

Arguments very similar to those used in the proof of Lemma 4 can be used to establish that Ω_R is strictly decreasing and that $\Omega_R(\kappa^{FB}) = 0$.

As in the platform's optimal problem, not all innovation cost thresholds are necessarily implementable given that the regulator would never direct the platform to imitate when imitation is welfare-decreasing. Let $\alpha_0^R \equiv r k_P / \Delta w$, and define

$$k_0^R \equiv \max\left\{ \underline{k}, \mu_E^M + \Delta\mu_E \int_{\{\alpha > \alpha_0^R\}} dF(\alpha)\, \alpha \right\}.$$

Then $\kappa^{FB} \in [k_0^R, \overline{k}]$. Note that under Assumption 5, $\alpha_0^R < \alpha_0$ and so $k_0^R \leq k_0$. Thus in particular $\kappa^* \geq k_0^R$, and so Ω_R may always be evaluated at κ^*.

The difference between $\Omega(\kappa)$ and $\Omega_R(\kappa)$ may then be written

$$\Delta\Omega(\kappa) = \frac{G'(\kappa)}{G(\kappa)} \left[\nu^M + \mu_E^M + \Delta\Lambda(\mu_E^M - r\kappa) \right] - r\Delta\Lambda,$$

where

$$\Delta\Lambda \equiv \frac{\Delta\nu + \Delta\mu_E}{|\Delta\mu_E|}.$$

Lemma 4 established that $\Omega(\kappa^*) = 0$, so that $\Delta\Omega(\kappa^*) = \Omega_R(\kappa^*)$. Then if $\Delta\Omega(\kappa^*) > 0$, it must be that $\kappa^{FB} > \kappa^*$ given that Ω_R is decreasing in κ and $\Omega_R(\kappa^{FB}) = 0$. And on the other hand if $\Delta\Omega(\kappa^*) < 0$, it must be that $\kappa^{FB} < \kappa^*$.

Note that κ^* satisfies

$$\frac{G'(\kappa^*)}{G(\kappa^*)} \left[\mu_P^M + \Lambda(\kappa^*)(\mu_E^M - r\kappa^*) + R(\alpha^*(\kappa^*)) \right] = r\Lambda(\kappa^*),$$

where $\Lambda(\kappa^*) > 0$. It must therefore be that

$$\frac{G'(\kappa^*)}{G(\kappa^*)}(\mu_E^M - r\kappa^*) < r.$$

Hence $\Delta\Omega(\kappa^*)$ is strictly decreasing in $\Delta\Lambda$, and therefore in ν^D. Further, when $\nu^D =$

$\nu^M - \Delta\mu_E = \underline{\nu}^D$, $\Delta\Lambda = 0$ and so $\Delta\Omega(\kappa^*) > 0$, while for large $\Delta\Lambda$, $\Delta\Omega(\kappa^*) < 0$. This implies existence of a threshold $\overline{\nu} > \underline{\nu}^D$ such that $\Delta\Omega(\overline{\nu}) = 0$, at which point $\kappa^{FB} = \kappa^*$.

It remains only to show that κ^{FB} is decreasing in ν^D. To see this, consider $\Omega_R(\kappa; \nu^D)$ as a function of ν^D in addition to κ, and notice that $\Omega_R(\kappa^{FB}(\nu^D); \nu^D) = 0$ implies that

$$\frac{G'(\kappa^{FB}(\nu^D))}{G(\kappa^{FB}(\nu^D))}(\mu_E^M - r\kappa^{FB}(\nu^D)) < r.$$

Hence $\Omega_R(\kappa^{FB}(\nu^D); \nu')$ is strictly decreasing in ν' so long as $\kappa^{FB}(\nu^D) \geq k_0^R(\nu')$, which is required for $\kappa^{FB}(\nu^D)$ to lie in the domain of Ω_R. Note that $\alpha_0^R(\nu^D)$ and therefore $k_0^R(\nu^D)$ is decreasing in ν^D, so that indeed $\kappa^{FB}(\nu^D) \geq k_0^R(\nu')$ for $\nu' > \nu^D$. It follows that $\Omega_R(\kappa^{FB}(\nu^D); \nu') < 0$ and so $\kappa^{FB}(\nu') < \kappa^{FB}(\nu^D)$ for any $\nu' > \nu^D$. This argument holds for any ν^D, yielding the claimed monotonicity.

A.8 Proposition 3

The regulation impacts welfare in two ways: it changes the threshold entry cost for E, and it changes the time and demand states at which the platform imitates. We first show that the second force always decreases welfare, holding fixed an entry threshold. We then complete the proof by deriving conditions under which the first force further decreases welfare.

Fix a threshold entry type $\kappa \geq k_0$ and consider the impact on welfare of a change from no regulation to regulation. Total welfare with and without regulation is

$$U_R(\kappa) = r^{-1}\left(w^M + \int_{\alpha^*(\kappa)}^{\overline{\alpha}} dF(\alpha)\,(\alpha\Delta w - rk_P)\right),$$

$$\widetilde{U}_R(\kappa) = r^{-1}\left(w^M + (\Delta w - rk_P)e^{-r\widetilde{T}(\kappa)}\right),$$

where $\widetilde{T}(\kappa)$ satisfies

$$\mu_E^M + \Delta\mu_E e^{-r\widetilde{T}(\kappa)} = r\kappa.$$

The discounted entry time may be equivalently written as a fictitious entry threshold $\widetilde{\alpha}(\kappa)$, conditioning on a dummy market demand state which does not enter either player's payoff function. This payoff threshold satisfies

$$\mu_E^M + \int_{\widetilde{\alpha}(\kappa)}^{\overline{\alpha}} dF(\alpha)\Delta\mu_E = r\kappa.$$

So we may write welfare under the regulation as

$$\widetilde{U}_R(\kappa) = r^{-1}\left(w^M + \int_{\widetilde{\alpha}(\kappa)}^{\overline{\alpha}} dF(\alpha)\,(\Delta w - rk_P)\right).$$

Now, note that

$$\int_{\alpha^*(\kappa)}^{\overline{\alpha}} dF(\alpha)\alpha = \int_{\widetilde{\alpha}(\kappa)}^{\overline{\alpha}} dF(\alpha) = \frac{\mu_E^M - r\kappa}{|\Delta\mu_E|}.$$

So the difference in total welfare across the two regimes is just

$$U_R(\kappa) - \widetilde{U}_R(\kappa) = rk_P \int_{\widetilde{\alpha}(\kappa)}^{\alpha^*(\kappa)} dF(\alpha).$$

Lemma A.1. *$\alpha^*(\kappa) > \widetilde{\alpha}(\kappa)$ for every $\kappa \geq k_0$.*

Proof. Define

$$\pi(\alpha) \equiv \int_{\alpha}^{\overline{\alpha}} dF(\alpha)\,\alpha, \quad \widetilde{\pi}(\alpha) \equiv \int_{\alpha}^{\overline{\alpha}} dF(\alpha).$$

Note that $\pi(\underline{\alpha}) = \widetilde{\pi}(\underline{\alpha}) = 1$ while $\pi(\overline{\alpha}) = \widetilde{\pi}(\overline{\alpha}) = 0$. Meanwhile,

$$\frac{d}{d\alpha}(\pi(\alpha) - \widetilde{\pi}(\alpha)) = F'(\alpha)(1-\alpha),$$

which is strictly positive if $\alpha < 1$ and strictly negative if $\alpha > 1$. It follows that $\pi(\alpha) > \widetilde{\pi}(\alpha)$ for every $\alpha \in (\underline{\alpha}, \overline{\alpha})$. Further, both π and $\widetilde{\pi}$ are strictly decreasing in α. Then since

$$\pi(\alpha^*(\kappa)) = \widetilde{\pi}(\widetilde{\alpha}(\kappa)) = \frac{\mu_E^M - r\kappa}{|\Delta\mu_E|}$$

and $\alpha^*(\kappa) \geq \alpha_0 > \underline{\alpha}$ whenever $\kappa \geq k_0$, it must be that $\alpha^*(\kappa) > \widetilde{\alpha}(\kappa)$. □

In light of the previous lemma, $U_R(\kappa) > \widetilde{U}_R(\kappa)$, establishing the claim that the regulation must decrease welfare, holding fixed an entry threshold.

As for the first force, suppose first that the regulation increases the threshold innovation cost. It is therefore enough to identify conditions under which the regulator would prefer less innovation absent the regulation. Lemma 5 establishes that this is so whenever ν^D is sufficiently large. Suppose alternatively that the regulation decreases the threshold innovation cost. In this case it is sufficient to identify conditions under which the regulator would prefer more innovation absent the regulation, and Lemma 5 establishes that this is so whenever ν_D is sufficiently close to $\underline{\nu}^D$.

Finally, for the third part of the proposition, consider a change from the regulated to the unregulated regime. As shown in the proof of Proposition 5,

$$U_R(\kappa^{**}) - \widetilde{U}_R(\kappa^{**}) = rk_P \int_{\widetilde{\alpha}(\kappa^{**})}^{\alpha^*(\kappa^{**})} dF(\alpha),$$

which is independent of ν^D. It is therefore sufficient to show that $U_R(\kappa^*) - U_R(\kappa^{**})$ becomes unboundedly negative as ν^D becomes large, as this implies that $U_R(\kappa^*) - \widetilde{U}_R(\kappa^{**})$ becomes negative as well.

We have

$$\begin{aligned} U_R(\kappa^*) - U_R(\kappa^{**}) = G(\kappa^*) &\left(w^M + \int_{\alpha^*(\kappa^*)}^{\overline{\alpha}} dF(\alpha)\, (\alpha \Delta w - rk_P) \right) \\ &- G(\kappa^{**}) \left(w^M + \int_{\alpha^*(\kappa^{**})}^{\overline{\alpha}} dF(\alpha)\, (\alpha \Delta w - rk_P) \right). \end{aligned}$$

This may be written as

$$U_R(\kappa^*) - U_R(\kappa^{**}) = C + \Delta w \left(G(\kappa^*) \int_{\alpha^*(\kappa^*)}^{\overline{\alpha}} dF(\alpha)\, \alpha - G(\kappa^{**}) \int_{\alpha^*(\kappa^{**})}^{\overline{\alpha}} dF(\alpha)\, \alpha \right),$$

where C is a collection of terms which does not depend on ν^D. Since Δw is positive and linearly increasing in ν^D, the desired result follows if we can show that

$$G(\kappa^*) \int_{\alpha^*(\kappa^*)}^{\overline{\alpha}} dF(\alpha)\, \alpha < G(\kappa^{**}) \int_{\alpha^*(\kappa^{**})}^{\overline{\alpha}} dF(\alpha)\, \alpha.$$

Using the fact that α^* satisfies the entrepreneur's FOC, this inequality is equivalent to

$$G(\kappa^*)(\mu_E^M - r\kappa^*) < G(\kappa^{**})(\mu_E^M - r\kappa^{**}).$$

Define $\zeta(k) \equiv G(k)(\mu_E^M - rk)$. Then

$$\zeta'(k) = G(k)\left(\frac{G'(k)}{G(k)}(\mu_E^M - rk) - r\right).$$

The second term in this expression is strictly decreasing in k. Therefore if $\zeta'(\kappa^{**}) < 0$, the same inequality holds for all $k > \kappa^{**}$. And $\zeta'(\kappa^{**}) < 0$ is equivalent to

$$\frac{G'(\kappa^{**})}{G(\kappa^{**})}(\mu_E^M - r\kappa^{**}) < r,$$

which is implied by the FOC satisfies by κ^{**}. So ζ is strictly decreasing in k above κ^{**}. The assumption that $\kappa^* > \kappa^{**}$ then implies that $\zeta(\kappa^*) < \zeta(\kappa^{**})$, proving the result.