Insider Imitation*

Erik Madsen[†] Nikhil Vellodi[‡]

[Current Draft]: September 15, 2022

Abstract

We study how regulating data usage impacts innovation in digital markets. Platforms commonly use proprietary data about third-party sellers to inform their own competing offerings, dampening incentives for innovation. We model this interaction and characterize how data usage restrictions reshape these incentives. We apply our results to show that a ban on data usage has an ambiguous impact on innovation dictated by the thickness of the right tail of demand for new products. We then generalize our analysis to more flexible regulations controlling *when* and *what* data is made available, showing how each of these additional levers should be used to improve the effectiveness of regulation. Our results contribute to an ongoing policy discussion regarding competition in digital markets.

JEL Classification: D42, D82, L42

Keywords: Innovation, digital markets, data privacy, data regulation

1 Introduction

The internet has become an important venue for economic transactions, a large fraction of which take place via dominant digital platforms. In many cases, such platforms act as both market gatekeepers and participants. For instance, Amazon offers a variety of its own products on its marketplace under brands such as "AmazonBasics and "Amazon

*We would like to thank Panle Barwick, Özlem Bedre-Defolie, Giulia Brancaccio, Olivier Compte, Jeff Ely, Andrei Hagiu, Philippe Jehiel, Justin Johnson, Alessandro Lizzeri, Leslie Marx, Kathryn Spiers, Larry White and various seminar and conference participants for insightful discussions and helpful feedback. Vellodi acknowledges support from the EUR grant ANR-17-EURE-0001.

[†]Department of Economics, New York University (emadsen@nyu.edu).

[‡]Paris School of Economics (nikhil.vellodi@psemail.eu).

Essentials"; Apple offers its own apps alongside third-party offerings via its App Store; and Meta products compete for advertising spots on Facebook and Instagram.

This blend of vertical and horizontal relations opens the door to potentially anti-competitive business practices. Existing research has highlighted the possibility of increased marketplace fees and steering of buyers toward private-label products. In this paper we focus on a less-studied practice: the use of proprietary market data to target the introduction of competing products.

This practice has been well-documented in actual marketplaces. Media and legislative investigations indicate that Amazon has routinely utilized its marketplace data in this manner.[1] On Apple's digital app store, the practice is similarly commonplace and has been termed "Sherlocking".[2] The practice has also attracted significant scrutiny from lawmakers and regulators. The European Commission has issued a preliminary view that Amazon's conduct breaches EU antitrust law (European Commission 2020), and the recently-approved Digital Markets Act explicitly bans the practice.[3] A US Congressional investigation has recommended that both Amazon and Apple's behavior be classified as anticompetitive (U.S. House Committee on the Judiciary 2020, p.16). Additionally, the UK's Competition and Markets Authority (CMA) has launched an investigation into Facebook's use of advertiser data to target competing ads for Facebook Marketplace and Facebook Dating.[4]

Is this practice harmful? On the one hand, consumers might benefit ex post through increased product diversity or lower prices; on the other, they may be harmed ex ante if targeted competition suppresses the development of innovative third-party products. The potential for such harm is particularly salient in digital markets, which are distinguished from traditional brick-and-mortar retail by low barriers to entry and the opportunity for sellers to explore niche markets. Under these conditions, targeted competition saps the gains from successful exploration and undermines digital markets as proving grounds for new products.

---

[1] See Mattioli (2020), Miranda (2018), and Soper (2016) for media reports, and U.S. House Committee on the Judiciary (2020, pp.274-282) for the findings of a US Congressional investigation.

[2] Sherlocking has been observed in, for instance, flashlight and search app categories. See `https://tinyurl.com/2tj3pea9`.

[3] See article 6(2) of the draft regulation available at `https://tinyurl.com/3drtceuu`.

[4] See `https://tinyurl.com/4ubf3s49`.

To study this issue, we develop a theoretical framework general enough to accommodate a variety of digital markets. In our model, an entrepreneur can incur a fixed cost to develop a new product (*innovate*). The entrepreneur then sells their product via a platform, who can develop a competing version of the product (*imitate*). Innovation is an inherently risky activity, as demand for the product is uncertain and can be discovered only by selling it. Once the product is introduced, the platform can observe its demand and choose if and when to imitate. We assume that the platform can commit to a policy conditioning imitation on realized demand prior to the entrepreneur's innovation decision.[5]

Our main results explore the impact of data regulation on innovation.[6] Regulation restricts the platform's access to third-party seller data, and thereby its ability to condition imitation on demand. Our methodological contribution is a first-order condition characterizing innovation under an arbitrary regulation (Theorem 1). This result reveals that data regulation impacts innovation through two channels. First, it reduces the *average* profitability of imitation by limiting the platform's ability to target imitation toward products with high demand. This average demand channel diminishes the total profit that the platform earns from innovators and encourages it to restrict innovation. Second, regulation alters the *marginal* profitability of additional imitation. The directional effect of this marginal demand channel is ambiguous and could either reinforce, partly offset, or overturn the first channel, depending on the details of the regulation as well as market conditions.[7]

We apply this general characterization to analyze specific data regulations. We first study a total ban on data usage, which prohibits "insider imitation" entirely. This regulation does not ban vertical integration outright, but instead erects a data firewall between the platform's marketplace and product divisions. It forms a natural benchmark because it is simple to implement, similar to existing legal firewalls in other contexts such as financial institutions, and appears in recent legislation as well as industry proposals for regulating online marketplaces.[8]

[5]We show that the platform can achieve credibility through a simple, realistic commitment to an overall budget for imitation. See Section 9.

[6]We use innovation as a proxy for consumer surplus in markets where innovation is an important source of consumer welfare relative to competition. See Appendix E for a formal statement of this equivalence.

[7]Our main results assume that the platform does not adjust its fees in response to data regulation. We discuss fee adjustment in Section 10.

[8] On the legislative side, the recently-adopted Digital Markets Act governing platform competition in the

We find that the effect of a data ban hinges on the thickness of the right tail of demand for new products. When the demand distribution has a thick right tail, the marginal demand channel is strongly stimulative and overwhelms the average demand channel, boosting innovation. By contrast, when the right tail of demand is thin, the two channels work together to stifle innovation. These tail conditions have a natural interpretation in terms of "experimental" and "incremental" product categories, with the former featuring products with a large chance of superstar demand while the latter generate mostly products with middling demand.

We next extend our analysis by allowing the regulator to decide *when* the platform gains access to demand data, a novel class of regulatory policies we call *data patents*. We show that a sufficiently long data patent always stimulates more innovation than does a data ban. Moreover, we identify mild regularity conditions on the demand distribution under which a properly-calibrated data patent stimulates innovation over and above the unregulated outcome.

Finally, we explore what can be achieved by a regulator with the ability to design arbitrary data usage restrictions. In this setting, the regulator decides not just when the platform gains access to data, but additionally *what* data it can observe. Implementing such restrictions would entail significant practical challenges as compared to enforcing a data ban or patent. Nonetheless, we view the unconstrained information design problem as an important theoretical benchmark tracing the outer envelope of outcomes implementable through data regulation. We show that an optimal regulation exhibits *lower censorship*, releasing data on sufficiently successful products immediately (and without garbling) while permanently protecting data on unsuccessful products.

Our work contributes to a recent literature modeling competition in vertically-integrated marketplaces. Existing work has focused on how vertical integration might impact market outcomes through changes in marketplace fees, downstream product-market competition, or consumer steering. Papers in this vein include Anderson and Bedre-Defolie (2021), Etro (2021, 2022), Hagiu, Teh, and Wright (2022), and Kang and Muir (2022). The main regu-

European Union prohibits all use of non-public seller data for competitive purposes (see footnote 3). On the industry side, Amazon has proposed to settle a pending European Commission lawsuit regarding its use of third-party seller data by pledging, among other actions, to voluntarily cease using non-public seller data for five years (https://tinyurl.com/jc8e847z).

latory question of interest in these papers is whether vertical integration should be allowed at all.

Our analysis takes vertical integration as given and studies the complementary question of whether particular data practices facilitated by vertical integration should be regulated. Miller (2008) and Jiang, Jerath, and Srinivasan (2011) show that regulation can mitigate hold-up when the platform lacks commitment power to avoid imitating successful products.[9] Our paper abstracts from hold-up and identifies circumstances under which data regulation may nonetheless improve market outcomes.[10]

Finally, an adjacent literature studies the usage and regulation of consumer data on platforms. See, for instance, Fainmesser, Galeotti, and Momot (2022), Gomes and Pavan (2022), Hidir and Vellodi (2021), Ichihashi (2019, 2022), and Kirpalani and Philippon (2021). Our work demonstrates that usage of *seller* as well as consumer data has important regulatory implications.

The remainder of the paper is structured as follows. We introduce the model in Section 2 and characterize the platform's optimal imitation policy in the absence of regulation in Section 3. Section 4 defines data privacy regulations, and in Section 5 we characterize how the platform's optimal imitation policy responds to regulation. In Section 6, study the impact of a ban on data usage. We then extend our analysis to data patents in Section 7 and to general data privacy regulations in Section 8. We discuss the substantive content of the commitment assumption in Section 9. In Section 10 we discuss how platform fees might adjust in response to data regulation. We offer concluding remarks and directions for future work in Section 11. We collect all proofs of results from the main text in the Appendix.

[9] As Hagiu, Teh, and Wright (2022) demonstrate, regulation may also be helpful for overcoming hold-up even when demand is known and data usage plays no role.

[10] Our economic analysis is complemented by several recent studies of the legal implications of marketplace data usage. See, e.g., Johnson (2020) and Khan (2017).

2 Model

2.1 The Environment

Our model is comprised of two players, a platform and an entrepreneur. The platform publicly commits to an imitation policy, which we discuss further in Section 2.3.[11] The entrepreneur observes the platform's policy and chooses whether to *innovate*. If she chooses not to, the game ends and both players make zero profits. Otherwise, the entrepreneur pays a fixed innovation cost $k > 0$ to develop the product, which she then sells over an infinite, continuous time horizon on a marketplace owned by the platform. At any time, the platform can choose to *imitate* the entrepreneur's product, paying a fixed imitation cost $k_P > 0$.

We model profits in a flexible, reduced-form manner. Both players share a common discount rate $r > 0$. So long as the entrepreneur's product is the only one present in the market, she earns a flow of profits $r\Pi_E^M(\alpha) > 0$, where α is a *demand state* whose effect on profits we describe in Section 2.2.1. Meanwhile the platform enjoys a flow of profits $r\Pi_P^M(\alpha) > 0$.

Once the platform has introduced an imitating product, the entrepreneur's flow profits fall to $r\Pi_E^D(\alpha) \in (0, r\Pi_E^M(\alpha))$, while the platform's flow profits rise to $r\Pi_P^D(\alpha) > r\Pi_M^D(\alpha)$. We will refer to $\Pi_i^M(\cdot)$ and $\Pi_i^D(\cdot)$ as the *monopoly* and *duopoly profits* of each player, respectively.

Both players maximize the net present value of lifetime profits. If the entrepreneur chooses to innovate and the platform chooses to imitate at time T, the entrepreneur's lifetime profits are

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

while the platform's lifetime profits are

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

[11] In Section 9 we discuss the extent of commitment required to achieve the platform's preferred outcome.

2.2 Uncertainty

Two important forms of uncertainty arise in the model, regarding the demand for a newly introduced product and the entrepreneur's cost of innovating.

2.2.1 Demand

We assume that, ex ante, both players view the demand state α for a new product as uncertain. Prior to introduction of the entrepreneur's product, they share a common prior that $\alpha \sim F$. We assume that F is continuous and has support on a (possibly unbounded) interval $[\underline{\alpha}, \overline{\alpha}] \subset \mathbb{R}_+$.[12] 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 each player. Specifically, we will assume that profits are *linear* in demand:

$$\Pi_i^s(\alpha) = \alpha \mu_i^s$$

for each player $i = P, E$ and market structure $s = M, D$, where $\mu_i^s > 0$ are constants. Under linear profits, a demand distribution with mean $\mathbb{E}[\alpha] = A$ yields average profits $A\mu_i^s$ to each firm i in each market structure s. Without loss, we normalize average demand so that $\mathbb{E}[\alpha] = 1$ and μ_i^s represent average profits for a given player and market structure. This normalization in particular implies $\overline{\alpha} > 1 > \underline{\alpha}$.

We let $\Delta\mu_i \equiv \mu_i^D - \mu_i^M$ denote the change in a player's average profits following imitation. Given our assumptions above, these profit changes satisfy $\Delta\mu_P > 0 > \Delta\mu_E$. We further assume that the platform finds imitation profitable on average:

Assumption 1. $\Delta\mu_P > k_P$.

2.2.2 Innovation costs

We assume that, while the entrepreneur observes her innovation cost k prior to deciding whether to innovate, from the platform's perspective k is unobserved and uncertain. At the

[12] We maintain these assumptions to streamline exposition, but they are not essential for our main results.

time it formulates an imitation policy, the platform believes that $k \sim G$, where the distribution function G has support $[\mu_E^D, \mu_E^M]$ and is continuously differentiable on its support.[13] We impose the following regularity conditions on G:

Assumption 2. *The cost hazard rate $G'(k)/G(k)$ is nonincreasing on $[\mu_E^D, \mu_E^M]$, and $G'(\mu_E^M) = 0$.*

The first part of the assumption is a monotone hazard rate condition ensuring that the platform's profit function is single-peaked. It is equivalent to log-concavity of G, a property satisfied by a wide class of standard distributions. The second part of the assumption rules out the possibility of a corner solution in which the platform prefers not to imitate the entrepreneur's product at all.

2.3 Imitation policies

Prior to the entrepreneur's innovation decision, the platform publicly commits to an *imitation policy* specifying the time at which it will introduce an imitation product. This policy cannot condition on the entrepreneur's innovation cost, which is private information. However, absent regulation it can condition on the market demand state α, which is observed by the platform as soon as the entrepreneur's product is introduced.[14] In Section 4, we formally define how data regulations shape the set of allowed imitation policies.

2.4 Regulator's Objective

We evaluate data usage regulations through the lens of a regulator concerned with maximizing consumer surplus, in markets where the welfare generated by entry of new products is much larger than the additional welfare from platform competition.[15] It can be shown that

[13] The upper bound $\kappa = \mu_E^D$ on the support of G is without loss, since entrepreneurs with costs above μ_E^M do not innovate regardless of the platform's behavior. The lower bound $\kappa = \mu_E^D$ on the support is maintained to avoid corner solutions when characterizing optimal imitation. It could be relaxed to accommodate entrepreneurs with costs below μ_E^D without qualitatively changing our results.

[14] Because the cost of innovation does not impact the entrepreneur's willingness to be imitated, the platform does not benefit from offering a menu of imitation policies. It also does not benefit from more general policies which dynamically adjust when an imitation product is offered, for instance by promising to eventually withdraw the product depending on demand conditions. See Appendix A for a proof.

[15] Remarks by scholars and regulators suggest that many digital marketplaces may satisfy this condition. "Report by the Committee for the Study of Digital Platforms Market Structure and Antitrust Subcommittee" (2019, pp. 12–13) stresses that innovation is among the most important determinants of consumer surplus.

in such markets, data regulations which induce more innovation generate greater consumer surplus. (See Appendix E for a proof.) We will therefore summarize the welfare impact of a regulation by its effect on the entrepreneur's rate of innovation, and we will suppose that the regulator's objective is to maximize innovation.

2.5 Model Discussion

Our reduced-form approach to modeling profits allows us to capture a range of settings, while abstracting from the details of parametric models of competition. For instance, $\mu_P^M > 0$ might reflect commissions the platform collects from the entrepreneur's sales, while $\Delta\mu_P > 0$ might capture additional profits the platform collects from its own sales. The magnitude of the ratio $\mu_P^D/(\mu_E^D + \mu_P^D)$ could reflect consumer willingness to substitute between the two products, the extent to which the platform can direct demand to its own product (for instance by biasing search results), and other factors determining the outcome of competition between the two players. And while we refer to the markets with and without platform entry as duopoly and monopoly, respectively, total profits under each structure could embed competition from other third-party products on the platform. Finally, the platform's payoffs may reflect not just short-run profits from the current market, but more general gains from increasing the set of products available on its marketplace, for instance by attracting a larger customer base.

Our linear payoff structure allows us to study the role of demand uncertainty on innovation while holding fixed the average profitability of the market. Much of our analysis could be generalized straightforwardly to accommodate nonlinear profit functions. However, comparative statics involving the distribution of demand would become more complex. A linear structure is also consistent with a microfoundation in which α represents uncertainty about the size of the market, variable production costs exhibit constant returns to scale, and the platform charges a proportional fee on sales. (See Appendix B for details.)

The imitation cost bound imposed in Assumption 1 focuses attention on markets in

Similarly, the recently-adopted Digital Markets Act governing platform competition in the European Union emphasizes "preserving and fostering innovation" on digital platforms as a key regulatory objective (See paragraphs 4, 25, 32, 57, 59, and 107 of the whereas statement in the draft available at https://tinyurl.com/3drtceuu).

which data regulation issues are nontrivial. If it were violated, the platform would refrain from imitating entirely in the absence of demand data. In particular, a ban on data usage would maximize innovation. Our analysis can be understood as applying to product categories in which public data (for instance, sales rank and reviews) has made the platform optimistic about the average profitability of imitation.

3 Laissez-Faire Outcome

We begin our analysis by studying the platform's imitation decision absent regulation. In an unregulated environment, the platform can condition its imitation time T on the demand state α arbitrarily.[16] Since an entrepreneur's profit from innovating decreases in her innovation cost, every imitation policy induces entrepreneurs with innovation costs below some *innovation threshold* $\kappa(T)$ to enter.

The platform's design problem can therefore be analyzed in two steps. First, given a target innovation threshold $\kappa \in [\mu_E^D, \mu_E^M]$, the platform maximizes profits among all policies achieving the target. Second, the platform optimizes over the innovation threshold using the family of optimal imitation policies determined in the first step.

The following lemma solves the first step by characterizing the optimal imitation policy implementing a given innovation threshold κ.

Lemma 1. *For each $\kappa \in [\mu_E^D, \mu_E^M]$, there exists a unique*[17] *optimal imitation policy $T^*(\kappa)$ implementing κ. It satisfies*

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

where $\alpha^*(\kappa)$ is the unique solution on $[\underline{\alpha}, \overline{\alpha}]$ to

$$\mu_E^M + \Delta\mu_E \mathbb{E}[\alpha \mathbf{1}\{\alpha^*(\kappa) \geq \alpha\}] = \kappa$$

[16] If the imitation time T is a deterministic function of the demand state, it can be considered a function from $[\underline{\alpha}, \overline{\alpha}]$ to $[0, \infty]$. To allow for policies which additionally condition on the realization of a randomization device, we will generally leave the dependence of T on α implicit rather than writing $T(\alpha)$.

[17] Uniqueness is up to zero-probability sets of demand states.

and is continuous and increasing in κ.

The lemma shows that the optimal imitation policy implementing a given innovation threshold takes a simple form: The platform imitates immediately in demand states above a marginal state $\alpha^*(\kappa)$, and avoids imitating entirely in all other states. This structure is optimal because the *conversion rate* of profits

$$\Lambda(\alpha) \equiv \frac{\alpha \Delta \mu_P - k_P}{\alpha |\Delta \mu_E|},$$

which measures how efficiently the platform can convert the entrepreneur's profits into its own, is higher when demand is larger. Effectively, the platform can spread its fixed cost of imitation over a larger profit gain when demand is high, boosting the net return to expropriating the entrepreneur's profits through imitation. As a result, the platform optimally concentrates all imitation in the largest demand states, generating a threshold imitation structure.[18] The marginal demand state $\alpha^*(\kappa)$ is pinned down by the requirement that the marginal entrepreneur (with innovation cost κ) makes zero expected profits from innovating.

Using Lemma 1, the platform's problem can be reduced to the choice of an optimal innovation threshold. This optimal threshold satisfies a standard first-order condition balancing the marginal gains from stimulating innovation against the marginal losses from scaling back imitation:

Proposition 1. *The optimal innovation threshold κ^* satisfies $\kappa^* \in (\mu_E^D, \mu_E^M)$ and is the unique solution to the first-order condition*

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

where

$$\hat{\alpha}(\kappa) \equiv \frac{\mathbb{E}[\alpha \mathbf{1}\{\alpha \geq \alpha^*(\kappa)\}]}{\mathbb{E}[\mathbf{1}\{\alpha \geq \alpha^*(\kappa)\}]}.$$

In order to marginally increase innovation, the platform must slightly scale back im-

[18] This imitation profile is consistent with the empirical finding of Zhu and Liu (2018) that Amazon introduces private-label products more often in product categories with high sales ranks and prices.

itation. Lemma 1 proves that it optimally does so in the marginal demand state $\alpha^*(\kappa)$, reducing marginal profits extracted from each innovator by the amount $\Lambda(\alpha^*(\kappa))$. On the other hand, boosting innovation yields marginal profits equal to the baseline profits μ_P^M generated by each additional innovator plus net profits from imitation. Since imitation occurs precisely in states above $\alpha^*(\kappa)$, the platform's net imitation profits are $(\mu_E^M - \kappa)\Lambda(\hat{\alpha}(\kappa))$, where $\mu_E^M - \kappa$ are the entrepreneur's lost profits from imitation, $\Lambda(\hat{\alpha}(\kappa))$ is the conversion rate of those losses into platform profits, and $\hat{\alpha}(\kappa) = \mathbb{E}[\alpha \mid \alpha \geq \alpha^*(\kappa)]$ is the *average demand state* conditional on imitation occurring. Note that $\hat{\alpha}(\kappa) \geq \alpha^*(\kappa)$, and so the platform expropriates profits more efficiently on average than it does at the margin. This distinction between the average and marginal demand state is a central theme of our analysis and will determine the impact of a data regulation on innovation.

The first-order condition provides a necessary condition for optimality supposing that the optimum is interior. In our setting Assumption 2 ensures that it is additionally sufficient. Specifically, the monotone hazard rate condition on innovation costs ensures that the slope of the platform's profit function crosses zero at most once, while the boundary condition on G' ensures that the slope is negative for sufficiently large κ. Since profits are zero at $\kappa = \mu_E^D$ (in which case no entrepreneurs innovate), profits cannot be optimized at the lower boundary, ensuring a single solution to the first-order condition which characterizes the optimal innovation threshold.

4 Data Regulations

We now formally define the class of data regulations we study.

4.1 Data Privacy Policies

We allow the regulator to influence the platform's imitation decision by restricting its access to information about demand for the entrepreneur's product. Formally, the platform may enforce a *data privacy policy* $\mathcal{D}$ (or "data policy" for short) specifying a filtration $\mathbb{F}(\mathcal{D}) = (\mathcal{F}(\mathcal{D})_t)_{t\geq 0}$ of the probability space $(\mathbb{R}_+, \mathscr{L}(\mathbb{R}_+), F)$ supporting the random variable α. Under the data policy $\mathcal{D}$, the platform may condition its imitation decisions only on information about α available under $\mathbb{F}(\mathcal{D})$. Informally, a data policy specifies the flow of information made available to the platform over time.

Our notion of data privacy policies embeds two key assumptions. First, the platform cannot be forced to forget any information it has earlier been given. And second, it cannot be blocked from imitating entirely. Data privacy policies are therefore not arbitrary restrictions on the space of implementable imitation policies. Rather, they constrain only the platform's ability to condition its imitation time on α.[19]

Under a given data policy $\mathcal{D}$, any deterministic imitation policy chosen by the platform must be an $\mathbb{F}(\mathcal{D})$-stopping time. We additionally allow the platform to commit to randomized policies which condition on a randomization device independent of α. The full set of imitation policies permitted under $\mathcal{D}$ is therefore the set of stopping times with respect to the information about α permitted under $\mathcal{D}$ along with the realization of the randomization device. We will denote the set of permissible policies $\mathcal{T}(\mathcal{D})$.

Examples of possible data privacy policies include:

- The laissez-faire policy $\mathcal{D}^{LF}$, defined by $\mathcal{F}(\mathcal{D}^{LF})_t = \mathscr{L}(\mathbb{R}_+)$ for all t. Under this policy, which is equivalent to the unregulated environment studied in Section 3, the platform observes the exact value of α at all times.

- A *data ban* $\mathcal{D}^{DB}$, defined by $\mathcal{F}(\mathcal{D}^{DB})_t = \{\mathbb{R}_+, \emptyset\}$ for all t. Under this policy, the platform learns nothing about α.

- A *data patent* $\mathcal{D}^P(T^P)$ of length T^P, defined by $\mathcal{F}(\mathcal{D}^P(T^P))_t = \{\mathbb{R}_+, \emptyset\}$ for all $t < T^P$ and $\mathcal{F}(\mathcal{D}^P(T^P))_t = \mathscr{L}(\mathbb{R}_+)$ for all $t \geq T^P$. Under a data patent, the platform learns the value of α only at the patent expiration time T^P, and receives no information about demand prior to this time.

4.2 Why Regulate Data?

Data privacy policies constitute a broad, flexible category of regulations. However, as noted above, they are not all-encompassing. In particular, data policies do not nest traditional

[19] In particular, data policies cannot implement blanket patents which block imitation entirely. See Section 4.2 for a further discussion of the relationship between data regulation and patent law.

intellectual property patents, which directly restraint the platform's ability to imitate novel designs. Nonetheless, we see several compelling reasons to study data regulations.

First, as previously discussed, restrictions on platform usage of seller data have been proposed by practitioners and have recently been adopted by the European Union (see footnote 8). Since data regulations are being implemented in practice, it is important to understand their implications for market outcomes, regardless of the available alternatives.

Second, they represent an attractive alternative to strengthening traditional intellectual property protections. One reason is that data regulations are by nature targeted to reach the digital markets at issue. Conversely, extending patent protections would simultaneously impact a variety of unrelated markets. Another reason is that competition on many digital markets is inherently asymmetric: A dominant platform can bias outcomes toward even imperfect substitutes through, for instance, self-preferencing in search results and product recommendations. As a result, patent protections restraining introduction of closely related products may not effectively protect innovators in digital markets. Data regulations address this novel problem by restricting the platform's ability to target particular products for competition.

5 Regulated Outcome

We now generalize the analysis of Section 3 to account for restrictions on the platform's freedom to pursue "insider imitation". Theorem 1, the main methodological development of the paper, derives a first-order condition characterizing optimal innovation under any given data privacy policy. In later sections, we will deploy it as a tool for analyzing specific data regulations.

5.1 Cost Minimization

Unlike in the laissez-faire benchmark, the optimal imitation policy implementing a given κ cannot be explicitly derived in general.[20] Nonetheless, the *total cost* to the platform of implementing a given innovation threshold κ, measured by its expenditure of fixed imitation

[20]Regardless of the data privacy policy, any κ in the range $[\mu_E^D, \mu_E^M]$ is implementable, since the set of imitation policies which do not condition on demand spans this range.

costs, can be partly characterized. This cost turns out to be the sole feature of an optimal imitation policy needed to characterize the platform's optimal innovation threshold.

Given a data privacy policy $\mathcal{D}$ and imitation policy $T \in \mathcal{T}(\mathcal{D})$, the platform's expected profits are

$$U_P(T) = G(\kappa)\left(\mu_P^M + \Delta\mu_P \mathbb{E}\left[\alpha e^{-rT}\right] - k_P \mathbb{E}\left[e^{-rT}\right]\right).$$

Any policy which implements innovation threshold κ must additionally satisfy the marginal entrepreneur's participation constraint

$$\mu_E^M + \Delta\mu_E \mathbb{E}\left[\alpha e^{-rT}\right] = \kappa,$$

making her indifferent between innovating or not. This constraint pins down the quantity $\mathbb{E}\left[\alpha e^{-rT}\right]$, and so maximizing U_P subject to an innovation target is equivalent to minimizing $\mathbb{E}\left[e^{-rT}\right]$. An optimal imitation policy therefore solves the cost-minimization problem

$$\inf_{T \in \mathcal{T}(\mathcal{D})} \mathbb{E}\left[e^{-rT}\right] \quad \text{s.t.} \quad \mathbb{E}\left[\alpha e^{-rT}\right] = \frac{\mu_E^M - \kappa}{|\Delta\mu_E|}.$$

Letting $T^*(\kappa; \mathcal{D})$ be any optimal innovation policy, we will define

$$C(\kappa; \mathcal{D}) \equiv \mathbb{E}\left[e^{-rT^*(\kappa;\mathcal{D})}\right]$$

to be the platform's *cost function.*[21] The cost function captures the minimal frequency of imitation $\mathbb{E}[e^{-rT}]$ required for the platform to extract a given revenue through imitation. As data regulations tighten, the platform's costs increase due to its diminished ability to target imitation toward products with strong demand (which yield a higher profit conversion rate, as discussed in Section 3).

The following lemma establishes that C is a monotone, convex function of the innovation threshold lying between certain bounds. Since C is convex, its left- and right-hand derivatives exist everywhere, which we will denote by C'_- and C'_+.

Lemma 2. *Given any data privacy policy $\mathcal{D}$, the cost function $C(\cdot; \mathcal{D})$ is decreasing, convex,*

[21] In case no optimizer exists, we will define $C(\kappa; \mathcal{D})$ more generally to be the minimized value of the objective.

and continuous. It additionally satisfies $C(\mu_E^D;\mathcal{D}) = 1$, $C(\mu_E^M;\mathcal{D}) = 0$, *and*

$$C(\kappa;\mathcal{D}^{LF}) \leq C(\kappa;\mathcal{D}) \leq \frac{\mu_E^M - \kappa}{|\Delta\mu_E|}$$

for all κ.

Monotonicity of the cost function reflects the fact that higher innovation thresholds require less imitation to implement, while convexity reflects the platform's ability to use randomized imitation policies to implement intermediate thresholds. Both the lower and upper bounds on the cost function are tight: The lower bound is achieved under the laissez-faire regime, while (as we will show in Section 6) the upper bound is achieved under a complete ban on data usage.

5.2 Optimal Innovation

Using the cost function, we can generalize the analysis of Section 3 to accommodate the restrictions imposed by a data privacy policy. The platform's optimal profits from implementing a given innovation threshold are

$$U_P(\kappa;\mathcal{D}) = G(\kappa)\left(\mu_P^M + \Delta\mu_P\frac{\mu_E^M - \kappa}{|\Delta\mu_E|} - k_P C(\kappa;\mathcal{D})\right).$$

The following theorem establishes that the unique maximizer of this function satisfies a first-order condition analogous to the one reported in Proposition 1, under an appropriate generalization of the concepts of average and marginal imitated demand. It is the main methodological development of our analysis.

Theorem 1. *Given a data privacy policy* $\mathcal{D}$, *the platform's optimal innovation threshold* $\kappa^*(\mathcal{D})$ *satisfies* $\kappa^*(\mathcal{D}) \in (\mu_E^D, \mu_E^M)$ *and is the unique solution to the first-order condition*

$$\Lambda(\alpha_+^*(\kappa;\mathcal{D})) \geq \frac{G'(\kappa)}{G(\kappa)}\left(\mu_E^M + (\mu_E^M - \kappa)\Lambda(\hat{\alpha}(\kappa;\mathcal{D}))\right) \geq \Lambda(\alpha_-^*(\kappa;\mathcal{D})),$$

where

$$\hat{\alpha}(\kappa;\mathcal{D}) \equiv \frac{\mu_E^M - \kappa}{|\Delta\mu_E| C(\kappa;\mathcal{D})}, \quad \alpha_\pm^*(\kappa;\mathcal{D}) \equiv \frac{1}{\Delta\mu_E C'_\pm(\kappa;\mathcal{D})}$$

are each positive and nondecreasing, and $\hat{\alpha}(\kappa;\mathcal{D}) \geq \max\{\alpha^*_\pm(\kappa;\mathcal{D}), 1\}$ *for all* $\kappa \in (\mu_E^D, \mu_E^M)$.

As in the laissez-faire benchmark, the optimal innovation threshold is determined by a tradeoff between the marginal profits of additional imitation and the marginal losses of reduced innovation by the entrepreneur. (Also as in that benchmark, Assumption 2 ensures that the first-order condition has a unique solution which characterizes the optimum.) Both of these quantities are controlled by the platform's ability to extract profits through imitation, as dictated by the cost function and its slope.

The quantity $\hat{\alpha}(\kappa;\mathcal{D})$ functions as the *effective* average (imitated) demand, while $\alpha^*(\kappa;\mathcal{D})$ functions as the effective marginal (imitated) demand.[22] It can be shown straightforwardly that $\hat{\alpha}(\kappa;\mathcal{D}^{LF}) = \hat{\alpha}(\kappa)$ and $\alpha^*(\kappa;\mathcal{D}^{LF}) = \alpha^*(\kappa)$, where $\hat{\alpha}(\kappa)$ and $\alpha^*(\kappa)$ are as defined in Section 3. Thus the effective average and marginal demands are a generalization of the average and marginal demands defined in the laissez-faire benchmark, and the first-order condition in Theorem 1 is a generalization of its unregulated analog in Proposition 1.

The effective average demand can be more directly connected to its unregulated analog by using the marginal entrepreneur's participation constraint to rewrite it as

$$\hat{\alpha}(\kappa;\mathcal{D}) = \frac{\mathbb{E}[\alpha \exp(-rT^*(\kappa;\mathcal{D}))]}{\mathbb{E}[\exp(-rT^*(\kappa;\mathcal{D}))]}.$$

Therefore, as in the laissez-faire setting, the effective average demand is the proportion of demand appropriated by imitation, conditional on imitation occurring. However, unlike in the laissez-faire case, the effective marginal demand cannot be interpreted as the threshold demand state used to compute average imitated demand. Under a general data policy, the effective marginal demand instead captures the demand state in which imitation is reduced in order to implement a slightly higher innovation threshold.

5.3 Marginal versus Average Demand

Theorem 1 identifies two general channels shaping innovation under a data regulation. The first channel reflects the regulation's impact on the effective average demand. Recall that

[22] Since C may exhibit kinks, the marginal imitated demand may differ depending on whether the platform is contemplating increasing or decreasing imitation. At a kink, $\alpha^*_+(\kappa;\mathcal{D}) > \alpha^*_-(\kappa;\mathcal{D})$. Otherwise, the two quantities coincide. Whenever C is differentiable, we will write $\alpha^*(\kappa;\mathcal{D})$ for this common value.

Lemma 2 established the lower bound $C(\kappa; \mathcal{D}) \geq C(\kappa; \mathcal{D}^{LF})$ for the regulated cost function. It follows immediately that $\hat{\alpha}(\kappa; \mathcal{D}) \leq \hat{\alpha}(\kappa)$. That is, the effective average demand must fall under the regulation. Intuitively, regulation reduces the efficiency of imitation and make it more costly (in the form of increased expenditure on fixed imitation costs) to appropriate a given amount of profit from each entrepreneur. This channel makes admission of additional entrepreneurs less profitable and pushes toward a lower optimal innovation threshold.

In addition to diminishing effective average demand, regulating data usage also changes the effective marginal demand. Unlike the average demand channel, this marginal demand channel has an ambiguous impact on innovation. Indeed, as the following lemma establishes, under any nontrivial data regulation the directional impact of the marginal demand channel varies with κ. To identify the net effect of a data regulation, we must therefore identify its sign near $\kappa = \kappa^*(\mathcal{D})$.

Lemma 3. *Fix a data regulation $\mathcal{D}$ and a κ for which $C(\kappa; \mathcal{D}) > C(\kappa; \mathcal{D}^{LF})$. Then there exist $\kappa' < \kappa$ and $\kappa'' > \kappa$ at which C is differentiable and $\alpha^*(\kappa'; \mathcal{D}) < \alpha^*(\kappa')$ while $\alpha^*(\kappa''; \mathcal{D}) > \alpha^*(\kappa'')$.*

6 Should Data Be Available?

We now apply our general characterization of how data regulation impacts innovation to assess concrete interventions. In this section we consider the simplest possible regulation: a complete ban on the use of marketplace data. (This regulation is also of applied interest, as discussed above.) In Section 6.1 we characterize the platform's optimal imitation policy under a data ban, and in Section 6.2 we compare this outcome to the unregulated one. Our main result is a set of sufficient conditions on the right tail weights of the demand distribution under which a data ban increases or decreases innovation.

6.1 Innovation Under a Ban

Under a data ban, the platform is completely barred from insider imitation. Formally, this restriction amounts to the requirement that the platform's imitation policy be independent of α. We will denote this regulatory regime by $\mathcal{D}^{DB}$.

In this setting, the cost function $C(\kappa;\mathcal{D}^{DB})$ (as defined in Section 5) can be calculated in closed form. Indeed, when T and α are independent, the normalization $\mathbb{E}[\alpha]=1$ allows the marginal entrepreneur's participation constraint to be written

$$\mathbb{E}[e^{-rT}] = \frac{\mu_E^M - \kappa}{|\Delta\mu_E|},$$

which pins down the cost of any imitation policy satisfying the participation constraint. It can be satisfied by, for instance, a deterministic imitation time or randomization over $T=0$ and $T=\infty$. Therefore

$$C(\kappa;\mathcal{D}^{DB}) = \frac{\mu_E^M - \kappa}{|\Delta\mu_E|}.$$

This result implies in particular that a data ban generates a cost function saturating the upper bound derived in Lemma 2.

Under this cost function, the average and marginal imitated demands can be calculated straightforwardly as $\hat{\alpha}(\kappa;\mathcal{D}^{DB}) = \alpha^*(\kappa;\mathcal{D}^{DB}) = 1$. Simply put, when imitation cannot condition on demand, profits are extracted evenly across all demand states both on average and at the margin. Theorem 1 therefore implies that the optimal innovation threshold $\kappa^*(\mathcal{D}^{DB})$ is the unique solution to the first-order condition

$$\frac{G'(\kappa)}{G(\kappa)}\left(\mu_P^M + (\mu_E^M - \kappa)\Lambda(1)\right) = \Lambda(1).$$

Applying the insights developed in Section 5, the effect of a data ban can be decomposed into two channels. By lowering the average imitated demand, a ban tends to suppress innovation. Meanwhile by changing the marginal imitated demand, it may either suppress or stimulate innovation. Recall from Lemma 1 that in the absence of regulation, the marginal demand state $\alpha^*(\kappa)$ is increasing in κ and satisfies $\alpha^*(\mu_E^D)=0$ and $\alpha^*(\mu_E^M)=\infty$. Hence for sufficiently small κ, the marginal demand channel suppresses imitation, while for large κ the channel is stimulative.[23]

[23] A data ban is therefore an example of a policy where the direction of the marginal demand channel changes exactly once with κ. The lower bound on the possible number of switches derived in Lemma 3 is therefore tight.

6.2 When Does a Ban Stimulate Innovation?

To determine whether a data ban stimulates or suppresses innovation, we must identify whether the marginal demand channel is sufficiently stimulative to overcome the suppressive effect of the average demand channel. We now link the direction and magnitude of the marginal demand channel to the weight of the right tail of the demand distribution, following which we provide sufficient conditions on tail weights under which a data ban raises or lowers innovation.

We first formally classify demand distributions according to the direction and (if the effect is stimulative) magnitude of the marginal demand channel under a data ban.

Definition 1. Given $\alpha^\dagger > 1$, a demand distribution F is $\alpha^\dagger$-*experimental* if $\alpha^*(\kappa^*(\mathcal{D}^{DB})) \geq \alpha^\dagger$. It is *incremental* if $\alpha^*(\kappa^*(\mathcal{D}^{DB})) \leq 1$.

These definitions identify whether the (unregulated) marginal demand state inducing a particular level of innovation lies to the left or right of average demand (which recall we normalize to 1). The level of innovation used as a reference point is that arising under a data ban, which as we will show shortly is the right one for assessing the impact of a ban.

Under an incremental demand distribution, the marginal demand state is below the mean, corresponding to product categories in which innovation is sensitive to returns when demand is mediocre. For such distributions, the marginal demand channel works to suppress innovation. By contrast, under an $\alpha^\dagger$-experimental demand distribution, the marginal demand state lies above $\alpha^\dagger$, corresponding to product categories in which innovation is sensitive to returns when demand is strong. For such distributions, the marginal demand channel works to boost innovation, with the magnitude of the effect scaling with $\alpha^\dagger$.

The terms "experimental" and "incremental" are meant to be deliberately evocative of relevant real-world market conditions. Product categories in which innovation is sensitive to returns at demand states below the mean are naturally "incremental", in that entrepreneurs are focused mostly on their profits when their product performs passably. By contrast, categories in which innovation is sensitive to returns at demand states above the mean (potentially far above if $\alpha^\dagger$ is large) are naturally "experimental" in the sense that entrepreneurs are focused on the potential for their products to generate superstar returns.

To provide a practical method for determining whether a product category is experimental or incremental, we next establish that these properties hold whenever a sufficient portion of the total mass of demand lies in the right or left tail, respectively.

Lemma 4. *Fix all model parameters except the demand distribution F. There exists an $M \in (0,1)$ (independent of F) such that:*

- *F is incremental if and only if*

$$\int_{\underline{\alpha}}^{1} \alpha \, dF(\alpha) \geq 1 - M.$$

- *F is $\alpha^\dagger$-experimental for $\alpha^\dagger > 1$ if and only if*

$$\int_{\alpha^\dagger}^{\overline{\alpha}} \alpha \, dF(\alpha) \geq M.$$

We are now ready to state our main result regarding the impact of a data ban. We establish that a data ban reduces innovation for any incremental product category, and it conversely boosts innovation for any sufficiently experimental product category.

Proposition 2. *Fix all model parameters except the demand distribution F.*

1. *If F is incremental, then $\kappa^* > \kappa^*(\mathcal{D}^{DB})$.*
2. *There exists a demand state $\alpha^\dagger > 1$ (independent of F) such that if F is $\alpha^\dagger$-experimental, then $\kappa^* < \kappa^*(\mathcal{D}^{DB})$.*

This result echos the general observation in Section 5 that data regulation impacts innovation through a combination of marginal and average demand channels. If marginal demand at $\kappa^*(\mathcal{D}^{DB})$ in the absence of regulation is less than 1, then marginal demand rises under a ban (recall that $\alpha^*(\kappa; \mathcal{D}^{DB}) = 1$) and both channels work to suppress innovation. Hence all incremental product categories experience less innovation under a ban. By contrast, if marginal demand is sufficiently large absent regulation, the marginal demand channel overwhelms the average demand channel and innovation rises under a ban.

Note that this result is not a full characterization of when a data ban increases or decreases innovation. In particular, it is silent on the impact of a data ban for demand distributions which are only marginally experimental. This is because for such distributions, the impact of a data ban is sensitive to the dispersion of demand, for instance as measured by second-order stochastic dominance. When demand is concentrated around the mean, a data ban can be shown to have an ambiguous impact on innovation even for experimental product categories. (See Appendix C for a demonstration.) The hypothesis that the demand distribution is sufficiently experimental can be informally understood as imposing a combination of enough mass in the right tail and sufficient dispersion.

Finally, while right tail weight and the traditional second-order stochastic dominance dispersion measure are not entirely independent conditions, these two measures of spread are not in general equivalent. In Appendix D we demonstrate by a series of examples that increased dispersion can either increase or decrease particular right tail weights, and vice versa. Our analysis therefore uncovers a novel summary statistic of uncertainty which is useful for predicting market outcomes.

7 When Should Data Be Available?

We next analyze a more flexible class of data regulations allowing the regulator to control not just whether the platform can access marketplace data, but additionally *when* this data becomes available. We will refer to this class of regulations as *data patents*.

In Section 7.1 we characterize the impact of a data patent on innovation, and we demonstrate that sufficiently stringent (but still limited) data patents are guaranteed to outperform a permanent data ban. In Section 7.2 we additionally show that under mild regularity conditions, a properly calibrated data patent boosts innovation over the laissez-faire level.

7.1 Innovation Under A Data Patent

A data patent $\mathcal{D}^P(T^P)$ prohibits the platform from accessing data on new products for a specified *patent length* T^P. Note that the class of data patents nests a data ban as a special case, since a data ban corresponds to a data patent with patent length $T^P = \infty$.

We first characterize how the platform optimally implements a target innovation threshold under $\mathcal{D}^P(T^P)$. A distinctive aspect of optimal imitation in this environment is that, for small κ, the platform optimally commits to a randomized imitation time.

Lemma 5. *Suppose the regulator imposes a data patent of length $T^P \in [0, \infty)$. There exists an innovation cost threshold $\bar{\kappa}(T^P) \in [\mu_E^D, \mu_E^M)$ such that:*

- *If $\kappa \geq \bar{\kappa}(T^P)$, the platform's unique optimal imitation policy is*

$$T^*(\kappa; \mathcal{D}^P(T^P)) = \begin{cases} T^P, & \alpha \geq \alpha^*(\kappa; \mathcal{D}^P(T^P)) \\ \infty, & \alpha < \alpha^*(\kappa; \mathcal{D}^P(T^P)) \end{cases}$$

  where $\alpha^*(\kappa; \mathcal{D}^P(T^P))$ is the unique solution on $[\underline{\alpha}, \overline{\alpha}]$ to

$$\mu_E^M + e^{-rT^P} \Delta\mu_E \int_{\alpha}^{\overline{\alpha}} \alpha' \, dF(\alpha') = \kappa$$

- *If $\kappa < \bar{\kappa}(T^P)$, the platform's unique optimal imitation policy is*

$$T^*(\kappa; \mathcal{D}^P(T^P)) = \begin{cases} 0 & w.p. \;\; \rho(\kappa; T^P) \\ T^*(\bar{\kappa}(T^P); \mathcal{D}^P(T^P)) & w.p. \;\; 1 - \rho(\kappa; T^P) \end{cases}$$

  where $\rho(\kappa; T^P) \in (0, 1]$ is the unique solution to

$$\rho\mu_E^D + (1 - \rho)\bar{\kappa}(T^P) = \kappa.$$

$\bar{\kappa}$ is continuous, increasing, and satisfies $\bar{\kappa}(0) = \mu_E^D$ and $\lim_{T^P \to \infty} \bar{\kappa}(T^P) = \mu_E^M$.

Randomization arises because the cost of an optimal *deterministic* imitation policy is non-convex in κ. For large κ, the platform can implement the target innovation threshold by deferring all imitation until the data patent expires and then selectively imitating in high demand states. For small κ, however, even unconditional imitation upon patent expiration induces too much innovation. In that case, the only feasible deterministic policy achieving κ involves unconditional imitation at a time prior to T^P.

This qualitative change in imitation behavior generates a concave kink in the cost of the optimal deterministic imitation policy. Lower costs can therefore be achieved for small values of κ by using randomization to convexify the cost function, as demonstrated in Figure 1. One important property of the randomization threshold $\bar{\kappa}$ is that it increases with the patent length, indicating that randomization is used for a wider range of innovation thresholds as the patent length grows.

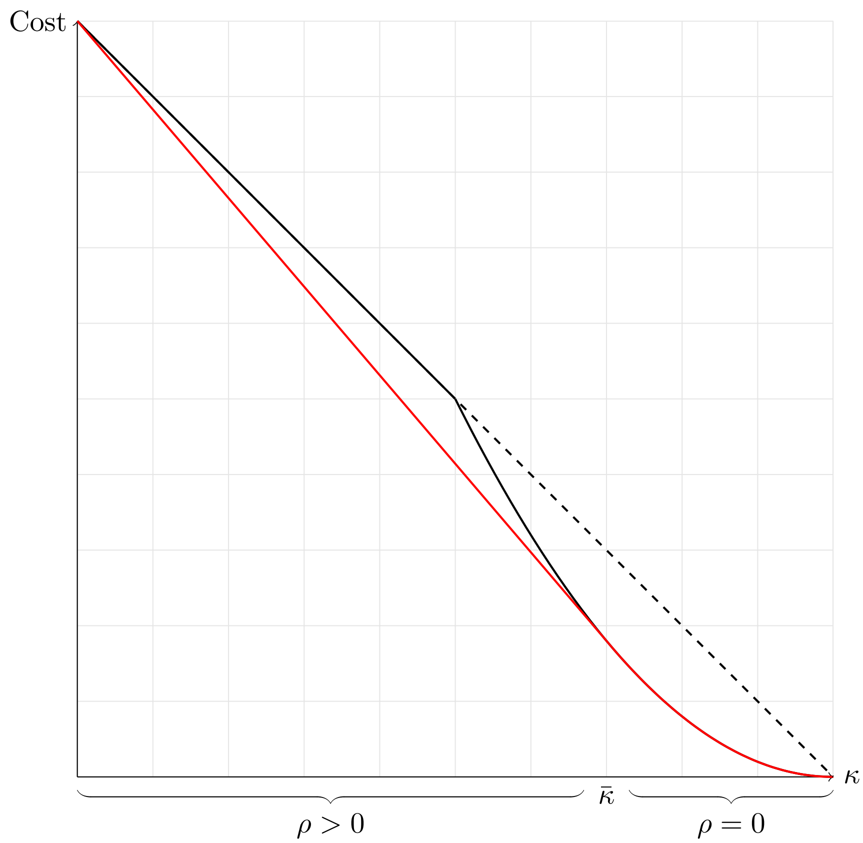

Figure 1: Constructing the cost function under a data patent. The solid black line denotes costs under the best deterministic policy, while the red line denotes the cost function.

Using the form of the platform's cost function plus the first-order condition of Theorem 1, we can partly characterize how the platform's optimal innovation threshold varies with the patent length. This characterization allows us to compare the performance of data patents and a data ban.

Proposition 3. *$\kappa^*(\mathcal{D}^P(T^P))$ is decreasing in T^P for sufficiently large T^P, and additionally*

$\lim_{T^P \to \infty} \kappa^*(\mathcal{D}^P(T^P)) = \kappa^*(\mathcal{D}^{DB})$.

This result is a consequence of the fact that the randomization regime of the platform's cost function expands to include the optimal innovation threshold when the patent is sufficiently long. For this portion of the cost function, increases in T^P lead to lower effect average demands and higher effective marginal demands, encouraging lower innovation thresholds. As a result, $\kappa^*(\mathcal{D}^P(T^P))$ is decreasing in T^P for long patents.

In the limit as $T^P \to \infty$, the cost function converges to the one arising under a data ban, implying as a consequence that $\kappa^*(\mathcal{D}^P(T^P))$ approaches $\kappa^*(\mathcal{D}^{DB})$. It can be shown that the associated optimal imitation policy, meanwhile, converges to a randomized policy which imitates immediately with probability $\rho = (\mu_E^M - \kappa^*(\mathcal{D}^{DB}))/|\Delta\mu_E|$ and refrains from imitating at all otherwise. (See Proposition F.1 in the Appendix.) This limiting policy represents one extreme of the spectrum of optimal policies under a data ban discussed in Section 5.

7.2 When Does a Data Patent Stimulate Innovation?

One important implication of Proposition 3 is that a sufficiently long data patent *always* improves on an outright data ban. However, it doesn't guarantee that data patents can improve on the unregulated outcome in settings where a data ban reduces innovation. We now show that under mild regularity conditions, sufficiently short data patents stimulate innovation.

Let $\hat{A}(\alpha') \equiv \mathbb{E}[\alpha \mid \alpha \geq \alpha']$ be average demand conditional on exceeding a threshold state α'. The quantity $\hat{A}(\alpha)$ is closely analogous to the (unregulated) average demand state $\hat{\alpha}(\kappa)$ characterized in Proposition 1, but taken as a function of the marginal demand state rather than the induced demand threshold. The two notions are linked via the identity $\hat{A}(\alpha^*(\kappa)) = \hat{\alpha}(\kappa)$.

Proposition 4. *Suppose that $\Lambda(\alpha)/\Lambda(\hat{A}(\alpha))$ is non-decreasing in α. Then $\kappa^*(\mathcal{D}^P(T^P))$ is single-peaked and is increasing in T^P for sufficiently small T^P.*

Under the regularity condition imposed in the proposition, sufficiently short data patents improve on the unregulated outcome. Further, single-peakedness of $\kappa^*(\mathcal{D}^P(T^P))$ plus the result of Proposition 3 imply that if $\kappa^*(\mathcal{D}^{DB}) \geq \kappa^*$, then *any* data patent improves innovation relative to the unregulated benchmark. Otherwise, sufficiently short data patents are stimulative, while long data patents harm innovation.

To understand this result, recall from Section 5 that, all else equal, raising the effective marginal demand induces more imitation by the platform. Increasing the length of the data patent turns out to decrease the effective marginal demand at a given innovation threshold, since delaying imitation increases the entrepreneur's profits and so imitation must occur in more demand states to compensate. This force tends to reduce imitation as the patent length increases. However, increasing the patent length also reduces effective average demand for the same reason, making exclusion of marginal entrepreneurs less costly and opposing the first force. The bound derived in Proposition 4 ensures that the first force dominates the second.

The requirement $\Lambda(\alpha)/\Lambda(\hat{A}(\alpha))$ is non-decreasing amounts to a growth rate bound on $\hat{A}(\alpha)$, which must not grow too much faster than α. We conclude our analysis by showing that this bound is satisfied by a number of common distributional families under *any* parameterization consistent with $\mathbb{E}[\alpha] = 1$.

Lemma 6. $\Lambda(\alpha)/\Lambda(\hat{A}(\alpha))$ *is non-decreasing in* α *if* F *follows a uniform, lognormal, Pareto, or gamma distribution.*

8 What Data Should Be Available?

We now give the regulator complete freedom to design a data policy and explore what can be achieved by tailoring *what* data is made available to the platform. We find that an optimal policy does not delay access to information, in contrast to the data patents studied in Section 7, but does selectively censor data on small markets.

More precisely, we show that there exists an optimal data regulation in the class of *lower censorship* policies $\mathcal{D}^C(\alpha^C)$, which immediately disclose the demand state provided it exceeds some censorship threshold α^C and otherwise provide no information about demand at any time. Formally, a lower censorship policy with censorship threshold α^C corresponds to the filtration $\mathcal{F}(\mathcal{D}^C(\alpha^C))_t = \sigma(\mathscr{L}([\alpha^C, \infty)) \cup \{[0, \alpha^C)\})$ for all t.

The following result establishes that every data privacy policy $\mathcal{D}$ can be weakly improved through lower censorship.

Lemma 7. *Fix any data privacy policy $\mathcal{D}$. Then there exists an $\alpha^C \in [\underline{\alpha}, \overline{\alpha}]$ such that $\kappa^*(\mathcal{D}^C(\alpha^C)) \geq \kappa^*(\mathcal{D})$.*

We prove this result by calibrating the lower censorship threshold α^C so that, by imitating in all demand states above α^C, the platform induces the innovation threshold $\kappa^*(\mathcal{D})$. For this choice of α^C, the platform incurs the same cost of implementing $\kappa^*(\mathcal{D})$ as under the laissez-faire policy. Since this cost is the lowest achievable under *any* data policy (recall Lemma 2), the effective average demand at $\kappa^*(\mathcal{D})$ must be weakly higher under $\mathcal{D}^C(\alpha^C)$ than under the original data policy. Simultaneously, this choice of α^C makes the cost function as steep as possible to the left of $\kappa^*(\mathcal{D})$, reducing the effective marginal demand at $\kappa^*(\mathcal{D})$ relative to the original policy. Figure 2 demonstrates this result graphically. As discussed in Section 4, raising effective average demand and lowering effective marginal demand both stimulate innovation. The optimal innovation threshold under this censorship threshold must therefore weakly exceed $\kappa^*(\mathcal{D})$.

Lemma 7 dramatically simplifies the search for an optimal data privacy policy by confining attention to the one-dimensional space of lower censorship policies. The following result characterizes the optimal censorship threshold and the corresponding maximum achievable quantity of innovation.

Proposition 5. *There exists a unique solution $\bar{\kappa}^* \in (\kappa^*, \mu_E^M)$ to the equation*

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

where

$$\underline{\alpha}^*(\kappa) \equiv \left(\Delta\mu_E \frac{C(\kappa; \mathcal{D}^{LF}) - 1}{\kappa - \mu_E^D}\right)^{-1}.$$

This solution satisfies $\bar{\kappa}^* \geq \kappa^*(\mathcal{D})$ for every data privacy policy $\mathcal{D}$, and the censorship threshold $\alpha^{C*} \equiv \alpha^*(\bar{\kappa}^*)$ satisfies $\bar{\kappa}^* = \kappa^*(\mathcal{D}^C(\alpha^{C*}))$. Under the data privacy policy $\mathcal{D}^C(\alpha^{C*})$, the platform optimally imitates immediately when $\alpha \geq \alpha^{C*}$ and never imitates otherwise.

The proof of this proposition leverages a key property of the censorship threshold char-

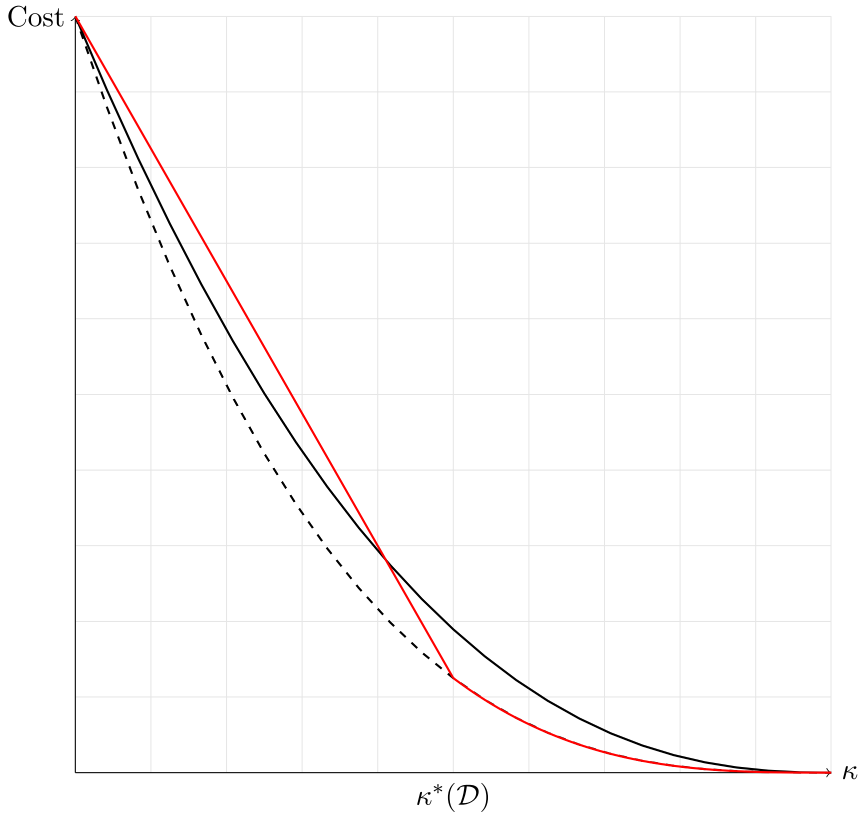

Figure 2: Constructing an improving lower censorship policy. $C(\cdot;\mathcal{D})$ is solid black, $C(\cdot;\mathcal{D}^{LF})$ is dashed black, and $C(\cdot;\mathcal{D}^C(\alpha^C))$ is red. α^C is calibrated so that $C(\cdot;\mathcal{D}^C(\alpha^C))$ is kinked at $\kappa = \kappa^*(\mathcal{D})$.

acterized in the proof of Lemma 7: It is calibrated so that the platform optimally imitates only in states at or above the threshold. As a result, attention can be restricted not just to lower censorship policies, but to those which persuade the platform not to imitate in censored demand states. This requirement implies an incentive-compatibility constraint which becomes more demanding as the censorship threshold rises and the average censored state improves. The optimal censorship threshold α^{C*} makes the platform just indifferent between imitating in censored states or not. It maximizes the amount of innovation which can be generated by any incentive-compatible lower censorship policy, and hence by any data privacy policy.

An important implication of Proposition 5 is that an optimal lower censorship policy strictly improves on the laissez-faire outcome. This result follows from the fact that, as illustrated in Figure 2, a lower censorship cost function with kink at any κ_0 is steeper to the left of κ_0 than is the unregulated cost function at κ_0. By construction, $\underline{\alpha}^*(\kappa_0)$ is the corresponding marginal demand under such a lower censorship policy. Thus $\underline{\alpha}^*(\kappa_0) < \alpha^*(\kappa_0)$ for all κ_0, and in particular $\underline{\alpha}^*(\bar{\kappa}^*) < \alpha^*(\bar{\kappa}^*)$. Comparing the first-order conditions in Propositions 1 and 5 therefore implies that $\bar{\kappa}^* > \kappa^*$. Similar reasoning can be used to conclude that $\bar{\kappa}^*$ additionally exceeds the maximum innovation possible under a data ban or any data patent.

While our results indicate that innovation can always be maximized through lower censorship, there exist other optimal policies as well. For instance, the regulator could further coarsen demand data so that the platform is informed only whether the demand state is above or below α^{C*}, without changing the resulting innovation threshold. Such a policy is effectively an imitation recommendation and is outcome-equivalent to lower censorship. We have focused on lower censorship policies in order to highlight that the binding incentive constraint involves downward deviations by the platform. Coarsening information about high demand states is not necessary for achieving the optimum, and the obedience constraint when the regulator recommends imitation is non-binding.

9 How Much Commitment Is Necessary?

Throughout our analysis, we have assumed that the platform has complete freedom to commit to an imitation policy. This assumption is consistent with the fact that dominant platforms often pursue long-term objectives that supersede short-term profit maximization. Nonetheless, our characterization of an optimal imitation policy in Theorem 1 shows that this policy leaves some money on the table ex post. Holding fixed the entrepreneur's innovation decision, the platform could generate additional profits $\Lambda(\alpha^*(\kappa^*(\mathcal{D});\mathcal{D})) > 0$ from each innovator by increasing imitation slightly in order to reduce the entrepreneur's profits by a dollar.

To implement the platform's optimal outcome, it must be able to credibly promise not to pursue any additional imitation. This promise need not entail commitment to all details of its imitation policy. Rather, it is sufficient that it commit to a *budget* for imitation, with the actual imitation policy chosen freely within that budget following the entrepreneur's innovation decision. The budget constraint ensures that the platform does not have the resources to imitate more than was initially promised.

This point can be made formally by interpreting the entrepreneur in our model as a continuum of potential innovators with heterogeneous innovation costs. In this context, consider a limited-commitment version of our model in which the platform publicly commits to an imitation budget; observes the innovation decisions of all entrepreneurs; and then chooses when and whether to imitate each innovator. If the platform commits to the budget $B = G(\kappa^*(\mathcal{D}))C(\kappa^*(\mathcal{D}))k_P$, then there exists an equilibrium of the resulting game in which entrepreneurs with costs below $\kappa^*(\mathcal{D})$ enter and the platform chooses an imitation policy which extracts revenues $\kappa^*(\mathcal{D})$ per entrepreneur on average.

To see why this outcome is an equilibrium, suppose that a mass $G(\kappa^*(\mathcal{D}))$ of entrepreneurs enter. The platform wishes to extract as much revenue from these entrepreneurs as possible subject to spending no more than $B/G(\kappa^*(\mathcal{D})) = C(\kappa^*(\mathcal{D}))k_P$ per entrepreneur on average.[24] This problem is the dual of the cost-minimization problem studied in Section 5, and so the platform can extract revenues of at most $\kappa^*(\mathcal{D})$ per entrepreneur on average. Hence, because the platform can apply the same imitation policy to all entrepreneurs (randomizing uniformly across the population in case the optimal policy involves randomization), there exists an optimal imitation decision which supports an equilibrium.

In practice, budgetary commitments are likely much easier for a platform to communicate and enforce than detailed commitments to particular imitation policies. Indeed, to the extent that actual capital budgeting in large organizations couples public commitments to divisional or project budgets with operational flexibility, the imitation policies we characterize require no novel commitment power. We therefore view our commitment assumption as likely to be satisfied in a broad range of applications.

[24]Formally, the platform should be constrained to spend no more than their budget in net present value. This formulation ensures that if the platform chooses to delay imitation, they can bank their budget and grow it at the risk-free rate in order to fund additional imitation in the future.

10 Flexible Fees

Our analysis so far has assumed that the platform's sole decision is how aggressively to imitate third-party products. In particular, we suppose that it cannot adjust the fees that it charges to third-party sellers in response to a data regulation. We now discuss the empirical validity of this assumption and generalize our analysis to accommodate fee flexibility.

Two pieces of empirical evidence suggest that, in practice, fees are unresponsive to changes in the market conditions. Thus data regulation in particular may not lead to substantial fee adjustment. First, prominent online platforms have rarely adjusted fees over time. For instance, the Apple app store has charged a uniform 30% fee on revenue throughout most of its existence.[25] Similarly, Amazon marketplace fees have varied little over time, even as it has increased private-label entry (Etro 2022). To the extent that market conditions have changed over time, this evidence suggests that platforms do not tend to adjust fees in response.

Second, platforms fees exhibit little variation across product categories. Apple essentially does not vary fees across categories at all,[26] while Amazon tailors fees across only a handful of expansive product categories (Etro 2022). To the extent that market conditions vary across product categories, this evidence further suggests that platforms do not calibrate fees in response. Additionally, this evidence suggests that even if fees can be adjusted, platforms may be constrained in practice to set a common fee level platform-wide. Since the extent of imitation is likely to vary significantly across product categories, fee adjustment in response to data regulation may be muted by the impact of such adjustment on categories involving little imitation.

Nonetheless, it is theoretically interesting to evaluate how fee-setting might respond to data regulation when such frictions do not completely rule out fee adjustment. Our model can be augmented to accommodate flexible fees by indexing the profit variables $\mu_i^s(f)$ with a fee parameter $f \in \mathbb{R}$.[27] We assume that $\mu_E^s(f)$ is differentiable and that $d\mu_E^s/df < 0$

[25] See https://tinyurl.com/2p84u2ac. One recent innovation is a small business program, instituted in 2021, which has reduced fees to 15% for small sellers.

[26] One exception is video rentals, which are exempted from fees.

[27] For simplicity, and in line with fee-setting in practice, we assume that fees do not vary with market structure. Our results would not change if such flexibility were allowed.

for all f under each market structure, reflecting the fact that fees are a transfer from the entrepreneur to the platform. Profits need not vary linearly with f if utility is imperfectly transferable. For instance, if the platform assesses fees as a flat percentage of revenue, then the entrepreneur's profits will typically vary nonlinearly with f. This framework therefore flexibly embeds a number of alternative fee-setting constraints.

Given an target innovation rate κ and fee f, let $U_P(\kappa, f)$ denote the platform's optimal profits under laissez-faire, with $U_P(\kappa, f; \mathcal{D})$ defined similarly under a data privacy policy $\mathcal{D}$. Not all (κ, f) pairs are feasible: For each κ in the support of G, there exists an imitation policy implementing κ under fee level f if and only if $(\mu_E^M(f) - \kappa)/|\Delta\mu_E(f)| \in [0, 1]$. We will say that κ is *feasible* if some such f exists. (Note that feasibility is independent of the ambient data regulation.) We will say that the fee level f is *κ-interior* for some feasible κ if $(\mu_E^M(f) - \kappa)/|\Delta\mu_E(f)| \in (0, 1)$.

In this augmented setting, data regulation leads to countervailing effects on fees and marginal demand. Intuitively, if imitation becomes more profitable on the margin, then the platform substitutes away from fees and toward imitation, and vice versa. The following proposition formalizes this result:

Proposition 6. *Fix a data privacy policy $\mathcal{D}$ and feasible innovation threshold κ. Suppose that $U_P(\kappa, f)$ and $U_P(\kappa, f; \mathcal{D})$ are quasiconcave in f and possess unique maximizers $f^*(\kappa)$ and $f^*(\kappa; \mathcal{D})$. Then:*

- *If $\alpha_+^*(\kappa, f^*(\kappa)) \leq \alpha_+^*(\kappa, f^*(\kappa); \mathcal{D})$, then $f^*(\kappa) \geq f^*(\kappa; \mathcal{D})$.*
- *If $\alpha_-^*(\kappa, f^*(\kappa)) \geq \alpha_-^*(\kappa, f^*(\kappa); \mathcal{D})$, then $f^*(\kappa) \leq f^*(\kappa; \mathcal{D})$.*

If additionally $C'(\kappa, f^*(\kappa))$ and $C'(\kappa, f^*(\kappa); \mathcal{D})$ exist and $f^*(\kappa)$ is κ-interior, then

$$\mathit{sign}(f^*(\kappa; \mathcal{D}) - f^*(\kappa)) = \mathit{sign}(\alpha^*(\kappa, f^*(\kappa)) - \alpha^*(\kappa, f^*(\kappa); \mathcal{D})).$$

The effect of fee adjustment on innovation is less straightforward, but it can be organized conceptually in a manner similar to our baseline analysis. One effect of fee flexibility is that the platform incurs lower losses from data regulation—trivially, since an expanded set of instruments can only raise profits. This *average profitability* channel makes innovation more valuable to the platform and stimulates more of it relative to the fixed-fee case.

Simultaneously, a change in fees induces a change in the marginal demand state whose direction can be pinned down with the aid of the following lemma:

Lemma 8. *Given any data privacy policy $\mathcal{D}$ and innovation threshold κ, the marginal demand states $\alpha^*_\pm(\kappa; \mathcal{D}, f)$ are nondecreasing in f.*

The direction of this *marginal profitability* channel depends on whether the regulation raises or lowers marginal demand at κ^*, holding fees fixed. If a regulation raises marginal demand, then fees drop (per Proposition 6), inducing a compensating drop in marginal demand. In this case the marginal profitability channel reduces the profitability of incremental imitation, reinforcing the average profitability channel and suggesting a rise in innovation as compared to the inflexible-fee benchmark.[28] On the other hand, if a regulation lowers marginal demand, then fees rise, generating a compensating rise in marginal demand. In this case the marginal and average profitability channels oppose one another, and the net effect of fee flexibility is ambiguous.

A full comparison of the interaction of these two effects is beyond the scope of the current paper. However, this analysis highlights how the basic forces identified in our main results remain central to understanding the effect of regulation in more general contexts.

11 Concluding Remarks

In this paper we study the welfare and regulatory implications of data usage for competition in digital markets. Our results indicate that severe restrictions on a platform's ability to access seller data (such as have been recently approved by the European Union under the Digital Markets Act) may backfire—in markets for incremental products, they may actually make the platform more aggressive and reduce innovative entry by third-party sellers. We show how to solve this problem through more tailored data-usage restrictions.

One simple, practical modification entails placing an expiration date on data protections. This tweak boosts innovation quite generally, and data patents can improve on the

[28] One subtlety is that the profit conversion rate $\Lambda(\alpha, f)$ changes with f for fixed α. Under suitable conditions, one can ensure that this effect does not offset changes in the marginal demand state. In particular, $\Lambda(\alpha, f)$ is increasing in f whenever it is positive so long as $\Delta\mu_P(f)$ and $\Delta\mu_E(f)$ are both increasing in f.

unregulated outcome in a wide variety of markets. A more ambitious change involves selectively protecting the data of products with low demand. Such protections can maximize the innovation stimulated by data regulation, although they require fine-grained analysis and control of seller data to implement.

While we have assumed that the demand for a new product is learned immediately in the absence of regulation, our framework can be readily adapted to accommodate gradual learning about demand. Formally, gradual learning can be embedded in our model by specifying a baseline filtration available to the platform in the absence of regulation, which can be further coarsened by the regulator. Analyzing the interaction of gradual learning about demand and optimal data regulation constitutes an interesting avenue for future work.

References

Anderson, S. P. and Ö. Bedre-Defolie (2021). "Hybrid Platform Model". Working paper.

Etro, F. (2021). "Product Selection in Online Marketplaces". *Journal of Economics & Management Strategy* 30, pp. 614–637.

Etro, F. (2022). "Hybrid Marketplaces with Free Entry of Sellers". 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, I. P., A. Galeotti, and R. Momot (2022). "Digital Privacy". *Management Science*, Forthcoming.

Gomes, R. and A. Pavan (2022). "Price Customization and Targeting in Matching Markets". *RAND Journal of Economics*, Forthcoming.

Hagiu, A., T.-H. Teh, and J. Wright (2022). "Should platforms be allowed to sell on their own marketplaces?" *RAND Journal of Economics* 53 (2), pp. 297–327.

Hidir, S. and N. Vellodi (2021). "Personalization, Discrimination and Information Disclosure". *Journal of the European Economic Association* 19 (2), pp. 1342–1363.

Ichihashi, S. (2019). "Online Privacy and Information Disclosure by Consumers". *American Economic Review* 110 (2), pp. 569–595.

Ichihashi, S. (2022). "Dynamic Privacy Choices". *American Economic Journal: Microeconomics*, Forthcoming.

Jiang, B., K. Jerath, and K. Srinivasan (2011). "Firm Strategies in the "Mid Tail" of Platform-Based Retailing". *Marketing Science* 30 (5), pp. 757–775.

Johnson, J. (2020). "Online Marketplaces and Vertical Integration: Prospects for Harm". *Antitrust* 35 (1), pp. 91–96.

Kang, Z. Y. and E. V. Muir (2022). "Contracting and vertical control by a dominant platform". Working paper.

Khan, L. M. (2017). "Amazon's Antitrust Paradox". *Yale Law Journal* 126 (3), pp. 710–805.

Kirpalani, R. and T. Philippon (2021). "Data Sharing and Market Power with Two-Sided Platforms". Working paper.

Mattioli, D. (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.

Miller, D. A. (2008). "Invention under uncertainty and the threat of ex post entry". *European Economic Review* 52 (3), pp. 387–412.

Miranda, L. (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.

"Report by the Committee for the Study of Digital Platforms Market Structure and Antitrust Subcommittee" (2019). *George J. Stigler Center for the Study of the Economy and the State*.

Soper, S. (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.

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.

Zhu, F. and Q. Liu (2018). "Competing with complementors: An empirical look at Amazon.com". *Strategic Management Journal* 39 (10), pp. 2618–2642.

A General Imitation Policies

In this appendix we show that it is without loss to restrict attention to imitation policies which do not withdraw an imitation product once it has been developed.

Define a *generalized imitation policy* $(X_t)_{t\geq 0}$ to be a $\{0,1\}$-valued stochastic process tracking whether the platform lists its own product at each time t. We require that X be a càdlàg process adapted to the filtration $\mathbb{F}^*(\mathcal{D})$ generated by the data privacy policy $\mathcal{D}$ and a randomization device. We call any such policy $\mathcal{D}$-feasible. We assume that flow profits for each player $i \in \{P, E\}$ are $r\mu_i^M$ when $X_t = 0$ and $r\mu_i^D$ when $X_t = 1$.[29]

Define $T(X) \equiv \inf\{t \,:\, X_t = 1\}$ to be the time at which the platform first develops an imitation product. We assume that the platform must pay the entire fixed cost k_P at time $T(X)$, following which it may freely withdraw and relist its product at no cost. Note that $T(X)$ is a stopping time with respect to $\mathbb{F}^*(\mathcal{D})$ under any $\mathcal{D}$-feasible policy X.

The following result shows that the platform can weakly improve on any feasible generalized imitation policy through a policy which never withdraws an imitation product.[30] Let $U_P(X)$ be the platform's expected profits given any data policy $\mathcal{D}$ and $\mathcal{D}$-feasible generalized imitation policy X. Since X is càdlàg, these profits are well-defined.

Proposition A.1. *Fix any data privacy policy $\mathcal{D}$ and $\mathcal{D}$-feasible generalized imitation policy X. Then there exists another $\mathcal{D}$-feasible policy X' such that $X'_t = \mathbf{1}\{t \geq T(X')\}$ for all t and $U_P(X') \geq U_P(X)$.*

Proof. Let

$$\kappa = \mu_E^M + r\Delta\mu_E \cdot \mathbb{E}\left[\int_0^\infty e^{-rt}\alpha X_t\, dt\right]$$

be the threshold innovation cost induced by the policy X. Given that $X_t = 0$ for $t < T(X)$ and $X_t \in \{0,1\}$ thereafter, there exists a unique $p \in (0,1]$ such that

$$\kappa = \mu_E^M + r\Delta\mu_E \cdot p \cdot \mathbb{E}\left[\int_{T(X)}^\infty e^{-rt}\alpha\, dt\right].$$

Define a policy X' by setting $X'_t = 0$ for all time with probability $1 - p$, independent of α, and $X'_t = \mathbf{1}\{t \geq T(X)\}$ for all t with the complementary probability. This policy is $\mathcal{D}$-feasible given that X is, and by construction $X'_t = \mathbf{1}\{t \geq T(X')\}$ for all t. Also by construction, X' induces innovation threshold κ. Further, X' imposes weakly lower imitation costs than X, since X incurs the same total

[29] Nothing would change if we allowed for more complex relationships between past competition and current payoffs, so long as the platform's flow profit gains and the entrepreneur's flow losses from imitation are weakly increasing in the total time the platform's product has been on the market.

[30] The proof can be strengthened to show that if the imitation product is not on the market "almost always" following its initial introduction, this improvement is strict.

costs as X with probability p and zero imitation costs with the complementary probability. Finally, expected imitation revenue is identical under X and X', given that

$$r\Delta\mu_P \cdot \mathbb{E}\left[\int_0^\infty e^{-rt}\alpha X_t\,dt\right] = r\Delta\mu_P \cdot p \cdot \mathbb{E}\left[\int_{T(X)}^\infty e^{-rt}\alpha\,dt\right] = r\Delta\mu_P \cdot \mathbb{E}\left[\int_0^\infty e^{-rt}\alpha X'_t\,dt\right].$$

It follows that $U_P(X') \geq U_P(X)$. □

B Microfounding Linear Profit and Surplus

In this appendix we present a simple model which microfounds the linear profit and surplus specification used throughout the paper. In the microfounded model, the entrepreneur and platform compete in prices[31], and marketplace buyers view the seller's products as differentiated goods, with total demand for each product given by $D_i(p_E, p_P, \alpha) = \alpha\bar{D}_i(p_E, p_P)$ for $i = P, E$. Each seller has a constant marginal cost c_i for producing the product, and the platform's fees are a fixed percentage t of the entrepreneur's revenues. In this model, the demand state α captures uncertainty about the number of buyers in the market, while all price elasticities are known.

We assume that conditional on the market structure, both sellers choose stage-game Nash equilibrium prices at each moment in time. We further assume that these prices are unique. If the entrepreneur is a monopolist, her stage-game profit function is

$$\pi_E^M(p_E, \alpha) = \alpha\bar{D}_E(p_E, \infty)((1-t)p_E - c_E).$$

Since this profit function is multiplicative in α, its maximum is independent of the demand state. Letting p_E^M denote the (unique, by assumption) price maximizing this profit function, flow profits under monopoly for each seller are

$$\Pi_E^M(\alpha) = \alpha\bar{D}_E(p_E^M, \infty)((1-t)p_E^M - c_E), \quad \Pi_P^M(\alpha) = \alpha t\bar{D}_E(p_E^M, \infty)p_E^M.$$

On the other hand, if the platform has entered, each seller's stage-game profit function is

$$\pi_E^D(p_E, p_P, \alpha) = \alpha\bar{D}_E(p_E, p_P)((1-t)p_E - c_E),$$
$$\pi_P^D(p_E, p_P, \alpha) = \alpha\left[\bar{D}_P(p_E, p_P)(p_P - c_P) + t\bar{D}_E(p_E, p_P)p_E\right]$$

Since these profit functions are each multiplicative in α, the set of Nash equilibria is independent of

[31] The model could alternatively be microfounded through quantity competition.

α. Letting (p_E^D, p_P^D) denote the (unique, by assumption) Nash equilibrium price vector, each seller's flow profits under duopoly are

$$\Pi_E^D(\alpha) = \alpha \bar{D}_E(p_E^D, p_P^D)((1-t)p_E^D - c_E), \quad \Pi_P^D(\alpha) = \alpha \left[\bar{D}_P(p_E^D, p_P^D)(p_P^D - c_P) + t\bar{D}_E(p_E^D, p_P^D)p_E^D\right]$$

This model therefore exhibits a linear profit structure.

Buyer demand can in turn be microfounded so that consumer surplus is also linear in the demand state. Assume that the market is made up of a continuum of atomistic consumers of measure α, each of whom has the quasilinear utility function $u(q_E, q_P, m) = v(q_E, q_P) + m$. Then the flow of consumer surplus in each state is

$$CS^M(\alpha) = \alpha \left[v(\bar{D}_E(p_E^M, \infty), 0) - \bar{D}_E(p_E^M, \infty)p_E^M\right],$$
$$CS^D(\alpha) = \alpha \left[v(\bar{D}_E(p_E^D, p_P^D), \bar{D}_P(p_E^D, p_P^D)) - \bar{D}_E(p_E^D, p_P^D)p_E^D - \bar{D}_P(p_E^D, p_P^D)p_P^D\right]$$

which is linear in α. Alternatively, the same result could be obtained from a random utility model in which each consumer purchases only one good and consumers experience iid random taste shocks to their utility for each good.

C Data Bans Under Concentrated Demand

In this appendix we show that data bans have an ambiguous impact on innovation under (marginally) experimental demand distributions when demand becomes concentrated around the mean.

We first construct a class of models featuring experimental demand in which a data ban increases innovation when demand becomes concentrated. Let F_Δ be the 2-point demand distribution with all demand concentrated at either 0 or $\overline{\alpha} = 1 + \Delta$.[32] To ensure unit demand on average, the probability of high demand must satisfy $\rho(\Delta) = 1/(1+\Delta)$.

Proposition C.1. *If $F = F_\Delta$ for any $\Delta > 0$, then $\kappa^* < \kappa^*(\mathcal{D}^{DB})$ and F is $\alpha^\dagger$-experimental for some $\alpha^\dagger > 1$.*

Proof. By construction, F_Δ is $(1+\Delta)$-experimental for all Δ. Thus it remains at least marginally experimental no matter how small Δ is taken. Further, for small Δ implementing $\kappa^*(\mathcal{D}^{DB})$ (which is interior and independent of Δ) must involve imitation only in the high demand state. Hence

[32] Formally, this class of demand distributions does not satisfy our requirements that F be continuous and supported on an interval. It could be approximated by one which does satisfy these assumptions without changing our results.

$\hat{\alpha}(\kappa^*(\mathcal{D}^{DB})) = \alpha^*(\kappa^*(\mathcal{D}^{DB})) = \overline{\alpha}$, since for small Δ. Then since $\kappa^*(\mathcal{D}^{DB})$ satisfies the FOC

$$\frac{G'(\kappa^*(\mathcal{D}^{DB}))}{G(\kappa^*(\mathcal{D}^{DB}))}(\mu_P^M + (\mu_E^M - \kappa^*(\mathcal{D}^{DB}))\Lambda(1)) = \Lambda(1),$$

we must have

$$\frac{G'(\kappa^*(\mathcal{D}^{DB}))}{G(\kappa^*(\mathcal{D}^{DB}))}(\mu_E^M - \kappa^*(\mathcal{D}^{DB})) < \Lambda(1)$$

and therefore

$$\begin{aligned}\Delta\bar{U}_P(\kappa^*(\mathcal{D}^{DB})) &= \frac{G'(\kappa^*(\mathcal{D}^{DB}))}{G(\kappa^*(\mathcal{D}^{DB}))}(\mu_P^M + (\mu_E^M - \kappa^*(\mathcal{D}^{DB}))\Lambda(\overline{\alpha})) - \Lambda(\overline{\alpha}) \\ &< \frac{G'(\kappa^*(\mathcal{D}^{DB}))}{G(\kappa^*(\mathcal{D}^{DB}))}(\mu_P^M + (\mu_E^M - \kappa^*(\mathcal{D}^{DB}))\Lambda(1) - \Lambda(1) = 0.\end{aligned}$$

Hence $\kappa^* < \kappa^*(\mathcal{D}^{DB})$ for all $\Delta > 0$. □

We next construct a class of models featuring experimental demand in which a data ban reduces innovation when demand becomes concentrated. Define $\kappa^0 \in (\mu_E^D, \mu_E^M)$ to be the unique solution to

$$\frac{G'(\kappa)}{G(\kappa)}(\mu_E^M - \kappa) = 1.$$

Note that κ^0 is independent of both μ_P^M and F. Let $\bar{F}_\Delta$ be the uniform demand distribution with support $[1 - \Delta/2, 1 + \Delta/2]$ for $\Delta \in (0, 2]$. The following proposition establishes that when κ^0 is sufficiently close to μ_E^M and μ_P^M is sufficiently small, $\bar{F}_\Delta$ is at least marginally experimental for all $\Delta > 0$ and simultaneously $\kappa^* > \kappa^*(\mathcal{D}^{DB})$ when Δ is small. Note that κ^0 may be taken close to μ_E^M by placing most of the mass of G near μ_E^M, so model parameterizations exist which satisfy these conditions.

Proposition C.2. *Fix all model parameters except for F and μ_P^M. Suppose that*

$$\frac{\mu_E^M - \kappa^0}{|\Delta\mu_E|} < 1/2.$$

Then for $\mu_P^M > 0$ sufficiently small, there exists a $\bar{\Delta} > 0$ such that if $F = \bar{F}_\Delta$ for $\Delta < \bar{\Delta}$, then $\kappa^* > \kappa^*(\mathcal{D}^{DB})$ and F is $\alpha^\dagger$-experimental for some $\alpha^\dagger > 1$.

Proof. Recall that $\kappa^* > \kappa^*(\mathcal{D}^{DB})$ if

$$\Delta\bar{U}_P(\kappa^*(\mathcal{D}^{DB})) > 0,$$

where

$$\Delta \bar{U}_P(\kappa) = \frac{G'(\kappa)}{G(\kappa)}(\mu_P^M + (\mu_E^M - \kappa)\Lambda(\hat{\alpha}(\kappa))) - \Lambda(\alpha^*(\kappa)).$$

Meanwhile, $\kappa^*(\mathcal{D}^{DB})$ satisfies the FOC

$$\frac{G'(\kappa^*(\mathcal{D}^{DB}))}{G(\kappa^*(\mathcal{D}^{DB}))}(\mu_P^M + (\mu_E^M - \kappa^*(\mathcal{D}^{DB}))\Lambda(1)) = \Lambda(1),$$

which may be equivalently written

$$\frac{G'(\kappa^*(\mathcal{D}^{DB}))}{G(\kappa^*(\mathcal{D}^{DB}))}(\mu_E^M - \kappa^*(\mathcal{D}^{DB})) = 1 - \frac{G'(\kappa^*(\mathcal{D}^{DB}))}{G(\kappa^*(\mathcal{D}^{DB}))}\Lambda(1)^{-1}\mu_P^M.$$

Inserting this expression into $\Delta U_P(\kappa^*(\mathcal{D}^{DB}))$ yields

$$\begin{aligned}\Delta \bar{U}_P(\kappa^*(\mathcal{D}^{DB})) = {} & (\Lambda(\hat{\alpha}(\kappa^*(\mathcal{D}^{DB}))) - \Lambda(1)) \\ & \times \left(\frac{\Lambda(\hat{\alpha}(\kappa^*(\mathcal{D}^{DB}))) - \Lambda(\alpha^*(\kappa^*(\mathcal{D}^{DB})))}{\Lambda(\hat{\alpha}(\kappa^*(\mathcal{D}^{DB}))) - \Lambda(1)} - \frac{G'(\kappa^*(\mathcal{D}^{DB}))}{G(\kappa^*(\mathcal{D}^{DB}))}\Lambda(1)^{-1}\mu_P^M \right).\end{aligned}$$

Since $\hat{\alpha}(\kappa) > 1$ for all $\kappa > \mu_E^D$ and Λ is an increasing function, this representation implies that $\Delta \bar{U}_P(\kappa^*(\mathcal{D}^{DB})) > 0$ if

$$\frac{\Lambda(\hat{\alpha}(\kappa^*(\mathcal{D}^{DB}))) - \Lambda(\alpha^*(\kappa^*(\mathcal{D}^{DB})))}{\Lambda(\hat{\alpha}(\kappa^*(\mathcal{D}^{DB}))) - \Lambda(1)} > \frac{G'(\kappa^*(\mathcal{D}^{DB}))}{G(\kappa^*(\mathcal{D}^{DB}))}\Lambda(1)^{-1}\mu_P^M.$$

Now, note that

$$\frac{\Lambda(\hat{\alpha}(\kappa)) - \Lambda(\alpha^*(\kappa))}{\Lambda(\hat{\alpha}(\kappa)) - \Lambda(1)} = \frac{1}{\alpha^*(\kappa)} \frac{\hat{\alpha}(\kappa) - \alpha^*(\kappa)}{\hat{\alpha}(\kappa) - 1}.$$

Let

$$M = \frac{\mu_E^M - \kappa^*(\mathcal{D}^{DB})}{|\Delta \mu_E|}$$

be as defined in Lemma 4. Then when $F = \bar{F}_\Delta$, we may explicitly compute

$$\alpha^*(\kappa^*(\mathcal{D}^{DB})) = \sqrt{(1 + \Delta/2)^2 - 2\Delta M}, \quad \hat{\alpha}(\kappa^*(\mathcal{D}^{DB})) = \frac{1 + \Delta/2 + \alpha^*(\kappa^*(\mathcal{D}^{DB}))}{2}.$$

Hence

$$\frac{\hat{\alpha}(\kappa^*(\mathcal{D}^{DB})) - \alpha^*(\kappa^*(\mathcal{D}^{DB}))}{\hat{\alpha}(\kappa^*(\mathcal{D}^{DB})) - 1} = \frac{1 + \Delta/2 - \sqrt{(1 + \Delta/2)^2 - 2\Delta M}}{\Delta/2 + \sqrt{(1 + \Delta/2)^2 - 2\Delta M} - 1}.$$

L'Hopital's rule can be used to compute

$$\lim_{\Delta \downarrow 0} \frac{\hat{\alpha}(\kappa^*(\mathcal{D}^{DB})) - \alpha^*(\kappa^*(\mathcal{D}^{DB}))}{\hat{\alpha}(\kappa^*(\mathcal{D}^{DB})) - 1} = \frac{M}{1 - M}.$$

Since trivially $\alpha^*(\kappa^*(\mathcal{D}^{DB})) \to 1$ in this limit, we therefore have

$$\lim_{\Delta \downarrow 0} \frac{\Lambda(\hat{\alpha}(\kappa^*(\mathcal{D}^{DB}))) - \Lambda(\alpha^*(\kappa^*(\mathcal{D}^{DB})))}{\Lambda(\hat{\alpha}(\kappa^*(\mathcal{D}^{DB}))) - \Lambda(1)} = \frac{M}{1-M}.$$

This work implies that $\kappa^* > \kappa^*(\mathcal{D}^{DB})$ for Δ sufficiently small if

$$\frac{M}{1-M} > \frac{G'(\kappa^*(\mathcal{D}^{DB}))}{G(\kappa^*(\mathcal{D}^{DB}))} \Lambda(1)^{-1} \mu_P^M.$$

Next, observe that as μ_P^M tends toward 0, $\kappa^*(\mathcal{D}^{DB})$ tends to κ^0. Hence the condition just derived is satisfied for small μ_P^M so long as $M^0/(1-M^0) > 0$, where

$$M^0 \equiv \frac{\mu_E^M - \kappa^0}{|\Delta \mu_E|}.$$

Since $\kappa^0 < \mu_E^M$, this condition is trivially satisfied. It follows that for μ_P^M sufficiently small, a data ban reduces innovation when $F = \bar{F}_\Delta$ and Δ is sufficiently small.

Finally, $\bar{F}_\Delta$ is incremental if and only if

$$\int_{1-\Delta/2}^{1} \alpha \, d\bar{F}_\Delta(\alpha) \geq 1 - M,$$

i.e., if

$$\frac{1}{2} + \frac{\Delta}{8} \leq M.$$

When this inequality is violated, $\bar{F}_\Delta$ must be experimental for some $\alpha^\dagger > 1$. Clearly it is violated for all $\Delta > 0$ whenever $M \leq 1/2$. Since M tends to M_0 as μ_P^M tends to 0, if $M_0 < 1/2$ then for sufficiently small μ_P^M we have $M < 1/2$, meaning $\bar{F}_\Delta$ is at least marginally experimental for all $\Delta > 0$. □

D Right Tail Weight versus SOSD

In this appendix we develop a series of simple examples to contrast the notion of right tail weight developed in Section 6 with the more familiar dispersion notion of second-order stochastic dominance, i.e., mean-preserving spreads. Given any demand distribution F and demand state $\alpha \geq 0$, define the α-tail weight $\Phi(\alpha; F)$ by

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

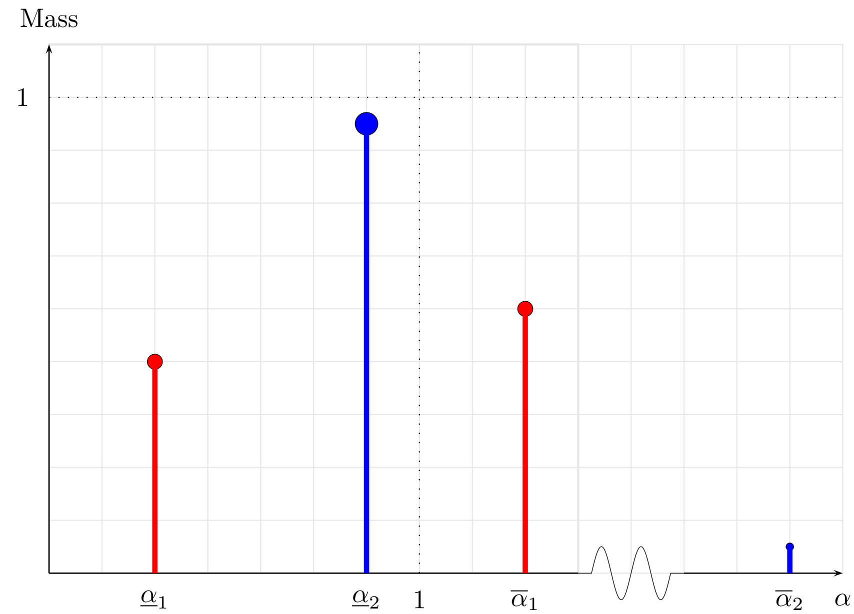

Figure 3: Increased upside demand potential doesn't imply a mean-preserving spread.

Lemma 4 says that a given demand distribution is $\alpha^\dagger$-experimental if and only if its $\alpha^\dagger$-tail weight is sufficiently large; and it is incremental if and only if its 1-tail weight is sufficiently small.

Our examples utilize the family of 2-point demand distributions, which have support on a pair of demand states $\{\overline{\alpha}, \underline{\alpha}\}$, where $\overline{\alpha} > 1 > \underline{\alpha} \geq 0$. Because we normalize average demand to 1, the locations of the high and low demand realizations fully characterize such a distribution. In particular, the probability ρ that $\alpha = \overline{\alpha}$ is uniquely pinned down by the condition $\mathbb{E}[\alpha] = 1$, yielding

$$\rho = \frac{1 - \underline{\alpha}}{\overline{\alpha} - \underline{\alpha}}.$$

We first show that a higher right-tail weight does not necessarily imply increased dispersion in the FOSD sense. Consider a pair of 2-point demand distributions F_1 and F_2 with demand realizations $\{\overline{\alpha}_1, \underline{\alpha}_1\}$ and $\{\overline{\alpha}_2, \underline{\alpha}_2\}$, where $\overline{\alpha}_2 > \overline{\alpha}_1 > \underline{\alpha}_2 > \underline{\alpha}_1$. This ordering of demand realizations implies that F_1 and F_2 cannot be ranked by second-order stochastic dominance, since the mass placed at $\underline{\alpha}_1$ by F_1 cannot be split to produce F_2, while the mass placed at $\overline{\alpha}_2$ by F_2 cannot be split to produce F_1. But since both F_1 and F_2 have mean 1, the ordering of demand realizations also implies that $\overline{\alpha}_2\rho_2 < \overline{\alpha}_1\rho_1$, where ρ_i is the probability that F_i places on the good demand realization. Therefore F_1 has a higher $\overline{\alpha}_1$-tail weight than F_2. Figure 3 graphically depicts this scenario.

We now show that increased dispersion in the second-order stochastic dominance sense need not

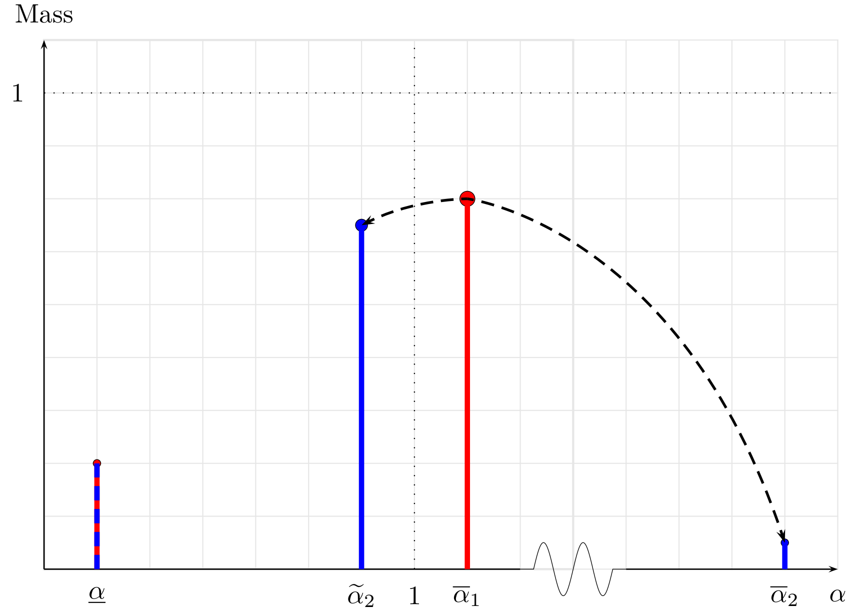

Figure 4: A mean-preserving spread doesn't imply greater upside demand potential.

imply an increased right-tail weight. Consider a 2-point demand distribution F_1 with demand realizations $\{\overline{\alpha}_1, \underline{\alpha}\}$, and define a 3-point distribution F_2 with demand realizations $\{\overline{\alpha}_2, \widetilde{\alpha}_2, \underline{\alpha}\}$ satisfying $\overline{\alpha}_2 > \overline{\alpha}_1 > 1 > \widetilde{\alpha}_2$, where F_1 and F_2 place the same probability on demand realization $\underline{\alpha}$, while the probabilities of $\widetilde{\alpha}_2$ and $\overline{\alpha}_2$ under F_2 are chosen so that F_2 is a mean-preserving spread of F_1. The mean-preserving spread relationship implies that $\overline{\alpha}_2 \rho_2 < \overline{\alpha}_1 \rho_1$, where ρ_i is the probability that F_i places on $\alpha = \overline{\alpha}_i$. Hence F_1 has a higher $\overline{\alpha}_1$-tail weight than F_2, and in fact there is no $\alpha > 1$ for which F_2 exhibits higher α-tail weight than F_1. Figure 4 graphically depicts this scenario.

E Innovation and Consumer Surplus

In this appendix, we establish conditions under which maximizing innovation can be formally viewed as the objective of a regulator concerned with maximizing consumer surplus, in the limit when innovation generates much more surplus than competition.

We model consumer surplus by assuming that consumers enjoy a flow surplus $rCS^s(\alpha)$ when the market structure is $s \in \{M, D\}$ and the demand state is α, where $CS^s(\alpha) = \alpha\nu^s$ and $\nu^D > \nu^M > 0$. If innovation does not occur, consumers receive a surplus normalized to zero. (See Appendix B for a microfoundation of this surplus structure.)

Let $U_R(\mathcal{D})$ be expected consumer surplus given a data policy $\mathcal{D}$, under the platform's corresponding optimal imitation policy. The following result establishes that, when ν^D is sufficiently close to ν^M, the directional impact on surplus of moving from one data policy to another to another tracks the impact of that change on innovation. Therefore, in the limit of small gains from competition, the regulator prefers to maximize innovation. Note that the bound on ν^D is independent of the pair of data policies being compared, and so this conclusion applies for any set of policies available to the regulator.

Proposition E.1. *Fix all model parameters except for ν^D. There exists $\bar{\nu} > \nu^M$ such that if $\nu^D < \bar{\nu}$, then given any data privacy policies $\mathcal{D}$ and $\mathcal{D}'$,*

$$sign\left(U_R(\mathcal{D}) - U_R(\mathcal{D}')\right) = sign\left(\kappa^*(\mathcal{D}) - \kappa^*(\mathcal{D}')\right).$$

Proof. We first derive an explicit expression for $U_R(\mathcal{D})$ given an arbitrary data policy $\mathcal{D}$. If an optimal imitation policy $T^*(\kappa^*(\mathcal{D}), \mathcal{D})$ exists under $\mathcal{D}$, then $U_R(\mathcal{D})$ may be written

$$U_R(\mathcal{D}) = G(\kappa^*(\mathcal{D}))\left(\nu^M + \Delta\nu \cdot \mathbb{E}\left[e^{-rT^*(\kappa^*(\mathcal{D}),\mathcal{D})}\alpha\right]\right),$$

where $\Delta\nu \equiv \nu^D - \nu^M > 0$. Additionally, $\kappa^*(\mathcal{D})$ must satisfy

$$\frac{\mu_E^M - \kappa^*(\mathcal{D})}{|\Delta\mu_E|} = \mathbb{E}\left[e^{-rT^*(\kappa^*(\mathcal{D}),\mathcal{D})}\alpha\right].$$

Hence $U_R(\mathcal{D})$ may be equivalently written $U_R(\mathcal{D}) = W(\kappa^*(\mathcal{D}))$, where

$$W(\kappa) \equiv G(\kappa)\left(\nu^M + \Delta\nu\frac{\mu_E^M - \kappa}{|\Delta\mu_E|}\right).$$

If no optimal imitation policy exists, a sequence of approximately optimal policies can be used to derive the same identity, which must therefore hold for all $\mathcal{D}$.

We next show that for ν^D sufficiently close to ν^M, the function $W(\kappa)$ is increasing on $[\mu_E^D, \bar{\kappa}^*]$, where $\bar{\kappa}^*$ is as defined in Proposition 5. The derivative of $W(\kappa)$ may be bounded below as

$$\begin{aligned}
W'(\kappa) &= G'(\kappa)\left(\nu^M + \Delta\nu\frac{\mu_E^M - \kappa}{|\Delta\mu_E|}\right) - G(\kappa)\frac{\Delta\nu}{|\Delta\mu_E|} \\
&\geq G'(\kappa)\nu^M - G(\kappa)\frac{\Delta\nu}{|\Delta\mu_E|} \\
&= G(\kappa)\left(\frac{G'(\kappa)}{G(\kappa)}\nu^M - \frac{\Delta\nu}{|\Delta\mu_E|}\right).
\end{aligned}$$

The FOC satisfied by $\bar{\kappa}^*$ implies that $G'(\bar{\kappa}^*)/G(\bar{\kappa}^*) > 0$. Then since $G'(\kappa)/G(\kappa)$ is nonincreasing on $[\mu_E^D, \bar{\kappa}^*]$ and bounded below by $G'(\bar{\kappa}^*)/G(\bar{\kappa}^*) > 0$, there exists a $\bar{\nu} > \nu^M$ such that

$$\frac{G'(\kappa)}{G(\kappa)}\nu^M - \frac{\Delta\nu}{|\Delta\mu_E|} \geq \frac{G'(\bar{\kappa}^*)}{G(\bar{\kappa}^*)}\nu^M - \frac{\Delta\nu}{|\Delta\mu_E|} > 0$$

for all $\kappa \leq \bar{\kappa}^*$ whenever $\nu^D < \bar{\nu}$. For such choices of ν^D, we have that $W'(\kappa) \geq 0$ for all $\kappa \leq \bar{\kappa}^*$, with $W'(\kappa) > 0$ when $\kappa > \mu_E^D$. Hence W is increasing on $[\mu_E^D, \bar{\kappa}^*]$ for such ν^D. This result, along with the fact every data policy induces an innovation threshold no higher than $\bar{\kappa}^*$, implies the desired result. □

F Proofs

F.1 Lemma 1

We first characterize the optimal imitation policy implementing an innovation threshold $\kappa \in (\mu_E^D, \mu_E^M)$, and return to extremal innovation thresholds afterward. When κ is interior, the platform must make the marginal entrepreneur just indifferent between innovating or not. Enforcing this participation constraint, the Lagrangian for the cost-minimization problem is

$$\mathscr{L} = \mathbb{E}[e^{-rT}] + \lambda\left(\mu_E^M + \Delta\mu_E \mathbb{E}[e^{-rT}\alpha] - \kappa\right),$$

with no constraints on the allowed choice of T aside from measurability. The Lagrangian may be equivalently written

$$\mathscr{L} = \mathbb{E}\left[e^{-rT}\left(1 - \lambda|\Delta\mu_E|\alpha\right)\right] + \lambda(\mu_E^M - \kappa).$$

If $\lambda \leq 0$, then the Lagrangian is uniquely minimized (up to measure-zero sets) by $T = \infty$. However, this policy does not satisfy the constraint given $\kappa < \mu_E^M$, and so at the optimum the Lagrange multiplier must be strictly positive. In that case the Lagrangian is uniquely minimized by

$$T(\lambda) = \begin{cases} 0, & \alpha \geq (|\Delta\mu_E|\lambda)^{-1} \\ \infty, & \alpha < (|\Delta\mu_E|\lambda)^{-1} \end{cases}$$

Note that this policy is measurable in α for any choice of λ, and so is admissible.

The marginal entrepreneur's profits are continuous and nonincreasing in λ under this family of policies, are decreasing in λ whenever profits lie in the interval $(\mu_E^D - \kappa, \mu_E^M - \kappa)$, and approach $\mu_E^M - \kappa > 0$ as $\lambda \to 0$, and approach $\mu_D^M - \kappa < 0$ as $\lambda \to \infty$. Hence there exists a unique $\lambda^*(\kappa) > 0$

satisfying the participation constraint with equality. Letting $\alpha^*(\kappa) \equiv (|\Delta\mu_E|\lambda^*(\kappa))^{-1}$, it follows that $\alpha^*(\kappa)$ uniquely satisfies the equation reported in the lemma statement. Note additionally that $T^*(\kappa) = T(\lambda^*(\kappa))$. Hence $(T^*(\kappa), \lambda^*(\kappa))$ is the unique saddle point of the Lagrangian which satisfies the zero-profit constraint, and so is the unique solution to the cost minimization problem. (No further constraint qualifications must be checked, since the objective and constraint are affine functions of the transformed choice variables $e^{-rT(\alpha)}$.)

We now return to the edge cases $\kappa \in \{\mu_E^D, \mu_E^M\}$. Since these extremal values are the boundaries of the support of the distribution of innovation costs, the marginal entrepreneur's participation constraint need only hold as an inequality. In particular, when $\kappa = \mu_E^D$ the marginal entrepreneur must make nonpositive profits, which is possible only when $T = 0$. And similarly, when $\kappa = \mu_E^M$ the marginal entrepreneur must make nonnegative profits, which is possible only when $T = \infty$. Letting $\alpha^*(\kappa) = \underline{\alpha}$ in the first case and $\alpha^*(\kappa) = \overline{\alpha}$ in the second, it follows that $T^*(\kappa)$ is trivially cost-minimizing in each case. Additionally, $\alpha^*(\kappa)$ is the unique threshold demand state on the range $[\underline{\alpha}, \overline{\alpha}]$ which satisfies the marginal entrepreneur's participation constraint with equality, as claimed.

It remains only to show that α^* is continuous and increasing in κ. Note that α^* is the inverse of the function

$$\Phi(\alpha) \equiv \mu_E^M + \Delta\mu_E \int_{\alpha}^{\overline{\alpha}} \alpha' \, dF(\alpha')$$

from $[\underline{\alpha}, \overline{\alpha}]$ to $[\mu_E^D, \mu_E^M]$. Since Φ is continuous and bijective, its inverse is also continuous. And since it is increasing, so is its inverse.

F.2 Proposition 1

This result is a corollary of Theorem 1. Lemma 1 implies that the unregulated cost function is $C(\kappa; \mathcal{D}^{LF}) = 1 - F(\alpha^*(\kappa))$, so that using the participation constraint the effective average demand may be written

$$\hat{\alpha}(\kappa; \mathcal{D}^{LF}) = \frac{\mu_E^M - \kappa}{|\Delta\mu_E|(1 - F(\alpha^*(\kappa)))} = \frac{\int_{\alpha^*(\kappa)}^{\overline{\alpha}} \alpha \, dF(\alpha)}{1 - F(\alpha^*(\kappa))} = \hat{\alpha}(\kappa).$$

As for the effective marginal demand, use the envelope theorem to compute $C'(\kappa; \mathcal{D}^{LF}) = -\lambda^*(\kappa)$, where $\lambda^*(\kappa)$ is the optimized Lagrange multiplier on the participation constraint in the cost-minimization problem. The proof of Lemma 1 established that $\lambda^*(\kappa) = (|\Delta\mu_E|\alpha^*(\kappa))^{-1}$. Hence the effective marginal demand is $\alpha^*(\kappa; \mathcal{D}^{LF}) = \alpha^*(\kappa)$. Theorem 1 therefore implies that $\kappa^* = \kappa^*(\mathcal{D}^{LF})$ satisfies the FOC stated in the proposition statement.

F.3 Lemma 2

We first establish the claimed bounds on C. Fix $\kappa \in [\mu_E^D, \mu_E^M]$ and let $\bar{t}(\kappa)$ be the unique time satisfying $e^{-r\bar{t}(\kappa)} = (\mu_E^M - \kappa)/|\Delta\mu_E|$. The imitation policy $T = \bar{t}(\kappa)$ does not condition on demand and so is feasible under any data privacy policy. It additionally satisfies the participation constraint at κ. Thus $C(\kappa; \mathcal{D}) \leq (\mu_E^M - \kappa)/|\Delta\mu_E|$. On the other hand, any imitation policy which is feasible under $\mathcal{D}$ is also feasible under $\mathcal{D}^{LF}$. Thus $C(\kappa; \mathcal{D}^{LF}) \leq C(\kappa; \mathcal{D})$.

Fix $\kappa < \mu_E^M$ and $\kappa' > \kappa$. Let $p \in [0, 1)$ be such that $p(\mu_E^M - \kappa) = \mu_E^M - \kappa'$. For every $\varepsilon > 0$ there exists an imitation policy T which satisfies $\mathbb{E}[e^{-rT}] < C(\kappa; \mathcal{D}) + \varepsilon$ and $\mathbb{E}[\alpha e^{-rT}] = (\mu_E^M - \kappa)/|\Delta\mu_E|$. Given the definition of p, the imitation policy T' which sets $T' = T$ with probability p and $T' = \infty$ otherwise satisfies $\mathbb{E}[\alpha e^{-rT'}] = (\mu_E^M - \kappa')/|\Delta\mu_E|$. Hence

$$C(\kappa') \leq \mathbb{E}[e^{-rT'}] \leq p(C(\kappa; \mathcal{D}) + \varepsilon).$$

Since this inequality holds for every ε, it follows that $C(\kappa'; \mathcal{D}) \leq pC(\kappa; \mathcal{D})$. Now, the characterization of the optimal laissez-faire policy in Lemma 1 implies that $C(\kappa; \mathcal{D}^{LF}) > 0$ given that $\kappa < \mu_E^M$. Then the bound $C(\kappa; \mathcal{D}) \geq C(\kappa; \mathcal{D}^{LF})$ established above implies that $C(\kappa; \mathcal{D}) > 0$. Combining this bound with $C(\kappa'; \mathcal{D}) \leq pC(\kappa; \mathcal{D})$ and $p < 1$ yields the relationship $C(\kappa'; \mathcal{D}) < C(\kappa; \mathcal{D})$, establishing monotonicity.

Next, fix κ, $\kappa' \neq \kappa$, and $p \in (0, 1)$. For every $\varepsilon > 0$, there exist imitation policies T and T' satisfying

$$\mathbb{E}[e^{-rT}] < C(\kappa; \mathcal{D}) + \varepsilon, \quad \mathbb{E}[e^{-rT'}] < C(\kappa'; \mathcal{D}) + \varepsilon$$

and

$$\mathbb{E}\left[\alpha e^{-rT}\right] = \frac{\mu_E^M - \kappa}{|\Delta\mu_E|}, \quad \mathbb{E}\left[\alpha e^{-rT'}\right] = \frac{\mu_E^M - \kappa'}{|\Delta\mu_E|}.$$

The imitation policy T'' which sets $T'' = T$ with probability p and $T'' = T'$ with probability $1 - p$ satisfies

$$\mathbb{E}[e^{-rT''}] < pC(\kappa; \mathcal{D}) + (1 - p)C(\kappa'; \mathcal{D}) + \varepsilon$$

and

$$\mathbb{E}\left[\alpha e^{-rT''}\right] = \frac{\mu_E^M - (p\kappa + (1 - p)\kappa')}{|\Delta\mu_E|}.$$

Hence

$$C(p\kappa + (1 - p)\kappa'; \mathcal{D}) < pC(\kappa; \mathcal{D}) + (1 - p)C(\kappa'; \mathcal{D}) + \varepsilon.$$

Since this argument holds for arbitrary $\varepsilon > 0$, it follows that

$$C(p\kappa + (1-p)\kappa'; \mathcal{D}) \leq pC(\kappa; \mathcal{D}) + (1-p)C(\kappa'; \mathcal{D}).$$

Since this argument holds for arbitrary κ, κ', and p, it follows that $C(\cdot; \mathcal{D})$ is convex.

Given convexity, continuity of C is immediate except at the endpoints $\kappa = \mu_E^D, \mu_E^M$. Note that the characterization of the optimal laissez-faire imitation policy in Lemma 1 implies that $C(\kappa; \mathcal{D}^{LF})$ is continuous everywhere and satisfies $C(\mu_E^D; \mathcal{D}^{LF}) = 1$ and $C(\mu_E^M; \mathcal{D}^{LF}) = 0$. The bounds on $C(\kappa; \mathcal{D})$ established above therefore imply that it is continuous at both endpoints and satisfies $C(\mu_E^D; \mathcal{D}) = 1$ and $C(\mu_E^M; \mathcal{D}) = 0$.

F.4 Theorem 1

Throughout this proof, we extend $\hat{\alpha}$ and α_+^* to $\kappa = \mu_E^D$ using their defining expressions, taking $\alpha_+^*(\mu_E^D; \mathcal{D}) = 0$ in case $C'_+(\mu_E^D; \mathcal{D}) = -\infty$. We similarly extend α_-^* to $\kappa = \mu_E^M$, taking $\alpha_-^*(\mu_E^D; \mathcal{D}) = \infty$ in case $C'_-(\mu_E^M; \mathcal{D}) = 0$. And given the monotonicity of $\hat{\alpha}$ established below, we will extend the definition of that function to $\kappa = \mu_E^M$ by defining $\hat{\alpha}(\mu_E^M; \mathcal{D}) \equiv \lim_{\kappa \uparrow \mu_E^M} \hat{\alpha}(\kappa; \mathcal{D})$. When we do not state the domain over which a property involving these functions holds, the implied domain is the largest one over which all functions are defined.

We begin by establishing several useful auxiliary results about $\hat{\alpha}$ and $\alpha_\pm^*$. We first note that $\alpha_+^* \geq \alpha_-^*$, since C'_- is the smallest subderivative of the convex function C at a given point while C'_+ is the largest, and so $C'_+ \geq C'_-$.

We next show that $\hat{\alpha}$ is continuous, α_+^* is right-continuous, and α_-^* is left-continuous. The continuity claim for $\hat{\alpha}$ follows from continuity of C, as established in Lemma 2. The continuity claims for $\alpha_\pm^*$ follow if we can show that C'_- is left-continuous and C'_+ is right-continuous. We prove the first result, with the second following from very similar work. Since $C'_-(\kappa)$ is a subderivative of C, $C(\kappa') - C(\kappa) \geq C'_-(\kappa)(\kappa' - \kappa)$ for all κ and κ'. Now take $\kappa \uparrow \kappa_0$ to obtain $C(\kappa') - C(\kappa_0) \geq C'_-(\kappa_0-)(\kappa' - \kappa_0)$ for all κ_0 and κ'. For $\kappa' < \kappa_0$, this inequality may be rearranged to read

$$\frac{C(\kappa_0) - C(\kappa')}{\kappa_0 - \kappa'} \leq C'_-(\kappa_0-),$$

and taking $\kappa' \uparrow \kappa_0$ implies $C'_-(\kappa_0) \leq C'_-(\kappa_0-)$ for all $\kappa_0 > \mu_E^D$. But since C'_- is nondecreasing, also $C'_-(\kappa') \leq C'_-(\kappa_0)$ for all $\kappa' < \kappa_0$, and so $C'_-(\kappa_0-) \leq C'_-(\kappa_0)$. Combining these two inequalities yields $C'_-(\kappa_0) = C'_-(\kappa_0-)$ for all $\kappa_0 > \mu_E^D$, as desired.

Finally, we show that $\alpha_-^*(\mu_E^M; \mathcal{D}) \geq 1 \geq \alpha_+^*(\mu_E^D; \mathcal{D})$. These limits follow from the fact, established in Lemma 2, that $C(\kappa;\mathcal{D})$ is bounded above by a line with slope $1/\Delta\mu_E$ everywhere and intersects that line at its endpoints. As a result, its slope at its left endpoint must be no greater than $1/\Delta\mu_E$ while its slope at its right endpoint must be no less than $1/\Delta\mu_E$. These requirements are equivalent to the stated bounds.

We now establish the properties of $\hat{\alpha}$ and $\alpha^*_\pm$ stated in the theorem statement, extended to apply over their respective maximal domains. Positivity follows from positivity of C and negativity of $C'_\pm$ for interior κ. Monotonicity of $\alpha^*_\pm$ follow from the corresponding monotonicity of $C'_\pm$ given that they are subderivatives of a convex function. Meanwhile, monotonicity of $\hat{\alpha}$ on the domain $[\mu_E^D, \mu_E^M)$ follows from the fact that

$$\frac{C(\kappa;\mathcal{D})}{\mu_E^M - \kappa} = -\frac{C(\mu_E^M;\mathcal{D}) - C(\kappa;\mathcal{D})}{\mu_E^M - \kappa}$$

is nonincreasing given convexity of C, as established in Lemma 2. Monotonicity over $[\mu_E^D, \mu_E^M]$ follows from continuity of $\hat{\alpha}$ at $\kappa = \mu_E^M$. Convexity of C implies that for all $\kappa < \mu_E^M$,

$$C'_+(\kappa;\mathcal{D}) \leq \frac{C(\mu_D^M;\mathcal{D}) - C(\kappa;\mathcal{D})}{\mu_E^M - \kappa} = -\frac{C(\kappa;\mathcal{D})}{\mu_E^M - \kappa},$$

which when rearranged is equivalent to $\hat{\alpha}(\kappa;\mathcal{D}) \geq \alpha^*_+(\kappa;\mathcal{D})$. (Since α^*_+ is not defined at $\kappa = \mu_E^M$, this inequality need not be established there.) Next, Lemma 2 established that $C(\kappa;\mathcal{D}) \leq (\mu_E^M - \kappa)|/\Delta\mu_E|$ for all κ. Hence $\hat{\alpha}(\kappa;\mathcal{D}) \geq 1$ for all $\kappa < \mu_E^M$. Continuity of $\hat{\alpha}$ at $\kappa = \mu_E^M$ then implies that $\hat{\alpha}(\mu_E^M;\mathcal{D}) \geq 1$ as well.

In the remainder of the proof, we characterize the platform's optimal innovation threshold. The platform's optimal profits among all imitation policies inducing a given innovation threshold κ are

$$U_P(\kappa;\mathcal{D}) = G(\kappa)\left(\mu_P^M + \Delta\mu_P \frac{\mu_E^M - \kappa}{|\Delta\mu_E|} - k_P C(\kappa;\mathcal{D})\right).$$

Since G and C are continuous, a maximizer must exist. Further, U_P has left- and right-hand derivatives everywhere given by $U'_{P,\pm}(\kappa;\mathcal{D}) = G(\kappa)\bar{U}'_{P,\pm}(\kappa;\mathcal{D})$, where

$$\bar{U}'_{P,\pm}(\kappa;\mathcal{D}) \equiv \frac{G'(\kappa)}{G(\kappa)}\left(\mu_P^M + (\mu_E^M - \kappa)\frac{\Delta\mu_P}{|\Delta\mu_E|} - k_P C(\kappa;\mathcal{D})\right) - \left(\frac{\Delta\mu_P}{|\Delta\mu_E|} + k_P C'_\pm(\kappa;\mathcal{D})\right).$$

Regrouping terms allows this expression to be equivalently written

$$\bar{U}'_{P,\pm}(\kappa;\mathcal{D}) \equiv \frac{G'(\kappa)}{G(\kappa)}\left(\mu_P^M + (\mu_E^M - \kappa)\Lambda(\hat{\alpha}(\kappa;\mathcal{D}))\right) - \Lambda(\alpha^*_\pm(\kappa;\mathcal{D})).$$

The system of inequalities in the theorem statement is therefore equivalent to the system $U'_{P,-}(\kappa;\mathcal{D}) \geq 0 \geq U'_{P,+}(\kappa;\mathcal{D})$, which constitute a necessary condition for optimality of an interior κ.

We first argue that $\kappa = \mu_E^D, \mu_E^M$ cannot be maximizers of U_P. Given the assumed support of G, $G(\mu_E^D) = 0$ and $G(\mu_E^M) = 1$. Additionally, Lemma Lemma 2 established that $C(\mu_E^M; \mathcal{D}) = 0$. Therefore $U_P(\mu_E^D; \mathcal{D}) = 0$ while $U_P(\mu_E^M; \mathcal{D}) = \mu_P^M > 0$. So $\kappa = \mu_E^D$ cannot be optimal. Meanwhile, Assumption 2 requires that $G'(\mu_E^M) = 0$, so that $U'_{P,-}(\mu_E^M; \mathcal{D}) = -\Lambda(\alpha_-^*(\mu_E^M; \mathcal{D}))$. Since $\alpha_-^*(\mu_E^M; \mathcal{D}) \geq 1$ and Assumption 1 implies that $\Lambda(1) > 0$, we have $U'_{P,-}(\mu_E^M; \mathcal{D}) < 0$. So $\kappa = \mu_E^M$ cannot be optimal.

To complete the proof, we establish that the system of inequalities $\bar{U}'_-(\kappa; \mathcal{D}) \geq 0 \geq \bar{U}'_+(\kappa; \mathcal{D})$ has a unique solution κ^* on (μ_E^D, μ_E^M). Since $G(\kappa) > 0$ on this range, this result implies that $U'_{P,-}(\kappa; \mathcal{D}) \geq 0 \geq U'_{P,+}(\kappa; \mathcal{D})$ has a unique solution on (μ_E^D, μ_E^M). We established earlier that a maximizer must exist and every maximizer is interior, and further any interior maximizer must satisfy $U'_{P,-}(\kappa; \mathcal{D}) \geq 0 \geq U'_{P,+}(\kappa; \mathcal{D})$. Hence κ^* must uniquely maximize U_P over $[\mu_E^D, \mu_E^M]$.

Let $\bar{\kappa} \equiv \inf\{\kappa \in (\mu_E^D, \mu_E^M) : \Lambda(\alpha_+^*(\kappa; \mathcal{D})) > 0\}$. Since $\alpha_-^*(\mu_E^M; \mathcal{D}) \geq 1$ and α_-^* is left-continuous, it must be that $\Lambda(\alpha_-^*(\kappa; \mathcal{D})) > 0$ for sufficiently large κ. The bound $\alpha_+^* \geq \alpha_-^*$ established above therefore implies $\bar{\kappa} < \mu_E^M$. For all $\kappa \in (\mu_E^D, \bar{\kappa}]$ we must have $\Lambda(\alpha_+^*(\kappa; \mathcal{D})) \leq 0$ and therefore $\bar{U}'_{P,+}(\kappa) > 0$. So any solution to the system must satisfy $\kappa > \bar{\kappa}$.

On the interval $(\bar{\kappa}, \mu_E^M]$, the expression $\phi(\kappa) \equiv -(\mu_E^M - \kappa)\Lambda(\hat{\alpha}(\kappa; \mathcal{D}))$ has right-hand derivative $\phi'_+(\kappa) = \Lambda(\alpha^* + (\kappa; \mathcal{D})) > 0$. Since ϕ may be equivalently written

$$\phi(\kappa) = -(\mu_E^M - \kappa)\frac{\Delta\mu_P}{|\Delta\mu_E|} + k_P C(\kappa; \mathcal{D}),$$

it is convex and the right-hand derivative is therefore a subderivative. It follows that ϕ is an increasing function on $(\bar{\kappa}, \mu_E^M)$. Meanwhile by assumption G'/G is a nonincreasing function which is positive for $\kappa < \mu_E^M$ given that G has support on $[\mu_E^D, \mu_E^M]$. Then since $\Lambda(\alpha_+^*(\kappa; \mathcal{D}))$ is nondecreasing, it follows that $\bar{U}'_{P,+}$ is decreasing on $(\bar{\kappa}, \mu_E^M)$.

Next consider the sign of $\bar{U}'_{P,+}(\bar{\kappa}; \mathcal{D})$. If $\bar{\kappa} > \mu_E^D$, then right-continuity of α_+^* implies that $\Lambda(\alpha_+^*(\bar{\kappa}; \mathcal{D})) = 0$ and therefore $\bar{U}'_{P,+}(\bar{\kappa}; \mathcal{D}) > 0$. Suppose instead that $\bar{\kappa} = \mu_E^D$. then $\alpha_+^*(\mu_E^D; \mathcal{D}) \leq 1$ ensures that $\Lambda(\alpha_+^*(\mu_E^D; \mathcal{D}))$ is bounded above. Meanwhile monotonicity of G'/G along with $G(\mu_E^D) = 0$ implies that $G'(\mu_E^D)/G(\mu_E^D) = \infty$, since otherwise

$$\int_{\mu_E^D}^{\kappa} \frac{G'(\kappa')}{G(\kappa')}\, d\kappa'$$

would be finite for all κ, a contradiction given that the integral is equal to $\log G(\kappa) - \log G(\mu_E^D)$, which is infinite. Hence in this case $\bar{U}'_{P,+}(\bar{\kappa}; \mathcal{D}) = \infty$. Thus in either case $\bar{U}'_{P,+}(\bar{\kappa}; \mathcal{D}) > 0$.

Now consider the sign of $\bar{U}'_{P,+}(\mu_E^M; \mathcal{D})$. By assumption $G'(\mu_E^M; \mathcal{D}) = 0$, meaning that $\bar{U}'_{P,+}(\mu_E^M; \mathcal{D}) = -\Lambda(\alpha_+^*(\mu_E^M; \mathcal{D}))$. Since $\bar{\kappa} < \mu_E^M$, we have $\Lambda(\alpha_+^*(\mu_E^M; \mathcal{D}))) > 0$ and therefore $\bar{U}'_{P,+}(\mu_E^M; \mathcal{D}) < 0$.

Now define $\kappa^* \equiv \inf\{\kappa \in [\bar{\kappa}, \mu_E^M] \ : \ U'_{P,+}(\kappa; \mathcal{D}) \le 0\}$. Since $U'_{P,+}$ is decreasing, positive at $\bar{\kappa}$, and negative at μ_E^M, it must be that $\kappa^* \in (\bar{\kappa}, \mu_E^M)$ and $U_{P,+}(\kappa; \mathcal{D}) > 0$ for all $\kappa < \kappa^*$ while $U_{P,+}(\kappa; \mathcal{D}) < 0$ for all $\kappa > \kappa^*$. Further, continuous differentiability of G, continuity of $\hat{\alpha}$, and right-continuity of α_+^* imply that $U'_{P,+}$ is right-continuous and therefore that $U'_{P,+}(\kappa^*; \mathcal{D}) \le 0$.

We finish the proof by arguing that κ^* uniquely satisfies the system $U'_{P,-}(\kappa; \mathcal{D}) \ge 0 \ge U'_{P,+}(\kappa; \mathcal{D})$ among all $\kappa \in [\bar{\kappa}, \mu_E^M]$. The fact that $U_{P,+}(\kappa; \mathcal{D}) > 0$ for all $\kappa < \kappa^*$ rules out all $\kappa < \kappa^*$ as solutions to the system. Meanwhile, the fact that $C'_\pm$ are subderivatives of a convex function implies that $C'_-(\kappa'; \mathcal{D}) \ge C'_+(\kappa; \mathcal{D})$ for all $\kappa' > \kappa$. Therefore $\alpha_-^*(\kappa'; \mathcal{D}) \ge \alpha_+^*(\kappa; \mathcal{D})$ for all $\kappa' > \kappa$ and hence $U'_{P,-}(\kappa'; \mathcal{D}) \le U'_{P,+}(\kappa; \mathcal{D})$ for all $\kappa' > \kappa$. So $U'_{P,-}(\kappa; \mathcal{D}) < \kappa$ for all $\kappa > \kappa^*$, ruling out all such κ as solutions to the system. The unique candidate solution is therefore κ^*. As observed above, this candidate satisfies $U'_{P,+}(\kappa^*; \mathcal{D}) \le 0$. Meanwhile, $\alpha_-^* \le \alpha_+^*$ implies that $U'_-(\kappa; \mathcal{D}) \ge U'_+(\kappa; \mathcal{D})) > 0$ for all $\kappa < \kappa^*$. Left-continuity of α_-^* implies the same property of U'_-, and therefore $U'_-(\kappa^*; \mathcal{D}) \ge 0$. Thus κ^* satisfies the system, as claimed.

F.5 Lemma 3

Since $C(\cdot; \mathcal{D})$ and $C(\cdot; \mathcal{D}^{LF})$ are convex, they are absolutely continuous. Write C' for their derivative wherever it exists. Then since $C(\mu_E^D; \mathcal{D}) = C(\mu_E^D; \mathcal{D}^{LF}) = 1$ while $C(\mu_E^M; \mathcal{D}) = C(\mu_E^M; \mathcal{D}^{LF}) = 0$, we must have

$$\int_0^1 C'(\kappa'; \mathcal{D})\, d\kappa' = \int_0^1 C'(\kappa'; \mathcal{D}^{LF})\, d\kappa'.$$

Now, by assumption there exists a κ such that $C(\kappa; \mathcal{D}) > C(\kappa; \mathcal{D}^{LF})$. Then

$$\int_0^\kappa C'(\kappa'; \mathcal{D})\, d\kappa' > \int_0^\kappa C'(\kappa'; \mathcal{D}^{LF})\, d\kappa',$$

implying that also

$$\int_\kappa^1 C'(\kappa'; \mathcal{D})\, d\kappa' < \int_0^\kappa C'(\kappa'; \mathcal{D}^{LF})\, d\kappa'.$$

The first inequality implies existence of a positive-measure subset of $[0, \kappa]$ on which $C'(\cdot; \mathcal{D})$ is defined and $C'(\kappa; \mathcal{D}) > C'(\kappa; \mathcal{D}^{LF})$. Meanwhile the second inequality implies existence of a positive-measure subset of $[\kappa, 1]$ on which $C'(\cdot; \mathcal{D})$ is defined and $C'(\kappa; \mathcal{D}) < C'(\kappa; \mathcal{D}^{LF})$. These facts imply the claim in the lemma.

F.6 Lemma 4

Define $\Phi(\alpha; F) \equiv \int_{\alpha}^{\overline{\alpha}} \alpha' \, dF(\alpha')$. Lemma 1 established that $\alpha^*(\kappa^*(\mathcal{D}^{DB}))$ is the unique solution on $[\underline{\alpha}, \overline{\alpha}]$ to $\Phi(\alpha; F) = M$, where

$$M \equiv \frac{\mu_E^M - \kappa^*(\mathcal{D}^{DB})}{|\Delta\mu_E|}.$$

Let $\bar{U}'_{P,\pm}(\kappa; \mathcal{D})$ be as defined in the proof of Theorem 1, and let $\bar{U}'_P(\kappa; \mathcal{D})$ denote the common value of $\bar{U}'_{P,\pm}(\kappa; \mathcal{D})$ when $C'(\kappa; \mathcal{D})$ exists. Since $\alpha^*(\cdot; \mathcal{D}^{DB}) = \hat{\alpha}(\cdot; \mathcal{D}^{DB}) = 1$ for any F, the function $\bar{U}'_P(\cdot; \mathcal{D}^{DB})$ does not depend on F. As $\kappa^*(\mathcal{D}^{DB})$ is the unique zero of this function, it does not depend on F either, and hence neither does M.

The bound $\kappa^*(\mathcal{D}^{DB}) \in (\mu_E^D, \mu_E^M)$ established in Theorem 1 implies that $M \in (0,1)$. Since $\Phi(\cdot; F)$ is decreasing, it follows that $\alpha^*(\kappa^*(\mathcal{D}^{DB})) \leq 1$ iff $\Phi(1; F) \leq M$, and this latter inequality is equivalent to the condition reported in the lemma statement. In the other direction, $\alpha^*(\kappa^*(\mathcal{D}^{DB})) \geq \alpha^\dagger$ iff $\Phi(\alpha^\dagger; F) \geq M$, which is the stated condition.

F.7 Proposition 2

Suppose that under some data privacy policy $\mathcal{D}$, the cost function $\mathcal{C}$ is differentiable everywhere. Then the proof of Theorem 1 establishes that $\text{sign}(\kappa^*(\mathcal{D}) - \kappa) = \text{sign}(\bar{U}'_P(\kappa; \mathcal{D}))$, where $\bar{U}'_{P,\pm}(\kappa; \mathcal{D})$ is as defined in that proof and where $\bar{U}'_P(\kappa; \mathcal{D})$ denotes the common value of $\bar{U}'_{P,\pm}(\kappa; \mathcal{D})$ when $C'(\kappa; \mathcal{D})$ exists.

We first derive conditions under which a data ban reduces innovation. Recall that $\hat{\alpha}(\kappa; \mathcal{D}^{DB}) = 1$ for all κ. Meanwhile for every $\kappa > \mu_E^D$ we have $\alpha^*(\kappa) > \underline{\alpha}$ and therefore

$$\hat{\alpha}(\kappa) = \mathbb{E}[\alpha \mid \alpha \geq \alpha^*(\kappa)] > \mathbb{E}[\alpha] = 1,$$

with the strict inequality following from the fact that F has full support on $[\underline{\alpha}, \overline{\alpha}]$. In particular, $\hat{\alpha}(\kappa^*(\mathcal{D}^{DB})) > \hat{\alpha}(\kappa^*(\mathcal{D}^{DB}); \mathcal{D}^{DB})$ given that $\kappa^*(\mathcal{D}^{DB}) > \mu_E^D$. Then $\kappa^*(\mathcal{D}^{DB}) > \kappa^*$ if $\alpha^*(\kappa^*(\mathcal{D}^{DB})) \leq 1$, since in that case

$$\bar{U}'_P(\kappa^*(\mathcal{D}^{DB}); \mathcal{D}^{LF}) > \bar{U}'_P(\kappa^*(\mathcal{D}^{DB}); \mathcal{D}^{DB}) = 0.$$

We now derive conditions under which a data ban boosts innovation. Define $\Delta\bar{U}'_P(\kappa) \equiv \bar{U}'_P(\kappa; \mathcal{D}^{LF}) - \bar{U}'_P(\kappa; \mathcal{D}^{DB})$. Since $\bar{U}'_P(\kappa^*(\mathcal{D}^{DB}); \mathcal{D}^{DB}) = 0$, it is sufficient to sign $\Delta\bar{U}'_P(\kappa^*(\mathcal{D}^{DB}))$. Note that

$$\Delta\bar{U}'_P(\kappa) = \frac{G'(\kappa)}{G(\kappa)} (\mu_E^M - \kappa) \left(\Lambda(\hat{\alpha}(\kappa)) - \Lambda(1)\right) - \left(\Lambda(\alpha^*(\kappa)) - \Lambda(1)\right).$$

Since $\hat{\alpha}(\kappa) < \infty$ for all κ, we may bound $\Delta\bar{U}'_P$ from above as $\Delta\bar{U}'_P(\kappa) < \phi(\kappa, \alpha^*(\kappa))$ for all $\kappa < \mu_E^M$, where

$$\phi(\kappa, \alpha) \equiv \frac{G'(\kappa)}{G(\kappa)}(\mu_E^M - \kappa)\,(\Lambda(\infty) - \Lambda(1)) - (\Lambda(\alpha) - \Lambda(1)).$$

(Note that $\Lambda(\infty) = \Delta\mu_P/|\Delta\mu_E|$ is finite, and $\Lambda(\infty) > \Lambda(1)$.)

The function ϕ has no direct dependence on F once κ and α are specified. Additionally, ϕ is decreasing in α and

$$\phi(\kappa, \infty) = \left(\frac{G'(\kappa)}{G(\kappa)}(\mu_E^M - \kappa) - 1\right)(\Lambda(\infty) - \Lambda(1))\,.$$

The FOC satisfied by $\kappa^*(\mathcal{D}^{DB})$ implies that

$$\frac{G'(\kappa^*(\mathcal{D}^{DB}))}{G(\kappa^*(\mathcal{D}^{DB}))}(\mu_E^M - \kappa^*(\mathcal{D}^{DB})) - 1 = -\frac{G'(\kappa^*(\mathcal{D}^{DB}))}{G(\kappa^*(\mathcal{D}^{DB}))}\frac{\mu_P^M}{\Lambda(1)} < 0.$$

Hence $\phi(\kappa^*(\mathcal{D}^{DB}), \infty) < 0$. Meanwhile trivially $\phi(\kappa^*(\mathcal{D}^{DB}), 1) > 0$. Since ϕ does not depend on F except through κ and α, and since (as noted in the proof of Lemma 4) $\kappa^*(\mathcal{D}^{DB})$ does not depend on F, it follows that there exists an $\alpha^\dagger > 1$ independent of F such that $\phi(\kappa^*(\mathcal{D}^{DB}), \alpha) \le 0$ iff $\alpha \ge \alpha^\dagger$. Given that $\phi(\kappa^*(\mathcal{D}^{DB}), \alpha^*(\kappa^*(\mathcal{D}^{DB})))$ is a strict upper bound on $\Delta\bar{U}_P'(\kappa^*(\mathcal{D}^{DB}))$, we may therefore conclude that $\Delta\bar{U}_P'(\kappa^*(\mathcal{D}^{DB})) < 0$ if $\alpha^*(\kappa^*(\mathcal{D}^{DB})) \ge \alpha^\dagger$.

F.8 Lemma 5

If $T^P = 0$, then the unique optimal imitation policy is as characterized in Lemma 1. Since $\bar{\kappa}^P(0) = \mu_E^D$, the imitation policy described in the lemma statement is equivalent to the policy derived in that lemma. Going forward, we assume $T^P > 0$.

We first characterize the platform's optimal deterministic imitation policy. Define

$$\underline{\kappa}(T^P) \equiv \mu_E^M + e^{-rT^P}\Delta\mu_E.$$

Note that $\underline{\kappa}(T^P) \in (\mu_E^D, \mu_E^M)$. Any deterministic imitation policy which delays imitation until time T^P can induce an innovation threshold no lower than $\underline{\kappa}(T^P)$. As a result, there exists a unique deterministic imitation policy implementing $\kappa < \underline{\kappa}(T^P)$. It involves imitating at the unique time $T^{D,*}(\kappa; T^P) < T^P$ satisfying

$$\mu_E^M + e^{-rT}\Delta\mu_E = \kappa.$$

Conversely, no deterministic imitation policy implementing $\kappa \ge \underline{\kappa}(T^P)$ can imitate prior to time T^P. Given that an optimal imitation policy delays imitation until patent expiration, logic very similar to the proof of Lemma 1 establishes that the unique optimal deterministic policy $T^{D,*}(\kappa; T^P)$ imitates at time T^P in states $\alpha \ge \alpha^{D,*}(\kappa; T^P)$ and never imitates otherwise, where $\alpha^{D,*}(\kappa; T^P)$ is the unique solution to

$$\mu_E^M + e^{-rT^P} \int_{\alpha}^{\overline{\alpha}} \alpha' \, dF(\alpha) = \kappa.$$

Note that $\alpha^{D,*}$ is continuous and increasing in κ and satisfies the boundary conditions

$$\alpha^{D,*}(\underline{\kappa}(T^P); T^P) = 0, \quad \alpha^{D,*}(\mu_E^M; T^P) = \infty.$$

It is also continuous and decreasing in T^P for fixed κ. Extend the definition of $\alpha^{D,*}$ to $\kappa < \underline{\kappa}$ by letting $\alpha^{D,*}(\kappa; T^P) = 1$.

Now let $C^D(\kappa; T^P)$ denote the cost of the optimal deterministic imitation policy implementing κ. This function is continuous and decreasing everywhere, and its derivative exists for all $\kappa \neq \underline{\kappa}(T^P)$ and satisfies

$$(C^D)'(\kappa; T^P) = \frac{1}{\Delta\mu_E \alpha^{D,*}(\kappa; T^P)}.$$

C^D therefore has a concave kink at $\kappa = \underline{\kappa}(T^P)$ and is convex on $[\underline{\kappa}(T^P), \mu_E^M]$. Define a function $\phi(\kappa; T^P)$ by

$$\phi(\kappa; T^P) = \frac{1 - C^D(\kappa; T^P)}{\kappa - \mu_E^D} + (C^D)'(\kappa; T^P).$$

Note that $\phi(\underline{\kappa}(T^P)+; T^P) = -\infty$ while $\phi(\mu_E^M; T^P) = 1/|\Delta\mu_E| > 0$, and $\phi(\cdot; T^P)$ is continuous and increasing on $(\underline{\kappa}(T^P), \mu_E^M]$ given that C^D is continuous, decreasing, and convex on that range. Hence there exists a unique $\bar{\kappa}(T^P) \in (\underline{\kappa}(T^P), \mu_E^M)$ satisfying $\phi(\kappa; T^P) = 0$.

Since $\alpha^{D,*}$ is continuous and decreasing in T^P for fixed $\kappa > \underline{\kappa}(T^P)$, it follows that $(C^D)'(\kappa; T^P)$ is continuous and decreasing in T^P for fixed κ and therefore $C^D(\kappa; T^P)$ is increasing in T^P for fixed κ. Thus $\phi(\kappa; T^P)$ is continuous and decreasing in T^P for fixed κ, implying that $\bar{\kappa}$ is continuous and increasing. Additionally, $C^D(\cdot; 0) = C(\cdot; \mathcal{D}^{LF})$ and so $C^D(\cdot; 0)$ is strictly convex on $[\mu_E^D, \mu_E^M]$, implying that $\phi(\kappa; 0) > 0$ for all $\kappa > \mu_E^D$ and therefore $\bar{\kappa}(\mu_E^D+) = \mu_E^D$. In the other extreme, $\underline{\kappa}(\infty) = \mu_E^M$ and so $\bar{\kappa}(\infty) = \mu_E^M$.

Define

$$\bar{C}(\kappa; T^P) \equiv \begin{cases} 1 + (C^D)'(\bar{\kappa}(T^P); T^P)(\kappa - \mu_E^D), & \kappa < \bar{\kappa}(T^P) \\ C^D(\kappa; T^P), & \kappa \geq \bar{\kappa}(T^P) \end{cases}$$

By construction, $\bar{C}(\cdot; T^P)$ is a convex function which lies below $C(\cdot; T^P)$ everywhere. Further, it is the largest such function. This is trivially so for $\kappa \geq \bar{\kappa}(T^P)$ and $\kappa = \mu_E^D$, while for $\kappa \in (\mu_E^D, \bar{\kappa}(T^P))$ any convex function $\widetilde{C} \leq C^D(\cdot; T^P)$ which satisfied $\widetilde{C}(\kappa) > \bar{C}(\kappa; T^P)$ for some $\kappa <\in (\mu_E^D, \bar{\kappa}(T^P))$

would have

$$
\begin{aligned}
\widetilde{C}(\mu_E^D) &\geq \widetilde{C}(\kappa) - \widetilde{C}'_+(\kappa)(\kappa - \mu_E^D) \\
&\geq \widetilde{C}(\kappa) - \frac{\widetilde{C}(\bar{\kappa}(T^P)) - \widetilde{C}(\kappa)}{\bar{\kappa}(T^P) - \kappa}(\kappa - \mu_E^D) \\
&> \bar{C}(\kappa; T^P) - \frac{C^D(\bar{\kappa}(T^P); T^P) - \bar{C}(\kappa; T^P)}{\bar{\kappa}(T^P) - \kappa}(\kappa - \mu_E^D) \\
&= \bar{C}(\kappa; T^P) - \frac{\bar{C}(\bar{\kappa}(T^P); T^P) - \bar{C}(\kappa; T^P)}{\bar{\kappa}(T^P) - \kappa}(\kappa - \mu_E^D) \\
&= \bar{C}(\mu_E^D; T^P) = 1,
\end{aligned}
$$

contradicting $\widetilde{C}(\mu_E^D) \leq C^D(\mu_E^D; T^P) = 1$. Thus $\bar{C}(\cdot; T^P)$ is the convex envelope of $C^D(\cdot; T^P)$, implying that $C(\cdot; T^P) = \bar{C}(\cdot; T^P)$.

The fact that $C(\kappa; T^P) = C^D(\kappa; T^P)$ for all $\kappa \geq \bar{\kappa}(T^P)$, combined with the strict convexity of $C(\cdot; T^P)$ on this range, implies that for all $\kappa \geq \bar{\kappa}(T^P)$ the unique optimal deterministic imitation policy remains uniquely optimal within the wider class of randomized policies. This result establishes the form and uniqueness of the optimal policy claimed in the lemma statement.

Now fix a target threshold $\kappa < \bar{\kappa}(T^P)$. The fact that $C(\kappa; T^P) < C^D(\kappa; T^P)$ for all $\kappa \in (\mu_E^D, \bar{\kappa}(T^P))$ implies that no randomized policy implementing κ and assigning positive probability to deterministic policies implementing thresholds in the range $(\mu_E^D, \bar{\kappa}(T^P))$ can be optimal. Additionally, strict convexity of $C(\cdot; T^P)$ on $[\bar{\kappa}(T^P), \mu_E^M]$ implies that no randomized policy implementing κ and assigning positive probability to deterministic policies implementing thresholds in the range $(\bar{\kappa}(T^P), \mu_E^M]$ can be optimal. Hence any optimal policy implementing κ must assigning positive probability only to deterministic policies implementing thresholds in the range $\{\mu_E^D, \bar{\kappa}(T^P)\}$. There exists a unique imitation policy implementing κ and satisfying this requirement. It involves randomization between the optimal deterministic policies implementing $\kappa = \mu_E^D$ and $\kappa = \bar{\kappa}(T^P)$, with probability $\rho(\kappa; T^P)$ assigned to the former policy, where $\rho(\kappa; T^P)$ is as defined in the statement of the lemma. This result establishes the form and uniqueness of the the optimal policy claimed in the lemma statement.

F.9 Proposition 3

Extend ρ to innovation thresholds above $\bar{\kappa}$ by letting $\rho(\kappa; T^P) = 0$. Additionally define $\rho^*(T^P) \equiv \rho(\kappa^*(\mathcal{D}^P(T^P)); T^P)$ to be the probability of immediate imitation under the platform's optimal innovation threshold. We prove the following technical result, which nests the results of Proposition 3.

Proposition F.1. *ρ^* is continuous, and there exists a patent length $\overline{T}^P \in (0,\infty)$ such that:*

- *If $T^P \leq \overline{T}^P$, then $\rho^*(T^P) = 0$*
- *If $T^P > \overline{T}^P$, then ρ^* is positive and increasing in T^P*

For $T^P \geq \overline{T}^P$, $\kappa^*(\mathcal{D}^P(T^P))$ is decreasing in T^P and $\lim_{T^P\to\infty} \kappa^*(\mathcal{D}^P(T^P)) = \kappa^*(\mathcal{D}^{DB})$ while $\lim_{T^P\to\infty} \rho^*(T^P) = (\mu_E^M - \kappa^*(\mathcal{D}^{DB}))/|\Delta\mu_E|$.

Proof. Note that $\rho(\kappa; T^P)$ is a jointly continuous function of (κ, T^P) given that $\bar{\kappa}$ is continuous. Hence ρ^* is continuous so long as $\kappa^*(\mathcal{D}^P(T^P))$ is continuous in T^P. Recall from Lemma 1 that $\kappa^*(\mathcal{D}^P(T^P))$ uniquely optimizes the function

$$U_P(\kappa; \mathcal{D}^P(T^P)) = G(\kappa)\left(\mu_P^M + \Delta\mu_P \frac{\mu_E^M - \kappa}{|\Delta\mu_E|} - k_P C(\kappa; \mathcal{D}^P(T^P))\right).$$

Continuity of $\kappa^*(\mathcal{D}^P(T^P))$ is therefore ensured by the maximum theorem so long as $C(\kappa; \mathcal{D}^P(T^P))$ is continuous in T^P for each κ.

First consider $T^P = 0$. Trivially $C(\kappa; \mathcal{D}^P(0)) = C(\kappa; \mathcal{D}^{LF})$ for all κ. Further, for any T^P we have the bounds $C^D(\kappa; T^P) \geq C(\kappa; \mathcal{D}^P(T^P)) \geq C(\kappa; \mathcal{D}^{LF})$ holding for all κ and T^P, where C^D is the cost function for deterministic policies characterized in the proof of Lemma 5. It is therefore sufficient to show that $\lim_{T^P\downarrow 0} C^D(\kappa; T^P) = C(\kappa; \mathcal{D}^{LF})$ for fixed κ. Using the characterization derived in the proof of Lemma 5, this cost function may be written

$$C^D(\kappa; T^P) = 1 - \int_{\mu_E^D}^{\kappa} \frac{1}{|\Delta\mu_E|\alpha^{D,*}(\kappa'; T^P)}\, d\kappa',$$

where

$$\alpha^{D,*}(\kappa; T^P) = \Phi^{-1}\left(e^{rT^P}\frac{\mu_E^M - \kappa}{|\Delta\mu_E|}\right)$$

and

$$\Phi(\alpha) = \int_{\alpha}^{\overline{\alpha}} \alpha'\, dF(\alpha').$$

Since Φ is decreasing in α, so is Φ^{-1}, meaning that $\alpha^{D,*}$ is decreasing in T^P for every κ. Additionally, $\alpha^{D,*}(\kappa; 0) = \alpha^*(\kappa)$ for all κ. The monotone convergence theorem then ensures that

$$\lim_{T^P\downarrow 0} C^D(\kappa; T^P) = 1 - \int_{\mu_E^D}^{\kappa} \frac{1}{|\Delta\mu_E|\alpha^*(\kappa')}\, d\kappa' = C(\kappa; \mathcal{D}^{LF}),$$

as desired.

Now consider any $T^P > 0$. Using the characterization derived in the proof of Lemma 5, $C(\kappa; \mathcal{D}^P(T^P))$

may be written

$$C(\kappa;\mathcal{D}^P(T^P)) = 1 - \int_{\mu_E^D}^{\kappa} \frac{1}{|\Delta\mu_E|\alpha^{D,*}(\max\{\kappa',\bar{\kappa}(T^P)\};T^P)}\,d\kappa'.$$

Since $\bar{\kappa}$ and $\underline{\kappa}$ are continuous and $\bar{\kappa}(T^P) > \underline{\kappa}(T^P)$, the quantity $\alpha^{D,*}(\max\{\kappa',\bar{\kappa}(T')\};T')$ is uniformly bounded away from zero for all $\kappa' \in [\mu_E^D, \mu_E^M]$ and T' on any sufficiently small neighborhood of T^P. The bounded convergence theorem and continuity of $\bar{\kappa}$ then imply that $C(\kappa;\mathcal{D}^P(T'))$ is continuous in T' at T^P if $\alpha^{D,*}$ is jointly continuous in both arguments at (κ, T^P). This joint continuity is ensured by the fact that $\alpha^{D,*}$ is a composition of (jointly) continuous functions.

Now, fix any T^P such that $\kappa^*(\mathcal{D}^P(T^P)) \le \bar{\kappa}(\mathcal{D}^P(T^P))$. Theorem 1 established that

$$\bar{U}_P'(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T^P)) = 0,$$

where $\bar{U}_P'(\kappa;\mathcal{D})$ is the common value of $\bar{U}_{P,\pm}'(\kappa;\mathcal{D})$ when $C'(\kappa;\mathcal{D})$ exists and $\bar{U}_{P,\pm}'$ are as defined in the proof of Theorem 1. (Note that $C(\kappa;\mathcal{D}^P(\kappa^P))$ is differentiable everywhere given the smooth-pasting condition satisfied by $\bar{\kappa}$.) Fix $T' > T^P$. We will show that

$$\bar{U}_P'(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T')) < \bar{U}_P'(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T')).$$

This inequality implies that $\bar{U}_P'(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T')) < 0$ and therefore

$$\kappa^*(\mathcal{D}^P(T')) < \kappa^*(\mathcal{D}^P(T^P)) \le \bar{\kappa}(T^P) < \bar{\kappa}(T').$$

In particular, $\rho^*(T') > 0$.

To establish the desired inequality, it is sufficient to show that

$$C(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T')) > C(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T^P)),$$
$$C'(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T')) > C'(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T^P)),$$

since these inequalities imply

$$\hat{\alpha}(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T')) < \hat{\alpha}(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T^P)),$$
$$\alpha^*(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T')) > \alpha^*(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T^P))$$

and therefore $\bar{U}_P'(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T')) < \bar{U}_P'(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T'))$. The first inequality is immediate from the fact that the unique policy achieving costs $C(\kappa^*(\mathcal{D}^P(T^P));\mathcal{D}^P(T^P))$ under a data patent of length T^P requires imitation at time T^P with positive probability (since $\kappa^*(\mathcal{D}^P(T^P)) > \mu_E^D$) and so is not implementable under a data patent of length $T' > T^P$. The first inequality, convexity of $C(\cdot;\mathcal{D}^P(T')))$, linearity of $C(\cdot;\mathcal{D}^P(T^P))$ on $[\mu_E^D, \bar{\kappa}(T^P)]$, and the boundary condition $C(\mu_E^D;\mathcal{D}^P(T')) = C(\mu_E^D;\mathcal{D}^P(T^P)) = 1$ then imply the second inequality.

We next prove that $\lim_{T^P\to\infty} \kappa^*(\mathcal{D}^P(T^P)) = \kappa^*(\mathcal{D}^{DB})$. Note that for all $\kappa \leq \bar{\kappa}(T^P)$ we have

$$C(\kappa;\mathcal{D}^P(T^P)) = 1 - \frac{\kappa - \mu_E^D}{|\Delta\mu_E|\alpha^{*,D}(\bar{\kappa}(T^P);T^P)}, \quad C'(\kappa;\mathcal{D}^P(T^P)) = -\frac{1}{|\Delta\mu_E|\alpha^{*,D}(\bar{\kappa}(T^P);T^P)}.$$

Given the bounds $C(\kappa;\mathcal{D}^{LF}) \leq C(\kappa;\mathcal{D}^P(T^P)) \leq (\mu_E^M - \kappa)/|\Delta\mu_E|$ and the limit $\bar{\kappa}(\infty) = \mu_E^M$, it follows that $\lim_{T^P\to\infty} C(\bar{\kappa}(T^P);\mathcal{D}^P(T^P)) = 0$. Evaluating the previous expression for $C(\kappa;\mathcal{D}^P(T^P))$ at $\bar{\kappa}(T^P)$ and taking a limit yields $\lim_{T^P\to\infty} \alpha^{*,D}(\bar{\kappa}(T^P);T^P) = 1$. Hence for every $\kappa < \mu_E^M$ we have

$$\lim_{T^P\to\infty} C(\kappa;\mathcal{D}^P(T^P)) = \frac{\mu_E^M - \kappa}{|\Delta\mu_E|} = C(\kappa;\mathcal{D}^{DB}), \quad \lim_{T^P\to\infty} C'(\kappa;\mathcal{D}^P(T^P)) = \frac{1}{\Delta\mu_E} = C'(\kappa;\mathcal{D}^{DB}).$$

It follows that for every $\kappa < \mu_E^M$ we have $\lim_{T^P\to\infty} \bar{U}_P'(\kappa;\mathcal{D}^P(T^P)) = \bar{U}_P'(\kappa;\mathcal{D}^{DB})$. For any $\kappa < \kappa^*(\mathcal{D}^{DB})$ we therefore have $\bar{U}_P'(\kappa;\mathcal{D}^P(T^P)) > 0$ for T^P sufficiently large, implying that $\kappa^*(\mathcal{D}^P(T^P)) > \kappa$ for T^P large. Since this argument holds for any $\kappa < \kappa^*(\mathcal{D}^{DB})$, it follows that $\liminf_{T^P\to\infty} \kappa^*(\mathcal{D}^P(T^P)) \geq \kappa^*(\mathcal{D}^{DB})$. A similar argument from the other direction, along with the fact that $\kappa^*(\mathcal{D}^{DB}) < \mu_E^M$, implies that $\limsup_{T^P\to\infty} \kappa^*(\mathcal{D}^P(T^P)) \leq \kappa^*(\mathcal{D}^{DB})$. Combining these two limits yields $\lim_{T^P\to\infty} \kappa^*(\mathcal{D}^P(T^P)) = \kappa^*(\mathcal{D}^{DB})$. Given continuity of ρ^*, this limit along with $\lim_{T^P\to\infty} \bar{\kappa}(T^P) = \mu_E^M$ implies the associated limiting expression for ρ^* in the proposition statement.

Define $\overline{T}^P \equiv \inf\{T^P \geq 0 \ : \ \rho^*(T^P) > 0\}$. The work so far proves that if $\rho(T^P) > 0$ for some T^P, then $\rho(T') > 0$ for all $T' > T^P$. Thus $\rho^*(T^P) = 0$ for all $T \leq \overline{T}^P$ while $\rho^*(T^P) > 0$ for all $T > \overline{T}^P$. We next establish that $\overline{T}^P \in (0,\infty)$. The bound $\overline{T}^P > 0$ follows from the facts that $\kappa^*(\mathcal{D}^P(T^P))$ is continuous in T^P, $\kappa^*(\mathcal{D}^P(0)) = \kappa^* > \mu_E^D$, and $\bar{\kappa}$ is continuous and satisfies $\bar{\kappa}(0) = \mu_E^D$. At the other extreme, $\lim_{T^P\to\infty} \bar{\kappa}(T^P) = \mu_E^M$ and $\lim_{T^P\to\infty} \kappa^*(\mathcal{D}^P(T^P)) = \kappa^*(\mathcal{D}^{DB}) < \mu_E^M$ implies that $\kappa^*(\mathcal{D}^P(T^P)) < \bar{\kappa}(T^P)$ for T^P sufficiently large. Hence $T^P < \infty$.

The definition of $\overline{T}^P$ and continuity of ρ^* and $\kappa^*(\mathcal{D}^P(T^P))$ imply that $\kappa^*(\mathcal{D}^P(T^P)) \leq \bar{\kappa}(T^P))$ for all $T^P \geq \overline{T}^P$. And we established above that if $\kappa^*(\mathcal{D}^P(T^P)) \leq \bar{\kappa}(T^P)$ then $\kappa^*(\mathcal{D}^P(T')) < \kappa^*(\mathcal{D}^P(T^P))$ for all $T' > T^P$. Hence $\kappa^*(\mathcal{D}^P(T^P))$ is decreasing in T^P for $T^P \geq \overline{T}^P$. Finally, note that ρ is decreasing in κ and increasing in $\bar{\kappa}(T^P)$, and therefore in T^P. Hence ρ^* is increasing in T^P for all $T^P \geq \overline{T}^P$. □

F.10 Proposition 4

To establish the result, it is sufficient to prove that under the hypothesis of the proposition, $\kappa^*(\mathcal{D}^P(T^P))$ is increasing in T^P on $[0, \overline{T}^P]$, where $\overline{T}^P$ is as defined in Proposition F.1. Combining this fact with the result that $\kappa^*(\mathcal{D}^P(T^P))$ is decreasing on $[\overline{T}^P, \infty)$ (established in Proposition F.1) implies the result.

Lemma 5 and Proposition 3 established that for $T^P \leq \overline{T}^P$, $\kappa^*(\mathcal{D}^P(T^P))$ maximizes the profit function

$$U_P^D(\kappa; T^P) \equiv G(\kappa)\left(\mu_P^M + (\mu_E^M - \kappa)\Lambda(\hat{\alpha}^D(\kappa; T^P))\right),$$

where

$$\hat{\alpha}^D(\kappa; T^P) \equiv \frac{\int_{\alpha^{D,*}(\kappa;T^P)}^{\overline{\alpha}} e^{-rT^P}\alpha\, dF(\alpha)}{\int_{\alpha^{D,*}(\kappa;T^P)}^{\overline{\alpha}} e^{-rT^P}\, dF(\alpha)} = \frac{\int_{\alpha^{D,*}(\kappa;T^P)}^{\overline{\alpha}} \alpha\, dF(\alpha)}{1 - F(\alpha^{D,*}(\kappa; T^P))} = \mathbb{E}[\alpha \mid \alpha \geq \alpha^{D,*}(\kappa; T^P)]$$

and $\alpha^{D,*}$ is as defined in the proof of Lemma 5. Using techniques very similar to those used in the proof of Theorem 1, this profit function can be shown to be maximized by the unique solution to the FOC $(\bar{U}^D)'(\kappa; T^P) = 0$, where

$$(\bar{U}^D)'(\kappa; T^P) \equiv \frac{G'(\kappa)}{G(\kappa)}\left(\mu_P^M + (\mu_E^M - \kappa)\Lambda\left(\hat{\alpha}^D(\kappa; T^P)\right)\right) - \Lambda\left(\alpha^{D,*}(\kappa; T^P)\right).$$

Fix $T^P < \overline{T}^P$ and $T' \in (T^P, \overline{T}^P]$, and let $\kappa = \kappa^*(\mathcal{D}^P(T^P))$. Rewrite $(\bar{U}^D)'$ as

$$(\bar{U}^D)'(\kappa; T') = \Lambda\left(\alpha^{D,*}(\kappa; T')\right)\left(\frac{G'(\kappa)}{G(\kappa)}\left(\frac{\mu_P^M}{\Lambda\left(\alpha^{D,*}(\kappa; T')\right)} + (\mu_E^M - \kappa)\frac{\Lambda\left(\hat{\alpha}^D(\kappa; T')\right)}{\Lambda\left(\alpha^{D,*}(\kappa; T')\right)}\right) - 1\right)$$

Given the form of $\hat{\alpha}^D$, we may also write

$$\frac{\Lambda\left(\hat{\alpha}^D(\kappa; T')\right)}{\Lambda\left(\alpha^{D,*}(\kappa; T')\right)} = \frac{\Lambda(\hat{A}(\alpha^{D,*}(\kappa; T')))}{\Lambda(\alpha^{D,*}(\kappa; T'))}.$$

Recall that $\alpha^{*,D}$ is decreasing in the length of the data patent. The hypothesis that $\Lambda(\alpha)/\Lambda(\hat{A}(\alpha))$ is nondecreasing therefore implies the inequality

$$\frac{\Lambda\left(\hat{\alpha}^D(\kappa; T')\right)}{\Lambda\left(\alpha^{D,*}(\kappa; T')\right)} \geq \frac{\Lambda\left(\hat{\alpha}^D(\kappa; T^P)\right)}{\Lambda\left(\alpha^{D,*}(\kappa; T^P)\right)}.$$

Since additionally $\Lambda\left(\alpha^{D,*}(\kappa;T')\right) < \Lambda\left(\alpha^{D,*}(\kappa;T^P)\right)$, we can bound $(\bar{U}^D)'(\kappa;T')$ below as

$$(\bar{U}^D)'(\kappa;T') > \Lambda\left(\alpha^{D,*}(\kappa;T')\right)\left(\frac{G'(\kappa)}{G(\kappa)}\left(\frac{\mu_P^M}{\Lambda\left(\alpha^{D,*}(\kappa;T^P)\right)} + (\mu_E^M - \kappa)\frac{\Lambda\left(\hat{\alpha}^D(\kappa;T^P)\right)}{\Lambda\left(\alpha^{D,*}(\kappa;T^P)\right)}\right) - 1\right)$$
$$= \frac{\Lambda\left(\alpha^{D,*}(\kappa;T')\right)}{\Lambda\left(\alpha^{D,*}(\kappa;T^P)\right)}(\bar{U}^D)'(\kappa;T^P) = 0.$$

Since $\bar{U}^D$ crosses zero once from above, this inequality implies that $\kappa^*(\mathcal{D}^P(T')) > \kappa^*(\mathcal{D}^P(T^P))$, as desired.

F.11 Lemma 6

We first derive a sufficient condition for monotonicity of $\Lambda(\alpha)/\Lambda(\hat{A}(\alpha))$. Note that

$$\frac{\Lambda(\alpha)}{\Lambda(\hat{A}(\alpha))} - 1 = \Lambda(\hat{A}(\alpha))^{-1}\left(\Lambda(\alpha) - \Lambda(\hat{A}(\alpha))\right).$$

Suppose that $\Lambda(\alpha) - \Lambda(\hat{A}(\alpha))$ is nondecreasing in α. Then since this expression is negative while $\Lambda(\hat{A}(\alpha))$ is positive and increasing in α, it follows that $\Lambda(\alpha)/\Lambda(\hat{A}(\alpha)) - 1$ is increasing in α, hence also $\Lambda(\alpha)/\Lambda(\hat{A}(\alpha))$ is increasing. Further,

$$\Lambda(\alpha) - \Lambda(\hat{A}(\alpha)) = \frac{k_P}{|\Delta\mu_E|}\hat{A}(\alpha)^{-1}\left(1 - \frac{\hat{A}(\alpha)}{\alpha}\right).$$

Since $\hat{A}$ is increasing in α and $\hat{A}(\alpha)/\alpha \geq 1$, it is therefore sufficient to establish that $\hat{A}(\alpha)/\alpha$ is nonincreasing in α. Note that $\hat{A}(\alpha)$ is independent of all model parameters except F, and so monotonicity of $\hat{A}(\alpha)/\alpha$ amounts to a condition on the demand distribution alone.

We now show that each family of distributions listed in the lemma statement satisfies the condition that $\hat{A}(\alpha)/\alpha$ is nonincreasing in α. For each family, we enforce appropriate restrictions on free parameters to ensure that demand has mean 1 and non-negative support.

Uniform: Suppose that demand is uniformly distributed on the interval $[1-\Delta, 1+\Delta]$ for $\Delta \in (0,1]$. Then

$$\hat{A}(\alpha) = \frac{\alpha + 1 + \Delta}{2}$$

and

$$\frac{\hat{A}(\alpha)}{\alpha} = \frac{1}{2}\left(1 + \frac{1+\Delta}{\alpha}\right),$$

which is decreasing in α.

Lognormal: Suppose that demand is lognormally distributed with log-variance $\sigma > 0$ and log-mean $\mu = -\sigma^2/2$. A standard result about conditional means of lognormal distributions is

$$\hat{A}(\alpha) = \frac{1 - \Phi\left(-\frac{\sigma}{2} + \frac{\log \alpha}{\sigma}\right)}{1 - \Phi\left(\frac{\sigma}{2} + \frac{\log \alpha}{\sigma}\right)},$$

where Φ is the distribution function of a standard normal distribution. Now, note that

$$\frac{\phi\left(-\frac{\sigma}{2} + \frac{\log \alpha}{\sigma}\right)}{\phi\left(\frac{\sigma}{2} + \frac{\log \alpha}{\sigma}\right)} = \alpha,$$

where ϕ is the density function of a standard normal distribution. Therefore

$$\frac{\hat{A}(\alpha)}{\alpha} = \frac{h\left(\frac{\sigma}{2} + \frac{\log \alpha}{\sigma}\right)}{h\left(-\frac{\sigma}{2} + \frac{\log \alpha}{\sigma}\right)},$$

where $h(x) \equiv \phi(x)/(1 - \Phi(x))$ is the hazard rate of a standard normal distribution. Now, $\hat{A}(\alpha)/\alpha$ is nonincreasing if and only if its logarithm is nonincreasing, and the representation for $\hat{A}(\alpha)/\alpha$ just derived implies that

$$\log \frac{\hat{A}(\alpha)}{\alpha} = \int_{-\sigma/2+\sigma^{-1}\log \alpha}^{\sigma/2+\sigma^{-1}\log \alpha} \tilde{h}'(x)\, dx,$$

where $\tilde{h}(x) \equiv \log h(x)$. This expression is nonincreasing in α if $\tilde{h}$ is concave. Direct computation yields $\tilde{h}'(x) = h(x) - x$, so that $\tilde{h}''(x) = h'(x) - 1$. A standard fact about $h(x)$ is that it is strictly convex with limiting slope $h'(\infty) = 1$. Thus $\tilde{h}$ is concave, implying that $\hat{A}(\alpha)/\alpha$ is nonincreasing.

Pareto: Suppose that demand is Pareto distributed with shape parameter $\beta > 1$ and scale parameter $\underline{x} = (\beta - 1)/\beta$. Then

$$\hat{A}(\alpha) = \frac{\beta}{\beta - 1}\alpha$$

and $\hat{A}(\alpha)/\alpha = \beta/(\beta - 1)$, which is nonincreasing in α.

Gamma: Suppose that demand is gamma distributed with shape parameter $\nu > 0$ and scale parameter $\theta = 1/\nu$. Then

$$\hat{A}(\alpha) = \nu^{-1}\frac{\Gamma(\nu + 1, \nu\alpha)}{\Gamma(\nu, \nu\alpha)},$$

where $\Gamma(s, x)$ is the upper incomplete gamma function. This function satisfies the identity

$$\Gamma(s + 1, x) = s\Gamma(s, x) + x^s e^{-x},$$

which can be used to rewrite $\hat{A}(\alpha)$ as

$$\hat{A}(\alpha) = 1 + \frac{\alpha^\nu e^{-\nu\alpha}}{\Gamma(\nu, \nu\alpha)}.$$

Using the derivative $\Gamma_x(s, x) = -x^{s-1}e^{-x}$, we have

$$\frac{d}{d\alpha}\left(\frac{\hat{A}(\alpha)}{\alpha}\right) = \frac{\zeta(\alpha)}{(\alpha\Gamma(\nu, \nu\alpha))^2},$$

where

$$\zeta(\alpha) \equiv -\Gamma(\nu, \nu\alpha)^2 - \nu^{-1}(\nu\alpha)^\nu e^{-\alpha\nu}(1 - \nu + \nu\alpha)\Gamma(\nu, \nu, \alpha) + \nu^{-1}(\nu\alpha)^{2\nu}\ e^{-2\nu\alpha}.$$

It is therefore enough to show that ζ is non-positive for all α. The standard limits

$$\lim_{x\to\infty} \Gamma(s, x) = \lim_{x\to\infty} x^s e^{-x} = \lim_{x\to\infty} x^{s+1}e^{-x} = 0$$

imply that $\lim_{\alpha\to\infty} \zeta(\alpha) = 0$, so it is sufficient to show that $\zeta'(\alpha) > 0$ for all α. Calculating this derivative yields

$$\zeta'(\alpha) = \nu(1 + \nu(\alpha - 1)^2)(\nu\alpha)^{\nu-1}e^{-\nu\alpha}\xi(\alpha),$$

where

$$\xi(\alpha) \equiv \Gamma(\nu, \nu\alpha) - \nu^{-1}\frac{\nu\alpha - \nu - 1}{1 + \nu(\alpha - 1)^2}(\nu\alpha)^\nu e^{-\nu\alpha}.$$

This expression has limit $\lim_{\alpha\to\infty} \xi(\alpha) = 0$, so $\zeta'(\alpha) > 0$ for all α so long as $\xi'(\alpha) < 0$ for all α. Differentiating and simplifying yields

$$\xi'(\alpha) = -2\nu^{-1}\frac{(\nu\alpha)^{\nu+1}e^{-\nu\alpha}}{(1 + \nu(\alpha - 1)^2)^2} < 0,$$

as desired.

F.12 Lemma 7

We first characterize $C(\cdot; \mathcal{D}^C(\alpha^C))$ for arbitrary α^C. Let $\kappa^C(\alpha^C)$ be the unique solution to $\alpha^C = \alpha^*(\kappa)$. Then for $\kappa \in [\kappa^C(\alpha^C), \mu_E^M]$, the platform has access to enough data to implement κ using the laissez-faire policy characterized in Lemma 1. Hence $C(\kappa; \mathcal{D}^C(\alpha^C)) = C(\kappa; \mathcal{D}^{LF})$ for all such κ. Meanwhile for $\kappa \in [\mu_E^D, \kappa^C(\alpha^C))$, monotonicity of the profit conversion rate in α implies that the lowest-cost feasible implementation of κ involves immediate imitation in all non-censored states, and delayed or random imitation uniformly across all censored states. The optimal policy involving delay chooses an imitation time $\underline{T}(\kappa)$ for censored states which satisfies

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

Therefore the discount factor $\exp(-r\underline{T}(\kappa))$ is linear in κ. Meanwhile, the cost function in this regime may be written in terms of $\underline{T}(\kappa)$ as

$$C(\kappa; \mathcal{D}^C(\alpha^C)) = 1 - \left(1 - e^{-r\overline{T}(\kappa)}\right) F(\alpha^C).$$

Hence C is linear in κ on $[\mu_E^D, \kappa^C(\alpha^C)]$.

Now, fix $\mathcal{D}$ as in the lemma statement. Let $\alpha^C = \alpha^*(\kappa^*(\mathcal{D}))$, so that $\kappa^C(\alpha^C) = \kappa^*(\mathcal{D})$. Then $C(\kappa^*(\mathcal{D}); \mathcal{D}^C(\alpha^C)) = C(\kappa^*(\mathcal{D}); \mathcal{D}^{LF}) \leq C(\kappa^*(\mathcal{D}); \mathcal{D})$. Meanwhile, Lemma 2 implies that $C(\mu_E^D; \mathcal{D}^C(\alpha^C)) = C(\mu_E^D; \mathcal{D}) = 1$, and convexity of the cost function implies that

$$C(\kappa^*(\mathcal{D}); \mathcal{D}) = 1 + \int_{\mu_E^D}^{\kappa^*(\mathcal{D})} C'(\kappa; \mathcal{D}) \, d\kappa \leq 1 + C'_-(\kappa^*(\mathcal{D}); \mathcal{D})(\kappa^*(\mathcal{D}) - \mu_E^D).$$

Suppose that $C'_-(\kappa^*(\mathcal{D}); \mathcal{D}) < C'_-(\kappa^*(\mathcal{D}); \mathcal{D}^C(\alpha^C))$. Then the previous inequality implies that

$$C(\kappa^*(\mathcal{D}); \mathcal{D}) < 1 + (\kappa^*(\mathcal{D}) - \mu_E^D) C'_-(\kappa^*(\mathcal{D}); \mathcal{D}^C(\alpha^C)).$$

Since C is linear on the interval $[\mu_E^D, \kappa^*(\mathcal{D})]$, the rhs of this inequality equals $C(\kappa^*(\mathcal{D}); \mathcal{D}^C(\alpha^C))$, a contradiction. So it must be that $C'_-(\kappa^*(\mathcal{D}); \mathcal{D}) \geq C'_-(\kappa^*(\mathcal{D}); \mathcal{D}^C(\alpha^C))$.

The bounds established in the previous paragraph imply that $\hat{\alpha}(\kappa^*(\mathcal{D}); \mathcal{D}^C(\alpha^C)) \geq \hat{\alpha}(\kappa^*(\mathcal{D}); \mathcal{D})$ and $\alpha^*_-(\kappa^*(\mathcal{D}); \mathcal{D}^C(\alpha^C)) \leq \alpha^*_-(\kappa^*(\mathcal{D}); \mathcal{D})$. Hence $\bar{U}_{P,-}(\kappa^*(\mathcal{D}); \mathcal{D}^C(\alpha^C)) \geq \bar{U}_{P,-}(\kappa^*(\mathcal{D}); \mathcal{D})$, where $\bar{U}_{P,-}$ is as defined in the proof of Theorem 1. Combining this inequality with the bound $\bar{U}_{P,-}(\kappa^*(\mathcal{D}); \mathcal{D}) \geq 0$ implied by optimality of $\kappa^*(\mathcal{D})$ yields $\bar{U}_{P,-}(\kappa^*(\mathcal{D}); \mathcal{D}^C(\alpha^C)) \geq 0$. Further, linearity of $C(\cdot; \mathcal{D}^C(\alpha^C))$ on $[\mu_E^D, \kappa^*(\mathcal{D})]$ implies that $\bar{U}'_P(\cdot; \mathcal{D}^C(\alpha^C))$ exists and is decreasing on $[\mu_E^D, \kappa^*(\mathcal{D})]$. Hence $\bar{U}_P(\kappa; \mathcal{D}^C(\alpha^C)) > 0$ for all $\kappa < \kappa^*(\mathcal{D})$, implying that $\kappa^*(\mathcal{D}^C(\alpha^C)) \geq \kappa^*(\mathcal{D})$.

F.13 Proposition 5

We proceed by deriving a data policy which maximizes innovation. The proof of Lemma 7 establishes that any data policy may be (weakly) improved by a lower censorship policy whose censorship threshold satisfies $\kappa^*(\mathcal{D}^C(\alpha^C)) \geq \kappa^C(\alpha^C)$, where $\kappa^C(\alpha^C)$ is the unique solution to $\alpha^C = \alpha^*(\kappa)$. So we may confine our search for optimal data policies within this class, which we will call *effective*.

Under any lower censorship policy (not necessarily effective), innovation thresholds $\kappa \geq \kappa^C(\alpha^C)$ are implementable using the laissez-faire imitation policy and the cost function therefore satisfies $C(\kappa; \mathcal{D}^C(\alpha^C)) = C(\kappa; \mathcal{D}^{LF})$. In particular, $\bar{U}_{P,+}(\kappa; \mathcal{D}^C(\alpha^C)) = \bar{U}_P(\kappa; \mathcal{D}^{LF})$ for all $\kappa \geq \kappa^C(\alpha^C)$, where $\bar{U}_{P,\pm}$ are as defined in the proof of Theorem 1 and $\bar{U}_P$ denotes the common value of $\bar{U}_{P,\pm}$ when C' exists.

Let $\bar{\alpha}^C \equiv \alpha^*(\kappa^*)$. Note that $\kappa^C(\bar{\alpha}^C) = \kappa^*$ by construction. We first establish that $\alpha^C \leq \bar{\alpha}^C$ implies $\kappa^*(\mathcal{D}^C(\alpha^C)) = \kappa^* \geq \kappa^C(\alpha^C)$. When $\alpha^C \leq \bar{\alpha}^C$, monotonicity of α^* implies that $\kappa^* \geq \kappa^C(\alpha^C)$. Hence $\bar{U}'_{P,+}(\kappa^*; \mathcal{D}^C(\alpha^C)) = \bar{U}'_P(\kappa^*; \mathcal{D}^{LF}) = 0$. Since $\bar{U}'_{P,-}(\kappa; \mathcal{D}^C(\alpha^C)) \geq \bar{U}'_{P,+}(\kappa; \mathcal{D}^C(\alpha^C))$ for all κ, we therefore have

$$\bar{U}'_{P,-}(\kappa^*; \mathcal{D}^C(\alpha^C)) \geq 0 = \bar{U}'_{P,+}(\kappa^*; \mathcal{D}^C(\alpha^C)).$$

In other words, κ^* satisfies the FOC under data policy $\mathcal{D}^C(\alpha^C)$. Since Theorem 1 establishes that the FOC has a unique solution which maximizes U_P, we conclude that $\kappa^*(\mathcal{D}^C(\alpha^C)) = \kappa^* \geq \kappa^C(\alpha^C)$.

This result implies that all lower censorship policies with $\alpha^C \leq \bar{\alpha}^C$ are effective, and that further the platform's innovation threshold is insensitive to α^C over this range. It is therefore without loss to confine attention to effective policies satisfying $\alpha^C \geq \bar{\alpha}^C$. For this range of censorship thresholds, we have $\kappa^* \leq \kappa^C(\alpha^C)$ and therefore $\bar{U}'_P(\kappa; \mathcal{D}^C(\alpha^C)) = \bar{U}'_P(\kappa; \mathcal{D}^{LF}) < 0$ for all $\kappa > \kappa^C(\alpha^C)$. It follows that $\kappa^*(\mathcal{D}^C(\alpha^C)) = \kappa^C(\alpha^C)$ for all effective policies satisfying $\alpha^C \geq \bar{\alpha}^C$. Since $\kappa^C(\alpha^C)$ is increasing in α^C, one optimal data policy must therefore be the lower censorship policy with the largest censorship threshold for which the policy remains effective.

For every $\alpha^C \geq \bar{\alpha}^C$, we have $\kappa^* \leq \kappa^C(\alpha^C)$ and $\bar{U}'_{P,+}(\kappa^C(\alpha^C); \mathcal{D}^C(\alpha^C)) = \bar{U}'_P(\kappa^C(\alpha^C); \mathcal{D}^{LF}) \leq 0$. Hence a policy with censorship threshold $\alpha^C \geq \bar{\alpha}^C$ is effective if and only if $\bar{U}'_{P,-}(\kappa^C(\alpha^C); \mathcal{D}^C(\alpha^C)) \geq 0$. This inequality may be equivalently written $\widehat{U}'_P(\kappa^C(\alpha^C)) \geq 0$, where

$$\widehat{U}'_P(\kappa) \equiv \frac{G'(\kappa)}{G(\kappa)} \left(\mu_E^M + (\mu_E^M - \kappa)\Lambda(\hat{\alpha}(\kappa))\right) - \Lambda(\underline{\alpha}^*(\kappa)).$$

The function $\widehat{U}'_P$ is continuous in κ. Since $C'(\kappa; \mathcal{D}^{LF}) = 1/(|\Delta\mu_E|\alpha^*(\kappa))$ and α^* is increasing, we have that $C(\cdot; \mathcal{D}^{LF})$ is strictly convex and therefore $\underline{\alpha}^*$ is increasing. Meanwhile, the first term of $\widehat{U}'_P(\kappa)$ is decreasing whenever $\Lambda(\alpha^*(\kappa)) > 0$. Strict convexity of $C(\cdot; \mathcal{D}^{LF})$ additionally that

$$\frac{C(\kappa; \mathcal{D}^{LF}) - 1}{\kappa - \mu_E^D} = \frac{C(\kappa; \mathcal{D}^{LF}) - C(\mu_E^D; \mathcal{D}^{LF})}{\kappa - \mu_E^D} < C'(\kappa; \mathcal{D}^{LF})$$

for all $\kappa > \mu_E^D$. Hence $\alpha^*(\kappa) > \underline{\alpha}^*(\kappa)$ for all $\kappa \geq \kappa^*$. Thus in particular $\widehat{U}'_P(\kappa)$ is decreasing at any $\kappa \geq \kappa^*$ such that $\Lambda(\underline{\alpha}^*(\kappa)) > 0$. It follows that for every $\kappa \in [\kappa^*, \mu_E^M)$, either $\widehat{U}'_P$ is decreasing in κ or it is positive. And at μ_E^M, we have $\widehat{U}'_P(\mu_E^M) = -\Lambda(\underline{\alpha}^*(\mu_E^M)) = -\Lambda(1) < 0$. So $\widehat{U}'_P$ crosses zero at most once, and any crossing is from above. Finally, $\alpha^*(\kappa^*) > \underline{\alpha}^*(\kappa^*)$ implies that $\widehat{U}'_P(\kappa^C(\bar{\alpha}^C)) = \widehat{U}'_P(\kappa^*) > \bar{U}'_P(\kappa^*; \mathcal{D}^{LF}) = 0$. Thus there exists a unique $\bar{\kappa}^* \in (\kappa^*, \mu_E^M)$ such that $\text{sign}(\widehat{U}'_P(\kappa)) = \text{sign}\,(\bar{\kappa}^* - \kappa)$ for $\kappa \in [\kappa^*, \mu_E^M]$.

An immediate implication of the construction of $\bar{\kappa}^*$ is that $\alpha^{C*} \equiv \alpha^*(\bar{\kappa}^*)$ is the largest censorship threshold for which $\mathcal{D}^{LF}(\alpha^C)$ is effective. This policy must therefore be optimal. Since the policy is effective, it satisfies $\kappa^*(\mathcal{D}^C(\alpha^{C*})) = \kappa^C(\alpha^{C*}) = \bar{\kappa}^*$. Optimality then implies that $\bar{\kappa}^* \geq \kappa^*(\mathcal{D})$ for any data privacy policy $\mathcal{D}$. Additionally, effectiveness implies that the platform optimally imitates immediately in all states above α^{C*} and never in states below α^{C*}, in order to implement the innovation threshold $\kappa^C(\alpha^{C*})$.

F.14 Proof of Proposition 6

For any data policy $\mathcal{D}$, the platform's expected profits may be written

$$U_P(\kappa, f; \mathcal{D}) = G(\kappa)\left(\mu_P^M(f)(1 - A^*(\kappa, f)) + \mu_P^D(f)A^*(\kappa, f) - k_P C(\kappa; f, \mathcal{D})\right),$$

where $A^*(\kappa, f) \equiv (\mu_E^M(f) - \kappa)/|\Delta\mu_E(f)|$. Let

$$\Delta U_P(\kappa, f) \equiv U_P(\kappa, f; \mathcal{D}) - U_P(\kappa, f; \mathcal{D}^{LF}).$$

Then we may compute

$$\Delta U_P(\kappa, f) = G(\kappa)k_P\left(C(\kappa; f, \mathcal{D}^{LF}) - C(\kappa; f, \mathcal{D})\right).$$

Next, define $\widetilde{C}(A; \mathcal{D})$ to be the value of the problem

$$\min_{T \in \mathcal{T}} \mathbb{E}[e^{-rT}] \quad \text{s.t.} \quad \mathbb{E}[\alpha e^{-rT}] = A$$

for $A \in [0, 1]$. Given the participation constraint

$$\kappa = \mu_E^M(f) + \Delta\mu_E(f)A^*(\kappa, f),$$

the cost function C is related to $\widetilde{C}$ via $C(\kappa; f, \mathcal{D}) = \widetilde{C}(A^*(\kappa, f); \mathcal{D})$. So we may equivalently write $\Delta U_P(\kappa, f)$ as

$$\Delta U_P(\kappa, f) = G(\kappa)k_P\left(\widetilde{C}(A^*(\kappa, f); \mathcal{D}^{LF}) - \widetilde{C}(A^*(\kappa, f); \mathcal{D})\right).$$

C and $\widetilde{C}$ additionally satisfy the inverse relationship $\widetilde{C}(A;\mathcal{D}) = C(\mu_E^M(f) + \Delta\mu_E(f)A; f, \mathcal{D})$. Hence

$$\widetilde{C}'_\pm(A^*(\kappa, f); \mathcal{D}) = C'_\pm(\kappa; f, \mathcal{D})\Delta\mu_E(f) = \left[\alpha^*_\pm(\kappa, f; \mathcal{D})\right]^{-1}.$$

So we may express the derivative of $\Delta U_P(\kappa, f)$ wrt f as

$$\frac{\partial}{\partial f_\pm}\Delta U_P(\kappa, f) = G(\kappa)k_P\left(\frac{1}{\alpha^*_\pm(\kappa, f)} - \frac{1}{\alpha^*_\pm(\kappa, f; \mathcal{D})}\right)\frac{\partial A^*}{\partial f}(\kappa, f).$$

Meanwhile, the derivative of A^* wrt f may be written

$$\frac{\partial A^*}{\partial f}(\kappa, f) = \frac{1}{\Delta\mu_E(f)}\left((1 - A^*(\kappa, f))\frac{d\mu_E^M}{df} + A^*(\kappa, f)\frac{d\mu_E^D}{df}\right).$$

Since $d\mu_E^s/df < 0$ under each market structure, we have $\partial A^*/\partial f < 0$. Hence

$$\operatorname{sign}\left(\frac{\partial}{\partial f_\pm}\Delta U_P(\kappa, f)\right) = \operatorname{sign}\left(\alpha^*_\pm(\kappa, f) - \alpha^*_\pm(\kappa, f; \mathcal{D})\right).$$

At $f = f^*(\kappa)$, it must be that $\frac{\partial}{\partial f_+}U_P(\kappa, f; \mathcal{D}^{LF}) \leq 0$. Therefore if $\alpha^*_+(\kappa, f) \leq \alpha^*_+(\kappa, f; \mathcal{D})$, then $\frac{\partial}{\partial f_+}U_P(\kappa, f; \mathcal{D}) \leq 0$. Quasiconcavity and uniqueness of the maximizer $f^*(\kappa; \mathcal{D})$ then implies that $f^*(\kappa; \mathcal{D}) \leq f^*(\kappa)$. A similar argument in the other direction establishes that if $\alpha^*_-(\kappa, f) \geq \alpha^*_-(\kappa, f; \mathcal{D})$, then $f^*(\kappa; \mathcal{D}) \geq f^*(\kappa)$.

If both $C'(\kappa, f^*(\kappa))$ and $C'(\kappa, f^*(\kappa); \mathcal{D})$ exist, then the reasoning of the previous paragraph may be strengthened. In particular, it must be that $\frac{\partial}{\partial f}U_P(\kappa, f; \mathcal{D}^{LF}) = 0$ at $f = f^*(\kappa)$, in which case $\alpha^*(\kappa, f) < \alpha^*(\kappa, f; \mathcal{D})$ implies that $\frac{\partial}{\partial f}U_P(\kappa, f; \mathcal{D}) < 0$ and therefore (supposing that $f^*(\kappa)$ is interior) $f^*(\kappa; \mathcal{D}) < f^*(\kappa)$. Similar reasoning applies in the other direction.

F.15 Proof of Lemma 8

Fix a data privacy policy $\mathcal{D}$. Since it is held fixed throughout this proof, we suppress it in all notation. Let $\widetilde{C}(A)$ be as defined in the proof of Proposition 6. Recall that C is related to $\widetilde{C}$ via $C(\kappa; f) = \widetilde{C}(A^*(\kappa, f))$, where $A^*(\kappa, f) \equiv (\mu_E^M(f) - \kappa)/|\Delta\mu_E(f)|$. as well as the inverse relationship $\widetilde{C}(A) = C(\mu_E^M(f) + \Delta\mu_E(f)A; f)$. The latter identity implies that $\widetilde{C}$ is convex given that C is.

Using the relationship $C(\kappa; f) = \widetilde{C}(A^*(\kappa, f))$, we may compute

$$C'_\pm(\kappa; f) = \frac{\widetilde{C}'_\pm(A^*(\kappa, f))}{\Delta\mu_E(f)},$$

allowing the marginal demand state to be written

$$\alpha^*_\pm(\kappa; f) = \left[\widetilde{C}'_\pm\left(A^*(\kappa, f)\right)\right]^{-1}.$$

Now, the definition of $A^*(\kappa, f)$ may be rearranged to read

$$\kappa = \mu_E^M(f)(1 - A^*(\kappa, f)) + \mu_E^D(f)A^*(\kappa, f).$$

Since μ_E^M and μ_E^D are each decreasing in f and $\Delta\mu_E(f) > 0$, this identity can be satisfied for all f only if $A^*(\kappa, f)$ is decreasing in f. Meanwhile, convexity of $\widetilde{C}$ implies that $\widetilde{C}'_\pm(A)$ are each nondecreasing in A. Therefore $\alpha^*_\pm(\kappa; f)$ are nondecreasing in f.