

Platform Traps^{*}

Andrei Hagiu[†] and Julian Wright[‡]

August 1, 2026

Abstract

A novel feature of platforms such as marketplaces and social networks is that non-participants may become worse off as others join. We show that in such settings, rational agents can be induced to join a platform despite being better off without it — a phenomenon we call a platform trap. We provide a theory of how platform traps emerge through one or more of the following additional features: (i) the ability of the platform to make dynamic price adjustments; (ii) the interplay between on-platform and off-platform negative externalities; (iii) favorable equilibrium selection in settings where participation decisions admit multiple equilibria.

1 Introduction

Many businesses increasingly rely on major platforms to reach customers, including third-party sellers on Amazon, hotels on Booking.com and Expedia, and restaurants on DoorDash and Uber Eats. As consumers shift to these platforms rather than shopping directly with sellers, concerns arise about sellers' growing dependence on, and exploitation by, platforms. This can involve higher or new types of fees, as well as an increasing need to advertise to be discovered on the platforms they depend on. But why would businesses join such platforms in the first place, knowing they will ultimately be held up this way?

^{*}We would like to thank seminar participants at Boston University, KU Leuven and Peking University, and conference participants at the Canon Institute for Global Studies workshop in Tokyo, MaCCI/EPoS Workshop on Digital Markets at the University of Mannheim, Melbourne IO day, SIDE Workshop, UIBE Workshop on Industrial Organization and Competition Policy, and in particular, Marshall Van Alstyne, Alexandre de Cornière, Lukas Hensel, Johannes Johnen, Bruno Jullien, Jan De Loecker, Volker Nocke, Martin Peitz, Markus Reisinger, Florian Schuett, Tat-How Teh, Frank Verboven and Allen Vong.

[†]Boston University Questrom School of Business, E-mail: ahagiu@bu.edu

[‡]Department of Economics, National University of Singapore, Singapore, E-mail: jwright@nus.edu.sg

In this paper, we provide a theory of why rational agents join a platform even when they end up worse off as a result. We call this a platform trap. The key departure from standard models of platforms is that each agent’s relevant outside option is not constant but deteriorates as each agent joins the platform. The theory applies to sellers participating on a marketplace that attracts consumers who previously searched directly for them, and to other settings: neighborhood shops opening outlets in suburban malls or brands supplying Walmart (which draws traffic away from their traditional channels); individuals joining an online social network rather than relying on direct connections (those remaining outside risk social exclusion); and news organizations signing deals with Facebook and Google (cannibalizing traffic on their own websites as users increasingly access news through these platforms).

Because each agent’s participation imposes a negative externality on other agents’ outside options, agents face a collective action problem: each finds it individually attractive to join but doesn’t take into account the resulting degradation of other agents’ outside option. The platform exploits this through higher prices, which can ultimately leave all agents worse off. This differs from the standard setting in which a monopoly platform extracts an increasing share of the value created as users join, due to positive network effects. There, the outside option is fixed, so the platform cannot leave users strictly worse off from joining. Moreover, in our setting the platform trap can arise absent network effects, including when on-platform externalities are negative.

Our theory has a finite number of strategic agents, such as sellers on a marketplace, who receive offers sequentially and decide whether to join, taking into account how their decision may affect later agents’ decisions. Both the benefit of joining the platform and the outside option can depend on how many agents join. The benefit of the outside option is weakly decreasing in the number of agents participating on the platform and strictly decreasing at some point. The difference between the benefit on the platform and the outside option is weakly increasing as more agents join. We microfound these properties in a marketplace model. As sellers join, they draw buyers from the direct channel to the marketplace and worsen the position of all sellers that are not on the marketplace. We provide conditions under which such a marketplace attracts sellers even though all sellers end up worse off; in some cases, all buyers are also worse off and total welfare is lower than without the platform.

Sequential participation decisions raise the possibility that some agents are pivotal: each of their decisions affects whether later agents can be profitably induced to participate. In this case, the platform must first attract a critical mass of participants before it can offer each later agent more value than the outside option. Once the platform

achieves this critical mass, each later agent knows that, regardless of its own decision, the platform can induce all remaining agents to join by adjusting prices, so it is willing to join even when charged a price based on all other agents joining.

In both the microfounded marketplace model and the more general setting that encompasses it, we derive equilibrium prices for the strategic agents and a sufficient condition for a platform trap. When no agents are pivotal, the platform extracts maximum surplus by setting a single price that leaves each agent with the same low surplus as its outside option when all other agents participate. At the other extreme, when all agents are pivotal and the platform provides enough value, it attracts all agents with increasing prices, leaving each agent with the outside option it would obtain when only earlier agents join the platform. In intermediate cases, the platform raises prices until the marginal joining agent is no longer pivotal, and then charges the same maximal price to all remaining agents as when no agents are pivotal.

In this baseline setting, whenever agents participate, all agents would be weakly better off without the platform, and at least some will be strictly better off. Agents would therefore be better off collectively committing to boycott the platform. Furthermore, since the platform may need to offer negative prices to attract the first pivotal agents, it may profitably attract all agents even when doing so is inefficient. In this case, both agent surplus and total surplus can be lower because the platform exists. Yet because the platform also attracts agents when doing so is efficient, the platform-trap problem is distinct from traditional concerns about monopoly platforms causing inefficiency, with neither implying the other.

The baseline logic suggests that dynamic price adjustments are essential for a platform trap. Surprisingly, even a platform restricted to a single price can generate a trap if each agent’s platform benefit is single-peaked and peaks before all agents participate.¹ The price is based on this maximum benefit, and if agents’ outside option falls fast enough with participation, all agents join and would be better off without the platform. Thus, negative externalities on and off the platform can generate a platform trap even without dynamic price adjustments.

We also examine factors such as simultaneous decisions, imperfect information, negotiated offers, heterogeneous agents and competing platforms, which can amplify or moderate the trap. Multiple equilibria can arise when agents move simultaneously at a constant price or do not observe others’ offers and decisions. With equilibrium selection favorable to the platform, the trap arises in a wider set of cases than in the baseline because agents can no longer be pivotal when their decisions are unobserved.

¹We show how this possibility can arise when a marketplace platform attracts competing sellers.

With unfavorable equilibrium selection, a platform trap can still arise, but only if on-platform externalities are negative. When the platform negotiates with each agent rather than making take-it-or-leave-it offers, the model can generate increasing price dynamics even without pivotal agents. Prices increase because each joining agent endogenously weakens the bargaining position of remaining agents. When one agent is larger than all others (a “superstar” agent), the platform must decide whether to attract that agent first or last. The optimal choice depends crucially on the curvature of the externality on the outside option. Finally, platform competition can mitigate but need not eliminate the trap.

2 Literature review

A platform trap is closely related to the collective traps introduced and examined by Bursztyn et al. (2025a). They provide experimental evidence that total consumer welfare from the availability of Instagram and TikTok is negative in their student sample, even though these individuals are willing to pay to use them. They propose a mechanism in which users impose a negative externality on non-users, for example through social exclusion or fear of missing out. As more people join, marginal users may participate to avoid worse non-participation outcomes, even if they experience negative overall utility from the platform. This mechanism is consistent with ours: the outside option weakens as platform participation increases. Bursztyn et al. offer a conceptual framework to interpret their results, but do not formalize how, or under what conditions, a collective trap can emerge through individual choices, nor how platforms might engineer such traps in a dynamic setting. Their focus is instead empirical: demonstrating and quantifying the collective trap rather than providing a theory of it.²

On the theory side, our paper relates to settings where a principal (the platform) contracts with multiple agents subject to cross-agent externalities. Segal (1999) provides the most general formulation, encompassing many applications. One application is Katz and Shapiro (1986), where a sponsored technology (the principal) licenses its technology and competes for users with an unsponsored technology (the outside option), with both technologies subject to network effects. Another application is Segal and Whinston (2000), where an incumbent signs exclusive contracts with agents, deny-

²Sullivan (2024) provides further empirical evidence consistent with a platform trap. In a counterfactual simulation, he finds that abolishing food delivery platforms in the United States would increase restaurants’ profits.

ing an entrant the scale needed to enter and harming agents who do not contract with the incumbent.³ These settings, including Segal’s general framework, differ from ours in various modeling choices, notably their focus on simultaneous agent decisions rather than our sequential baseline. But more fundamentally, they differ in their research question. They analyze inefficiencies in contracts offered by the principal, whereas we ask whether agents would be better off without the principal. For example, Segal and Whinston ask whether the incumbent’s exclusive contracts lead to inefficient outcomes, not whether agents would be better off without the incumbent. Indeed, agents would never be better off without the incumbent in their setting, since the entrant, assumed to have the same costs, would then become the monopolist. Thus, there is no platform trap in Segal and Whinston.

Our paper also connects to the literature on data externalities. When more users share personal data, a platform can infer information about users who do not share, or share less. This weakens the outside option of non-participation or limited data disclosure, allowing the platform to collect and monetize more data while providing less compensation for the resulting privacy nuisance (Choi et al., 2019). More broadly, this is an example of social data: individual data are informative about others, creating data externalities that affect the terms on which platforms acquire and use personal information (Bergemann et al., 2022).

A platform trap also differs fundamentally from inefficiencies in the classic literature on network externalities, such as Farrell and Saloner (1985) and Katz and Shapiro (1985). This literature shows that users can become locked into an inferior technology, as in the classic QWERTY example, or inefficiently adopt a new technology due to excess momentum (Farrell and Saloner, 1986; Katz and Shapiro, 1986). Even if only one technology is sponsored, analogous to the platform in our setting, these papers ask whether adoption is inefficient, not whether agents would be better off without the sponsored technology. As we show, whether a platform’s existence creates an inefficiency is distinct from whether it makes agents worse off; our setting permits scenarios where either, both, or neither occur.

Within the platform literature, most papers do not study pricing dynamics. Jullien

³Their setting differs in several respects: there are no on-platform externalities, and the incumbent can extract surplus from agents who do not sign up if it signs up enough agents exclusively to deter entry by the rival platform. This structure explains why they cannot obtain our result that the platform can trap all agents even under perfectly coalition-proof Nash equilibria with a simultaneous constant price offer when on-platform externalities are negative, or our result that it can be efficient, from a total surplus perspective, for agents to sign up even though they are necessarily worse off. Nevertheless, the notion of pivotal agents in the sequential version of their model is the same as in ours.

et al. (2021) survey some that do. Among them, Biglaiser et al. (2022) is arguably closest to our approach given their agents also make sequential participation decisions. But their focus differs. They study agents moving from an incumbent to an entrant platform to explain incumbency advantage. Our key contribution to platform pricing dynamics is to show how increasing prices can emerge from pivotal agents or from the platform’s endogenously improving bargaining position. These effects are absent in standard platform settings where agents’ outside option remains constant. A further novelty relative to platform models with network effects is that the optimal order of attracting agents should depend on the externality on other agents’ outside option.

Finally, a literature shows how platforms intensify competition between sellers relative to a direct channel when sellers set the same price across both channels (e.g., Baye and Morgan, 2001; Wang and Wright, 2020; and Gomes and Mantovani, 2025). In these settings, sellers may be collectively better off without the platform because it lowers their margins on both channels by intensifying competition. Our mechanism does not rely on same-price constraints across channels or on the commoditization of sellers that underpins seller harm in these settings. Moreover, these works rule out the dynamic decision-making that explains why sellers join despite ending up worse off, one of our key contributions.

3 Model

This section provides a fully worked out microfoundation of a platform trap arising for a two-sided marketplace platform. The agents that make strategic decisions are sellers. Buyers play a passive role, but help determine the payoffs for sellers. For this reason, the microfounded model below naturally maps to a more general version of the model that can also represent a one-sided platform (e.g., a social network), as Section 5 explains.

We assume there is a single platform, $N \geq 2$ independent sellers (local monopolists), and a continuum of measure one of buyers. Each buyer wants to buy from one independent seller, drawn uniformly across the N sellers. We will discuss the case with competing sellers below.

Buyers and sellers can participate on one of two channels: the platform or a non-strategic alternative, which we call the direct channel.⁴ Outside these channels, buyers and sellers have a fixed outside option normalized to zero. Let π_j be a seller’s profit per

⁴In an earlier version of the paper, we analyzed the case where sellers kept their direct channel when joining the platform, which made it easier for the platform to engineer a platform trap.

matched buyer on channel j ($j = D$ for the direct channel and $j = P$ for the platform) and let v_j be the buyer's indirect utility from purchasing from her matched seller. These primitives are consistent with many types of buyer-seller interactions on a given channel. For concreteness, we focus on a standard monopoly pricing interpretation with elastic demand. Suppose a buyer's willingness-to-pay for quantity q of her matched seller's product is $w_j(q)$ on channel j , implying a seller's demand from each matched buyer is $q_j(p)$ on channel j at the seller's price p . Then, if all sellers have marginal cost c and $p_j \equiv \arg \max_p \{(p - c) q_j(p)\}$, we have $\pi_j = (p_j - c) q_j(p_j)$ and $v_j = w_j(q_j(p_j)) - p_j q_j(p_j)$. We treat v_P , v_D , π_P and π_D as the model primitives in what follows.

Buyers differ in their cost of using each channel. This could be a one-time search, travel, or general mismatch cost. Once they incur this cost, they have access to all sellers present on that channel. Specifically, we assume each buyer draws a cost $s_j \in [\underline{s}_j, \bar{s}_j]$ of using channel $j \in \{P, D\}$. We allow $\underline{s}_j < 0$, so one or both channels may give some buyers fixed benefits rather than costs. The draws of s_D and s_P need not be independent. Rather, we only assume the difference in draws, $s = s_D - s_P$, has an induced cumulative distribution function G with support equal to a non-degenerate subinterval of $[\underline{s}_D - \bar{s}_P, \bar{s}_D - \underline{s}_P]$. We assume $v_j \geq \bar{s}_j - \underline{s}_k$ for $k \neq j$, so a buyer's maximum possible additional cost of using a channel is always less than their match value. We also assume all buyers prefer going to at least one of the channels rather than not going to either (and getting a zero payoff), regardless of how sellers are split among the two channels — a sufficient condition for this is provided below.

The game consists of a single period, so payoffs are only realized once. Within the period, there are many stages. In stage k of the first N stages, the platform offers seller $k \in \{1, \dots, N\}$ a fixed price (or fee) P^k to join the platform, which seller k accepts or rejects. Since sellers are homogeneous, their order does not matter. Each seller, in sequence, observes the platform's previous offers and the previous sellers' channel decisions, and makes its decision upon receiving its offer.⁵ After all sellers have made their channel decisions, buyers observe their draws of s_D and s_P , as well as the number of sellers on each channel (but don't observe which specific sellers are on each channel), and simultaneously decide which channel to use. As is the case on most marketplaces, we assume the platform does not charge buyers for access. Sellers set their prices simultaneously with buyers' channel choices. Finally, buyers observe all the sellers and their prices in their chosen channel, and make purchasing decisions. Payoffs are then realized. We assume throughout that if a seller is indifferent between

⁵As long as a seller can observe the previous sellers' decisions, the platforms' previous prices are irrelevant; we assume they are observed so that the game is one of perfect (and complete) information.

joining the platform and not joining, the seller joins the platform.

This timing implies that only sellers are strategic: each decides sequentially whether to join, anticipating later sellers’ responses and buyers’ eventual channel choices. Buyers, by contrast, are non-strategic decision makers, choosing their preferred channel after seller participation is determined. We look for a subgame perfect equilibrium of this game.

3.1 Discussion of model assumptions

The multi-stage single-period nature of our model can be interpreted as capturing sellers making irreversible participation decisions over time on the platform. The payoffs for buyers and sellers then represent the present discounted value of future payoffs, over many (possibly infinite) periods. The absence of discounting across the sellers’ decisions reflects that joining decisions occur over a short horizon relative to the longer horizon over which payoffs arise. Supplementary Appendix A provides an infinite period version of the model with discounting to illustrate how a very similar platform trap can still arise, although the analysis becomes much more complicated.

In the baseline setting, the platform can set a different price (or fee) for each seller, including negative prices if necessary. Section 6.1 considers the other extreme in which the platform cannot price discriminate, showing that whether a platform trap arises depends on additional payoff assumptions. Our timing implies the realistic feature that when making an offer to a given seller, the platform cannot commit to the prices it will charge to subsequent sellers. This shows that the platform trap does not depend on commitment power. Indeed, committing to a fixed sequence of prices independent of sellers’ participation decisions is often worse for the platform than retaining the flexibility to adjust its offers, although full commitment to participation-contingent pricing always maximizes the platform’s payoffs. The role of commitment is analyzed in Supplementary Appendix B.

In the baseline setting, the strategic agents (the sellers) decide sequentially whether to join with full information, ensuring equilibrium participation and payoffs are unique. This rules out platform traps arising merely from “bad” equilibrium selection by sellers. If instead multiple sellers decide simultaneously whether to participate on the platform (either based on private offers, or a single price offered to them all), multiple equilibria can emerge in their decision stage. In that case, whether a platform trap arises depends on equilibrium selection and the nature of sellers’ payoffs. Section 6.2 addresses this.

3.2 Preliminaries

Since buyers do not know whether their matched seller is on a particular channel until they choose that channel, they care about how many sellers they can reach on each channel. The more sellers that participate on the platform, the greater the payoff a buyer gets from going to the platform relative to the direct channel, since the buyer is more likely to be able to find her preferred seller there. Specifically, if n sellers participate on the platform, a buyer chooses it if their expected utility $\frac{n}{N}v_P - s_P$ exceeds their expected direct channel utility $\frac{N-n}{N}v_D - s_D$.⁶

Let $m(n)$ be the measure of buyers choosing the platform, so $1 - m(n)$ choose the direct channel:

$$m(n) = 1 - G\left(\frac{N-n}{N}v_D - \frac{n}{N}v_P\right). \quad (1)$$

Since $G(\cdot)$ is increasing over a subinterval of $[\underline{s}_D - \bar{s}_P, \bar{s}_D - \underline{s}_P]$, $m(n)$ is weakly increasing in n , with $m(0) = 0$ (since $v_D \geq \bar{s}_D - \underline{s}_P$) and $m(N) = 1$ (since $v_P \geq \bar{s}_P - \underline{s}_D$). We make the weak assumption that

$$m(N-1) > 0, \quad (2)$$

or equivalently, that $G(\frac{1}{N}v_D - \frac{N-1}{N}v_P) < 1$. This says that when all but one seller joins the platform, it attracts a positive measure of buyers. This holds for high enough N , reflecting that $G(-v_P) = 0$ given $v_P \geq \bar{s}_P - \underline{s}_D$.

Let $u(n) = \frac{1}{N}(1 - m(n))\pi_D$ be a seller's expected gross profit from selling on the direct channel and $b(n) = \frac{1}{N}m(n)\pi_P$ its expected gross profit from selling on the platform, when n sellers in total join the platform. These are gross profits since they are calculated before considering any platform prices. From above, $b(\cdot)$ is weakly increasing in n and strictly increasing for at least some n , capturing positive indirect network effects: sellers benefit more from joining when more sellers join because this attracts more buyers. By the same token, $u(n)$ is weakly decreasing in n and strictly decreasing for at least some n . Since $u(0) = \frac{1}{N}(1 - m(0))\pi_D = \frac{\pi_D}{N} > 0$, condition (2) implies $u(N-1) < u(0)$: a seller that stays on the direct channel is strictly worse off if all other sellers join the platform.

⁶A sufficient condition for our assumption that all buyers prefer one of the two channels to the zero outside option is

$$\min_{0 \leq n \leq N} \left\{ \max \left\{ \frac{n}{N}v_P - \bar{s}_P, \frac{N-n}{N}v_D - \bar{s}_D \right\} \right\} \geq 0.$$

This implies $v_D \geq \bar{s}_D$ (for $n = 0$) and $v_P \geq \bar{s}_P$ (for $n = N$).

Define the marginal surplus function

$$\Delta(n+1) \equiv b(n+1) - u(n),$$

which measures the incremental payoff to a seller from joining the platform rather than selling directly when n other sellers join. It determines a seller's willingness to pay for access and will be central to our analysis. Clearly, $\Delta(n+1) \geq \Delta(n)$ for all relevant n , with strict inequality for at least some n . This means the marginal seller's incentive to join the platform weakly increases with the number of participating sellers, implying (as we will show) that in equilibrium, either all sellers join the platform or none do.

We also assume $\Delta(N) > 0$, which, since $m(N) = 1$, is equivalent to

$$\pi_P > (1 - m(N-1)) \pi_D.$$

Thus, profit per buyer on the direct channel cannot be so high relative to profit per buyer on the platform that a seller prefers to remain alone on the direct channel when all other sellers join the platform. This is the weakest assumption ensuring that joining the platform yields strictly greater gross profit in that situation. Otherwise, if $\Delta(N) \leq 0$, the platform can never profitably attract sellers. Indeed, we assume that the platform must earn positive profits to operate.

Finally, in an equilibrium in which all N sellers join at the platform prices $(P^1, \dots, P^N)$, seller k would be weakly better off without the platform if and only if $b(N) - P^k \leq u(0)$.⁷

4 Analysis

A key step to characterize platform pricing and strategic agents' (i.e., sellers') decisions is to recognize that some sellers can be pivotal. Specifically, $k_0 \in \{0, \dots, N-1\}$ sellers are pivotal in the original game if the non-participation of any of the first k_0 sellers prevents the platform from profitably inducing any of the N sellers to participate.⁸

To build intuition, consider two extreme cases: no seller is pivotal in the original game, or every seller is pivotal.⁹ For example, with two sellers ($N = 2$), suppose the

⁷If the platform doesn't exist, all sellers use the direct channel, which attracts all buyers there given $v_D - \bar{s}_D > 0$. Thus, each seller obtains $u(0)$.

⁸In the proof of Proposition 1 below, we will show that either the first k_0 agents in the sequence are pivotal for some $1 \leq k_0 \leq N-1$, or no agents are pivotal.

⁹When we say "every seller is pivotal", we mean all sellers are pivotal except the last to decide,

platform needs both to join for it to offer each seller strictly higher expected gross profit than the direct channel, i.e., $\Delta(2) > 0 \geq \Delta(1)$. Then the first seller to receive an offer is pivotal: if (and only if) it rejects, the second seller cannot be profitably attracted. Thus, the maximum price the platform can charge the first seller is $b(2) - u(0)$, and the maximum profit it can extract in total is $2b(2) - u(0) - u(1)$. Now consider three sellers ($N = 3$) under the same assumptions, and suppose in addition $2b(2) > u(0) + u(1)$. The first seller is no longer pivotal, because even if it rejects, the platform can still profitably attract the remaining two. If the first seller accepts, the second seller is not pivotal because $\Delta(2) > 0$ implies the platform will attract the third seller regardless of the second seller's decision. However, if the first seller rejects, the situation reduces to the $N = 2$ case, making the second seller pivotal even though it is not pivotal in the original game.

For general N , if $\Delta(1) > 0$, no seller is ever pivotal because $\Delta(\cdot)$ weakly increasing implies $\Delta(n) > 0$ for all $n \geq 1$. The platform then creates higher expected gross profit for a seller than the direct channel, regardless of how many other sellers are expected to join. As a result, the platform can always profitably induce any seller to participate. Using the payoffs in Section 3.2, $\Delta(1) = \frac{1}{N}(m(1)\pi_P - \pi_D)$, which can be positive or negative in our setting.

When $\Delta(1) > 0$, the last seller can be induced to join even if no one else has joined. Since $\Delta(n)$ is weakly increasing, each seller expects all subsequent sellers to join regardless of its own decision. Thus, if $N' < N$ prior sellers have joined, the k -th seller will accept the platform's offer if and only if the price P satisfies

$$P \leq P(N') = \Delta(N' + N - k + 1).$$

A higher price causes rejection, eliminating positive revenue from that seller and reducing revenue from subsequent sellers. By setting $P = P(N')$ at each stage, the platform ensures that all N sellers join. Along the equilibrium path, $N' = k - 1$, so the common price is $P^* = \Delta(N)$.

This is the same platform price that would arise if all sellers decided simultaneously and coordinated on the equilibrium most favorable to the platform (i.e., where all sellers join whenever that is an equilibrium). In that case, each seller expects the other $N - 1$ sellers to join, yielding a payoff $b(N) - P$ if they also join, and $u(N - 1)$ if they do not. Setting $P = \Delta(N)$ makes each seller indifferent between joining and not, when they expect all others to participate.

which by definition can't be pivotal. Likewise for agents later on.

For any given $N > 1$, $\Delta(1) > 0$ is a sufficient condition for no seller to be pivotal. More generally, no seller is pivotal if and only if the platform can still profitably attract all $N - 1$ remaining sellers after the first seller rejects. Obtaining a simple closed-form condition for this in terms of the payoff primitives is difficult for general $N > 2$, as we discuss after Proposition 3 below.

Now consider the opposite extreme, where every seller is pivotal. Specifically, suppose

$$\Delta(N - 1) = b(N - 1) - u(N - 2) \leq 0 < \Delta(N),$$

which, given that $\Delta(\cdot)$ is weakly increasing, implies $\Delta(k) \leq 0$ for all $k \leq N - 1$. In this case, losing even one seller prevents the platform from profitably attracting any remaining seller: the last seller would prefer the direct channel at any positive price, so the platform has no incentive to attract it, and it cannot attract the second-to-last seller either, and so on. Thus, every seller is pivotal. To induce participation, the platform must offer each seller at least its expected direct-channel profit assuming no subsequent sellers join. Although all sellers join in equilibrium, the price increases along the sequence, from $P^1 = b(N) - u(0)$ for the first seller to $P^N = b(N) - u(N - 1) = \Delta(N)$ for the last. This is profitable for the platform if $b(N) > \frac{1}{N} \sum_{k=1}^N u(k - 1)$. In the marketplace setting, this condition is equivalent to $\pi_P > \sum_{k=1}^N \frac{1}{N} (1 - m(k - 1)) \pi_D$.

Between the two extremes, i.e. when $\Delta(1) \leq 0 < \Delta(N - 1)$, some but not all sellers may be pivotal. Specifically, the proof of the next proposition shows that there is a threshold $k_0 \geq 1$ such that the first $k \leq k_0$ sellers are pivotal and are charged $P^k = b(N) - u(k - 1)$. Once these sellers are on board, the remaining $N - k_0$ sellers are no longer pivotal and are each charged $P^k = b(N) - u(N - 1) = \Delta(N)$. These prices are consistent with the extreme cases above. A scenario with $1 \leq k_0 \leq N - 1$ can be interpreted as one in which the platform must first attract a critical mass k_0 of sellers before it can offer more expected profit than the direct channel to each subsequent seller irrespective of whether these sellers ultimately join. The proof considers a general reduced-form game of which the two-sided marketplace setting is a special case, and uses induction over games with increasing values of N . As with all proofs not covered in the text, it is relegated to the Appendix.

Proposition 1. *If no sellers are pivotal, the platform optimally attracts all N sellers with the constant price $P^k = \Delta(N) > 0$ for all $k \in \{1, \dots, N\}$. Otherwise, there exists a threshold seller $k_0 \in \{1, \dots, N - 1\}$ such that the profit-maximizing prices that induce*

participation on the platform are

$$P^k = \begin{cases} b(N) - u(k-1) & \text{if } 1 \leq k \leq k_0 \\ b(N) - u(N-1) & \text{if } k_0 + 1 \leq k \leq N \end{cases}. \quad (3)$$

If $\sum_{k=1}^N P^k > 0$, these are the platform's optimal prices and all N sellers join, with the first k_0 sellers pivotal. If instead $\sum_{k=1}^N P^k \leq 0$, the platform cannot profitably attract any sellers, so no one joins in equilibrium. In any equilibrium where sellers join, they would all be weakly better off without the platform, and at least some would be strictly better off. Buyers, however, may be better or worse off without the platform, depending on parameters. Finally, a sufficient condition for all sellers to join is

$$b(N) > \frac{1}{N} \sum_{k=1}^N u(k-1). \quad (4)$$

A platform trap is a scenario in which all strategic agents (i.e., the sellers) join the platform despite all being weakly worse off, and some strictly worse off, than if the platform did not exist. Sellers end up worse off because each seller that joins lowers the payoff remaining sellers obtain by not joining, which lets the platform extract more from all of them. That is, the platform exploits a collective action problem among sellers, inducing them to collectively cannibalize their direct channel. In the end, under the sufficient condition (4), all sellers join despite being better off without the platform, i.e., if they collectively boycotted it. The sufficient condition (4) for a platform trap holds in our baseline marketplace setting if $v_P \geq v_D$ and $\pi_P \geq \pi_D$. For example, this could arise if the platform provides additional transactional benefits relative to the direct channel, so $q_P(p) \geq q_D(p)$ for all p .

Seller and buyer surplus

In an equilibrium where sellers join the platform, all sellers join, so seller k gets $\frac{1}{N}m(N)\pi_P - P^k$. For non-pivotal sellers, this equals $\frac{1}{N}(1 - m(N-1))\pi_D$. Recalling that (2) requires $m(N-1) > 0$, we conclude that non-pivotal sellers obtain strictly less than the $\frac{1}{N}\pi_D$ they would get without the platform. For pivotal seller k , the equilibrium payoff is $\frac{1}{N}(1 - m(k-1))\pi_D$, which is weakly lower than $\frac{1}{N}\pi_D$, given $m(k-1) \geq 0$. The last seller to sign, pivotal or not, is strictly worse off because $m(N-1) > 0$. The first seller to sign is indifferent if it is pivotal, given $m(0) = 0$. Therefore, in general, all sellers are weakly worse off, and at least one seller is strictly worse off, due to the

platform's existence.¹⁰

To understand what the platform's existence implies for buyers' expected utility, note that if the platform doesn't exist (or sellers collectively boycott it), buyers each get $v_D - s_D$ from the direct channel, which under our assumptions is always non-negative. Similarly, when all sellers join the platform, buyers each get $v_P - s_P$ from the platform, which is always their preferred option given $v_P \geq \bar{s}_P - \min\{\underline{s}_D, 0\}$. Then under a platform trap, buyers for which $s > v_D - v_P$ are strictly better off with the platform, while those for which $s < v_D - v_P$ are strictly worse off due to the platform's existence (where recall $s = s_D - s_P$). The possibility that some buyers are worse off reflects buyers who prefer to buy from sellers directly but, after the platform trap, cannot access sellers directly. If $BS(n)$ represents aggregate buyer surplus when n sellers join the platform, we have $BS(N) = v_P - \mathbb{E}[s_P]$ and $BS(0) = v_D - \mathbb{E}[s_D]$. Thus, aggregate buyer surplus is lower under a platform trap if and only if $v_P - v_D + \mathbb{E}[s] < 0$.

The question remains whether all buyers can be worse off in equilibrium. This is certainly possible. For simplicity, suppose $v_P = v_D = v$ and $\pi_P = \pi_D = \pi$, so the two channels are equivalent for matched buyer-seller transactions. Assume $\bar{s}_D = \underline{s}_D = 0$ and $\bar{s}_P > \underline{s}_P = 0$, so buyers only face a cost of going to the platform (for instance, if it is costly for them to adopt the new channel). Then $s = -s_P$, which is distributed with full support on $[-\bar{s}_P, 0]$. In this case, (2) is equivalent to $N > 2$, which also implies (4) holds,¹¹ so the equilibrium in which the platform attracts all sellers still exists. Then all buyers are weakly worse off, and all buyers with $s_P > 0$ are strictly worse off, obtaining $v - s_P$ with the platform instead of v without it, while total surplus is lower.

Another way in which all buyers could be worse off in a platform trap is if the platform charges ad valorem fees to sellers, rather than fixed prices, and sellers pass these fees through into higher buyer prices. In Supplementary Appendix C, we show that when the platform charges ad valorem fees, all buyers and sellers can be worse off due to the platform's existence even when $s_P = 0$ and $s_D \geq 0$ (so buyers only incur costs of going to the direct channel) and $q_P(p) > q_D(p)$ (so the platform provides additional transactional benefits relative to the direct channel).

Total surplus

Having considered how a platform's existence affects buyers and sellers, it remains to

¹⁰If G is strictly increasing on the interval $[-v_P, v_D]$, then $b(n)$ is strictly increasing and $u(n)$ is strictly decreasing for all n , and so the platform trap would always involve all sellers being strictly worse off, with the exception of the first seller to join if it is pivotal.

¹¹Indeed, it is equivalent to $N > \sum_{k=1}^N G\left(\left(\frac{N-(k-1)}{N} - \frac{k-1}{N}\right)v\right)$, which holds because $G\left(\left(\frac{N-(k-1)}{N} - \frac{k-1}{N}\right)v\right) \leq 1$ for all $k = 1, \dots, N$, with strict inequality for $k = N$.

consider whether total surplus, including the platform's profit, can also be lower under a platform trap. Since the fixed prices paid by sellers to the platform are transfers, total surplus when n sellers join the platform is

$$W(n) = nb(n) + (N - n)u(n) + BS(n).$$

In any equilibrium, either no sellers join or all sellers join, and in the latter case there is a platform trap. Thus, to see whether total surplus is higher with the platform when it is viable, we only need to compare $W(N)$ with $W(0)$. Since $BS(N) - BS(0) = v_P - v_D + \mathbb{E}[s]$, total surplus is strictly lower with the platform if and only if

$$b(N) < u(0) - \frac{1}{N}(v_P - v_D + \mathbb{E}[s]). \quad (5)$$

For a one-sided platform (with just strategic agents), or for a marketplace where buyer surplus is ignored, the condition for the platform to decrease total surplus is simply $b(N) < u(0)$. The same condition applies for a marketplace where buyers view the two channels symmetrically, in the sense that $v_P = v_D$ and $\mathbb{E}[s] = 0$.

Comparing (5) with (4), the condition for lower total surplus differs sharply from the sufficient condition for a platform trap: the inequalities in $b(N)$ go in opposite directions. Total surplus can nevertheless be lower under a platform trap. A sufficient condition is:

$$u(0) - \frac{1}{N}(v_P - v_D + \mathbb{E}[s]) > b(N) > \frac{1}{N} \sum_{k=1}^N u(k-1),$$

which is possible if $u(n)$ decreases fast enough in n . However, total surplus can also be higher despite a platform trap if $b(N)$ is high enough. Moreover, because $v_P - v_D + \mathbb{E}[s]$ can be positive or negative, a platform that would be efficient if all sellers joined need not generate a platform trap, because it may not be able to profitably attract sellers. We summarize these findings in the following proposition.

Proposition 2. *The existence of a platform that attracts all sellers decreases total surplus if and only if condition (5) holds. If buyers view the two channels symmetrically (or their surplus is ignored), the condition is simply $b(N) < u(0)$, and in this case, if the existence of the platform increases total surplus, a platform trap must arise. More generally, there is no necessary relationship between whether a platform that attracts all sellers increases total surplus and whether a platform trap arises, in the sense that all four combinations are possible.*

Proposition 2 illustrates that the problem of a platform trap is quite distinct from traditional concerns about a monopoly platform generating inefficient outcomes, with neither concern necessarily implying the other.

The conclusions of Propositions 1 and 2 are driven by the negative externalities that platform participation has on the direct channel (or, more generally, the alternative option). If the direct channel instead has a fixed value to sellers or does not exist, then the existence of a platform will always make sellers weakly better off: sellers can always choose not to participate and obtain this fixed value. This is the case in the classic analysis of two-sided platforms with positive cross-side network effects that set fees to attract users on each side (e.g., Rochet and Tirole, 2003 and Armstrong, 2006). In those models, the alternative to a monopoly platform is not a direct channel where the two sides can interact, perhaps less efficiently, but no platform at all (i.e., a fixed outside option). So there is no sense in which the platform is attracting buyers and sellers away from something which itself is affected by the existence of the platform. As a result, buyers and sellers can never be worse off due to the platform's existence, and total surplus is always higher. Applying this to our baseline model, if we remove the direct channel option so that the only alternative to joining the platform is the fixed outside option giving a zero payoff, the platform would charge each seller k joining in sequence the same price $b(N)$.¹² There would no longer be any possibility of pivotal sellers with the corresponding increasing prices P^k that we characterized in (3), or any kind of platform trap.

Seller competition

So far we have focused on a marketplace setting where sellers do not compete. In Supplementary Appendix D, we allow sellers to compete. Seller competition counteracts the positive effect of having more sellers on a channel (i.e., attracting more buyers). However, if the competition effect is not too strong, the buyer-expansion effect dominates, so platform profit $b(n)$ can still increase and direct-channel profit $u(n)$ can still decrease in n , and all the reduced-form payoff conditions needed for Proposition 1 still hold. We illustrate this with an example in Supplementary Appendix D. On the other hand, if the competition effect is sufficiently strong on the platform channel, and stronger than on the direct channel, $b(n)$ may eventually decrease even while $u(n)$ is everywhere decreasing in n . Supplementary Appendix D provides such an example, in which $b(n)$ is single-peaked despite all the other reduced-form payoff conditions needed for the platform-trap result continuing to hold. This motivates the generalization in

¹²To make this comparison meaningful, we must assume $b(\cdot)$ is weakly increasing, so that $\Delta(\cdot)$ can be weakly increasing despite $u(\cdot)$ being constant and equal to zero.

Section 5, which does not require $b(n)$ to be increasing and allows it, in particular, to be single-peaked.

5 Generalizations

We now turn to a more general reduced-form game with a single platform and N strategic agents, under the same timing and pricing assumptions as in Section 3. The N agents play the same role as sellers in the marketplace setting, while buyers' decisions are subsumed into the agents' payoffs $b(\cdot)$ and $u(\cdot)$. This generalization also captures one-sided settings such as social networks. We state the assumptions for this general setup with N agents.

- (A1) *Fully flexible prices.* The platform can set a different price P^k for each agent $k \in \{1, \dots, N\}$.
- (A2) *Take-it-or-leave-it offers.* The platform makes take-it-or-leave-it offers.
- (A3) *Full information and sequential moves.* Agents make sequential participation decisions and observe all prior decisions.
- (A4) *Ex-ante homogeneous agents.* All agents face the same payoff functions $b(\cdot)$ and $u(\cdot)$.
- (A5) *Monotonicity of payoff functions.* The payoff functions $b(\cdot)$ and $u(\cdot)$ satisfy the following conditions: (a) $u(\cdot)$ is non-negative, weakly decreasing in its argument, with $u(N-1) < u(0)$; and (b) $\Delta(n) = b(n) - u(n-1)$ is weakly increasing in n , with $\Delta(N) > 0$.

Each of Assumptions A1-A4 will be relaxed in later sections. Unlike the baseline model of Section 3, we do not assume $b(\cdot)$ is increasing. In a social network, benefits may initially increase as more users join, but later decrease due to congestion. Supplementary Appendix D shows that $b(\cdot)$ may be single-peaked when sellers compete and the competition effect is stronger on the platform than in the direct channel.

We show that the platform trap result in Proposition 1 applies to this more general setting.

Proposition 3. *Assume there are N agents, satisfying Assumptions A1-A5. Then Proposition 1 applies to these agents (replacing “sellers” with “agents”), and ignoring the implications for buyers, if any.*

For any given primitives (the number of agents N , and their payoff functions $b(\cdot)$ and $u(\cdot)$) one of two outcomes arises. Either the platform attracts all agents, yielding a unique equilibrium outcome in which all join and prices satisfy the properties in Proposition 1; or it attracts none, in which case there is a continuum of equilibria (in prices) with no participation. Furthermore, when the platform attracts all agents, there exists $k_0 \in \{0, \dots, N-1\}$ such that the first k_0 agents are pivotal and each agent $k \leq k_0$ is charged $P^k = b(N) - u(k-1)$, while the remaining $N - k_0$ agents are non-pivotal and each agent $k \geq k_0 + 1$ is charged $P^k = b(N) - u(N-1) = \Delta(N) > 0$.¹³ This is because at any stage, the subgame has only two possible participation outcomes: either all remaining agents join or none do (as shown in the proof of Proposition 1).

While Proposition 1 and its generalization in Proposition 3 provide a sufficient condition for a platform trap, stating a necessary and sufficient condition and fully characterizing the number k_0 of pivotal agents becomes increasingly complex as N grows. Let $\Pi(n|k)$ denote the platform's continuation profit after attracting $k \leq N$ agents and conditional on having to induce all $n \leq N$ remaining agents to participate. Suppose $\Pi(N|0) > 0 > \Pi(N-1|0)$, so the platform can profitably attract all N agents, but the first agent is pivotal: if she rejects, the platform cannot earn positive profit from the remaining $N-1$ agents. If the first agent accepts the offer, $\Pi(N|0) > 0$ implies $\Pi(N-1|1) > 0$. Is the second agent pivotal? If the second agent rejects, continuation profit is $\Pi(N-2|1)$: there are $N-2$ agents left, and only one of the first two agents has accepted. $\Pi(N-2|1)$ can be positive or negative depending on how strongly the first agent's decision weakens later agents' outside options. If the externality is weak, $\Pi(N-2|1)$ is close to $\Pi(N-2|0)$, so $\Pi(N-2|1) < 0$ as $\Pi(N-2|0) \leq \Pi(N-1|0) < 0$, and in this case, the second agent is pivotal. If the externality is strong enough that $\Pi(N-2|1) > 0$, the second agent is not pivotal. Thus, there may be one or multiple pivotal agents. The same reasoning applies to any agent k , conditional on all previous agents being pivotal.¹⁴ Agents that are not pivotal in the original game may become pivotal off the equilibrium path. Thus, each subgame may have a different number of pivotal agents, making a full characterization of the equilibrium (in particular, k_0) unwieldy.

To illustrate this and give a sense of how prices relate to model primitives, Supplementary Appendix E characterizes the equilibrium for $N = 2$ and $N = 3$. It also shows that even with linear $b(\cdot)$ and $u(\cdot)$, additional restrictions on their slopes are

¹³The prices charged to pivotal agents may or may not be negative. For example, if $b(N-1) - u(N-2) \leq 0$, then all agents are pivotal, but it is still possible to have $b(N) - u(0) \geq 0$, so that all agents face non-negative prices. This can happen if b increases a lot when the last agent joins.

¹⁴We are grateful to an anonymous referee who provided this explanation.

needed for a closed form solution for general N .

We can, however, characterize the necessary and sufficient condition for a platform trap, and the associated number of pivotal agents, for the two extreme cases discussed earlier.

Corollary 1. *If $\Delta(1) > 0$ or $b(N-1) > \frac{1}{N-1} \sum_{k=1}^{N-1} u(k-1)$, there is a unique equilibrium outcome in which the platform attracts all N agents and charges $P^k = \Delta(N) > 0$ for all $k \in \{1, \dots, N\}$. Here $k_0 = 0$ so there are no pivotal agents. All agents would be strictly better off without the platform.*

Corollary 2. *If $b(N-1) - u(N-2) \leq 0$, the platform attracts all N agents if and only if (4) holds. Under these conditions, there exists a unique equilibrium outcome in which the platform charges $P^k = b(N) - u(k-1)$ to agent $k \in \{1, \dots, N\}$. Here $k_0 = N-1$ so the first $N-1$ agents are pivotal. All agents would be weakly better off without the platform, with some strictly better off.*

Corollary 1 provides two alternative sufficient conditions for when no agent is pivotal. When $N > 2$ and $b(\cdot)$ is weakly increasing, the condition $b(N-1) > \frac{1}{N-1} \sum_{k=1}^{N-1} u(k-1)$, which is weaker than $\Delta(1) > 0$, suffices to rule out any agent being pivotal. On the other hand, Corollary 2 gives a sufficient condition for the case all but the last agent are pivotal.

Having characterized the platform trap with a fixed number of agents, a natural question is how prices, agents' net payoffs and platform profits change as the total number of agents N increases? Using the general formulation of payoffs in which we assume $b(\cdot)$ and $u(\cdot)$ are invariant to N , we find:

Corollary 3. *The net surplus left to agents weakly decreases in N , while platform profit increases in N . If the platform is profitable in the game with N agents, the number of pivotal agents is weakly lower in the game with $N' > N$ agents than in the game with N agents.*

Proposition 1 shows that prices for non-pivotal agents weakly increase with N , while prices for pivotal agents may increase or decrease, depending on whether $b(\cdot)$ increases or decreases in n . However, net payoffs for all agents always weakly decline as N increases because the total negative externality on the outside option is larger. Furthermore, it is intuitive that as N increases, each agent becomes less critical, making it easier for the platform to attract participation and reducing the likelihood of pivotal

agents. These results suggest that as a platform becomes more well-known (so more agents enter the market), the trap deepens. Over time, the platform is less likely to use negative prices, its prices rise, and agents' net payoffs decline.

6 Additional mechanisms

Our main platform trap result in Section 5 relied on Assumptions A1-A5. In this section, we show that relaxing Assumptions A1 and A3 sometimes limits and sometimes expands the scope for platform traps. It can also give rise to new platform trap mechanisms.

6.1 Constant pricing

One might think dynamic price adjustments are essential for the platform trap to arise. Perhaps surprisingly, we show this is not always the case by relaxing Assumption A1 in this section. Even when the platform must offer the same price to all agents (and agents still decide sequentially), a platform trap may still emerge. Since the platform must earn a profit, its constant price must be positive. Yet a platform trap can still occur when negative externalities dominate after enough agents join.

Proposition 4. *Suppose $b(\cdot)$ is single-peaked and the platform must charge the same price to all agents. There is a unique equilibrium outcome in which the platform attracts all N agents at the price $P = \max_{1 \leq n \leq N} \{b(n)\} - u(0)$, provided it is positive. In this case, agents are indifferent about the platform's existence if $b(N) = \max_{1 \leq n \leq N} \{b(n)\}$ and would be strictly better off without the platform if $b(N) < \max_{1 \leq n \leq N} \{b(n)\}$. If $\max_{1 \leq n \leq N} \{b(n)\} \leq u(0)$, the platform attracts no agents.*

The platform trap disappears when the platform must charge the same price to all agents and $b(\cdot)$ is weakly increasing. In this case, at the price stated in Proposition 4, each agent is effectively pivotal because the platform cannot adjust prices to later agents based on the decisions of earlier agents. Thus, the platform can profitably attract all agents if and only if $b(N) > u(0)$ and all agents are indifferent about the platform's existence (in the baseline setting only the first pivotal agent was indifferent).

However, when $b(\cdot)$ is decreasing, or first increasing and then decreasing, the platform trap reemerges despite the platform's inability to dynamically adjust prices. To see this, suppose $b(\cdot)$ peaks at $n^* < N$. Then the first $N - n^* \geq 1$ agents are not pivotal: using the argument for increasing $b(\cdot)$, the platform can attract the last n^*

agents even if none of the first $N - n^* \geq 1$ agents has joined. Thus, the first agents are easier to attract, while later agents keep joining as their outside option worsens, even though each additional participant reduces others' benefits. After the first n^* agents join, the trap is driven by both $b(\cdot)$ and $u(\cdot)$ declining, with $u(\cdot)$ falling faster.

This new mechanism implies some agents may be worse off under constant pricing than under price discrimination. In the baseline, if agent k was pivotal, it paid $b(N) - u(k - 1)$. With a constant price, agent k pays $b(n^*) - u(0)$, where $n^* = \arg \max_n \{b(n)\} < N$. If $n^* < N$ and

$$b(n^*) - b(N) > u(0) - u(k - 1),$$

then agent k pays more and receives a lower net payoff under constant pricing.¹⁵ This requires that $b(n)$ declines sufficiently from its peak to $b(N)$. Only early agents that were pivotal in the baseline may be worse off when the platform cannot price discriminate. By contrast, non-pivotal agents, including the last, are always weakly better off. Indeed

$$b(n^*) - u(0) \leq \Delta(n^*) \leq \Delta(N) = b(N) - u(N - 1).$$

Turning to welfare, Proposition 2 shows that a platform attracting all agents decreases total surplus if and only if $b(N) < u(0)$. Depending on the shape of $b(n)$, even if the platform can profitably operate when it must set a constant price, it may make agents weakly worse off and raise total surplus, or strictly worse off and lower total surplus. The latter case arises when $b(n)$ exceeds $u(0)$ for some n , but falls below $u(0)$ at N .

Proposition 4 implies that when $b(\cdot)$ is weakly increasing, the platform trap relies on the ability to adjust prices for future agents if an agent deviates and rejects its offer. However, it may not require much price flexibility to restore the trap. Supplementary Appendix F shows that if the platform can change prices once, the condition $b(N - 1) > u(0)$ ensures a unique equilibrium outcome in which all agents join at the price $\Delta(N) > 0$ and would be strictly better off without the platform.

6.2 The role of equilibrium selection

In the baseline model, agents decide whether to join the platform sequentially with full information, which ensures that equilibrium participation and payoffs are unique. If agents observe only their own offers, then multiple participation equilibria can arise

¹⁵In particular, this always holds for the first agent ($k = 1$) if they were pivotal in the baseline, since $b(n^*) - u(0) > b(N) - u(0)$.

for given platform prices, and equilibrium selection matters. This holds whether agents still decide sequentially and the platform can set different prices (only Assumption A3 is relaxed), or the platform has to set a constant price and agents decide simultaneously (both Assumptions A1 and A3 are relaxed).

The following proposition characterizes the best and worst equilibria for the platform when it sets a single price in stage 1 and agents decide simultaneously in stage 2. Supplementary Appendix G shows the same results hold when the platform can set different prices and agents only observe their own offers (regardless of whether they decide simultaneously or sequentially), provided we restrict attention to strong perfect Bayesian equilibria with passive beliefs in the sequential case.

Proposition 5. *Suppose the platform charges a single price in stage 1 and all agents decide simultaneously in stage 2. There is a continuum of equilibria.*

1. *In the best equilibrium for the platform, each agent pays $P = \Delta(N)$, all N agents join, and platform profits are $N\Delta(N)$. All agents would be strictly better off without the platform.*
2. *If $\Delta(1) > 0$, in the worst equilibrium for the platform, each agent pays $P = \Delta(1)$, all N agents join, and platform profits are $N\Delta(1)$. All agents would be strictly worse off without the platform if $b(N) > b(1)$, indifferent if $b(N) = b(1)$, and strictly better off if $b(N) < b(1)$. If $\Delta(1) \leq 0$, in the worst equilibrium for the platform, no agents join and the platform profits are zero.*

Depending on equilibrium selection, the platform trap may arise more broadly than in the baseline (case 1 in the proposition) or be more limited (case 2 when $b(\cdot)$ is weakly increasing). More generally, the comparison depends on equilibrium selection, the number of pivotal agents in the baseline, and the shape of $b(\cdot)$.¹⁶ The key reason is that agents are no longer pivotal: one agent's rejection cannot influence others. As a result, in the best equilibrium, previously pivotal agents are strictly worse off than in the baseline, while the platform does better and the trap applies for a wider set of payoff functions.

More surprisingly, the full platform trap can arise even in the worst equilibrium for the platform. This occurs when the aggregate externalities on the platform are negative, i.e., $b(N) < b(1)$. This shows that a platform trap can arise even without

¹⁶For example, suppose $b(\cdot)$ is decreasing and in the baseline all agents are pivotal. Then total payoffs to agents in the baseline are $\sum_{k=1}^N u(k-1)$, whereas here they are $N(u(0) - (b(1) - b(N)))$ in the worst equilibrium. Either one could be higher.

any dynamics or price discrimination. In this worst equilibrium, agents may still benefit from the platform, as they face a relatively low price ($P = \Delta(1)$) when on-platform network effects are positive. However, if $b(N) < b(1)$, all agents would be better off without the platform. The trap arises because agents join the platform even though full participation reduces their benefit. Each agent pays $\Delta(1) = b(1) - u(0)$ to join, knowing they will receive $b(N)$ rather than $b(1)$, because rejecting yields $u(N - 1)$, which is worse.¹⁷

This contrasts with Proposition 2 in Segal and Whinston (2000). They find that when an incumbent makes simultaneous offers to buyers and cannot discriminate, exclusion cannot occur in any perfectly coalition-proof Nash equilibrium. In our setting, under the single price $P = \Delta(1) > 0$, joining is the only Nash equilibrium: agents prefer to join even if no others do, and even more so if others join, so the perfectly coalition-proof refinement has no bite.

7 Extensions

In this section, we relax in turn three key assumptions — take-it-or-leave-it offers, homogeneous agents, and monopoly platform — and show that platform traps persist in these richer settings, while also examining how platforms actively engineer them. For brevity, the proofs for this section are contained in Supplementary Appendix H.

7.1 Negotiated prices

The baseline assumed the platform makes take-it-or-leave-it offers to agents (Assumption A2). If the platform must negotiate, its ability to extract surplus, and thus the scope of the platform trap, is limited. To explore this and the implications for price dynamics, we now assume the platform engages in sequential Nash bargaining with each agent. We focus on the case with no pivotal agents, so $\Delta(1) > 0$, which ensures the platform can profitably attract all agents. This scenario yields the strongest form of the trap and serves as a benchmark: even here, negotiation can imply some agents are not made worse off by the platform’s existence. The assumption also ensures closed form solutions despite the recursive structure that arises because each participation decision affects all subsequent negotiations. Finally, assuming $\Delta(1) > 0$ shows that prices can still increase as more agents join even without pivotal agents, because each additional agent weakens the bargaining position of those who follow.

¹⁷Indeed, $u(N - 1) \leq b(N) - \Delta(1)$ since $\Delta(1) \leq \Delta(N)$.

Proposition 6. *Suppose $\Delta(1) > 0$ and the platform bargains sequentially with N agents, where each agent has bargaining power $0 \leq \alpha < 1$ and the platform has $1 - \alpha$. There is a unique equilibrium in which all agents participate and the price charged to agent k is*

$$P^k = (1 - \alpha) \left(\sum_{j=0}^{N-k} C_j^{N-k} (1 - \alpha)^j \alpha^{N-k-j} \Delta(k + j) \right)$$

for all $k \in \{1, \dots, N\}$, where C_j^{N-k} is the binomial coefficient. With linear externalities such that $\Delta(n) = \beta_0 + \beta_1 n$ and $\beta_1 > 0$, prices P^k increase in k when bargaining power is intermediate (i.e., $0 < \alpha < 1$).

Since $\sum_{j=0}^{N-k} C_j^{N-k} (1 - \alpha)^j \alpha^{N-k-j} = 1$ for any k , the optimal price to agent k is $(1 - \alpha)$ times a weighted average of $\Delta(n)$ for n ranging from k to N . If the externalities $u(\cdot)$ and $b(\cdot)$ are linear and produce $\Delta(n) = \beta_0 + \beta_1 n$ with $\beta_1 > 0$, prices are

$$P^k = (1 - \alpha) (\beta_0 + \beta_1 (N - \alpha(N - k))). \quad (6)$$

Given that $0 < \alpha < 1$ and $\beta_1 > 0$, prices are increasing in k . If $b(N) < u(0)$ or α is sufficiently low, all agents would be strictly better off without the platform. If $b(N) > u(0)$ and α is sufficiently high, all agents are better off with the platform. Finally, if $b(N) > u(0)$, there is an intermediate range of α , such that agents receiving early offers would be strictly worse off without the platform, while those receiving later offers would be strictly better off without it.

7.2 One superstar agent

So far all agents have been homogeneous (Assumption A4). We now relax this by adding one “superstar” agent alongside N regular (symmetric) agents. This raises the question: which type of agent should the platform attract first to engineer a platform trap?

We assume the superstar is equivalent to a coalition of $S > 1$ agents, with payoffs $Sb(\cdot)$ and $Su(\cdot)$, and the same impact on all agents’ payoffs as S regular agents.¹⁸ Thus, when the superstar agent and $n \leq N$ regular agents join the platform, an individual agent’s payoff on the platform is $b(S + n)$ and its outside option payoff is $u(S + n)$. If the superstar does not join, payoffs are the usual $b(n)$ and $u(n)$. We adjust the baseline

¹⁸Supplementary Appendix H.2 analyzes a slightly different version in which the superstar has the size and the payoffs of an individual agent, but an outsized impact on agents’ payoffs. The results are similar.

assumption that every agent receives positive surplus from the platform when all others join by imposing $b(N + S) > u(N)$, which also implies $b(N + S) > u(N + S - 1)$ since $S > 1$. To streamline the exposition, we also assume $u(\cdot)$ is strictly decreasing.

Focusing on the most relevant parameter ranges, we obtain the following proposition.

Proposition 7. *Suppose there are N regular agents and one superstar agent, of size S and equivalent in impact to $S > 1$ regular agents.*

1. *If $\min\{b(N), b(N + S - 1)\} > u(0)$, the platform attracts all agents and the order of offers is irrelevant.*
2. *If $b(N + S - 1) > u(0) > u(N - 1) > b(N)$, the platform attracts all agents and optimally offers to the superstar last.*
3. *If $b(S + N - 1) < u(S + N - 2)$ and $u(\cdot)$ is convex, the platform offers to the superstar last and attracts all agents if and only if $b(N + S) > \frac{Su(N) + \sum_{k=0}^{N-1} u(k)}{N + S}$.*
4. *If $b(S + N - 1) < u(S + N - 2)$ and $u(\cdot)$ is concave, the platform offers to the superstar first and attracts all agents if and only if $b(N + S) > \frac{Su(0) + \sum_{k=0}^{N-1} u(S + k)}{N + S}$.*

In cases (1) and (2), all agents would be strictly better off without the platform. In cases (3) and (4), all agents except the first would be strictly better off without the platform, and the first is indifferent.

When no agents are pivotal (first case in Proposition 7), each knows the platform can still attract all others regardless of their decision. Thus, the order of offers does not matter and the platform attains maximum profit by leaving each agent with their lowest outside option.

The platform can obtain the same profit even if only the superstar is pivotal (second case in Proposition 7) by approaching it last, eliminating its pivotal role. By first contracting with the regular (non-pivotal) agents, who anticipate the superstar will join, the platform extracts their full surplus. Their participation undermines the superstar's outside option, allowing the platform to charge it the highest feasible price at the end.

Finally, when all agents are pivotal (third and fourth cases in Proposition 7), the platform can no longer achieve the same maximum profit. The optimal timing for approaching the superstar now depends on the curvature of $u(\cdot)$: approach it last if $u(\cdot)$ is convex, and first if it is concave. Approaching the superstar early requires offering it a lower price (its outside option is higher), but once secured, its participation allows the platform to charge higher prices to more subsequent agents (their outside

options fall more). With convex $u(\cdot)$ the first effect dominates; with concave $u(\cdot)$ the second does. If $u(\cdot)$ is linear, the timing is irrelevant.

The general logic is that the platform should first attract non-pivotal agents to minimize the price discounts needed to attract pivotal agents. Afterward, it should attract those pivotal agents whose participation creates the greatest negative externality on others' outside options. An implication is that agents with a constant outside option should be attracted first. In reality, this might mean first attracting agents with very weak (or no) outside options, since their presence on the platform helps attract others, but not vice versa.

7.3 Platform competition

To what extent does platform competition weaken the platform trap? To explore this, suppose there are two platforms. The joining process remains sequential over N stages, with one agent selected at random in each stage. At each stage, both platforms simultaneously make offers to the selected agent.

Agents receive $b_1(n_1, n_2)$ or $b_2(n_1, n_2)$ from joining platform 1 or 2, respectively, when n_1 agents are on platform 1 and n_2 are on platform 2. We assume $b_1(n_1, n_2)$ is weakly increasing in n_1 and weakly decreasing in n_2 , and vice versa for $b_2(n_1, n_2)$, so platforms create network effects. The outside option is $u(n_1, n_2)$, strictly decreasing in both arguments, reflecting that more participation on either platform reduces the value of staying out.

We also assume $b_1(N, 0) > u(N - 1, 0)$ and $b_2(0, N) > u(0, N - 1)$, so both platforms offer positive surplus when attracting all agents, and $b_1(m, n) > b_2(n, m)$ for any (m, n) with $1 \leq m + n \leq N$, so platform 1 offers higher benefits at any configuration. Finally, we adopt the following tie-breaking rule: whenever agents are indifferent between the two platforms, they join the one offering higher gross benefits (and platform 1 if equal), and if indifferent between a platform and the outside option, they join a platform.

We characterize the outcome for $N = 2$ under two additional assumptions:

$$b_1(2, 0) > u(0, 0) \tag{7}$$

$$u(0, 1) \geq u(1, 0). \tag{8}$$

Assumption (7) ensures that if only platform 1 exists, it is more efficient for both agents to join platform 1 than for both to stay out. Assumption (8) states that the outside option is (weakly) worse when one agent joins platform 1 than when one agent

joins platform 2, reflecting that platform 1 is better and so an agent joining it has a greater negative effect on the outside option. Together with the assumptions on b_1 and b_2 above, these ensure that platform 1 attracts both agents.

The following proposition focuses on the most relevant cases.

Proposition 8. *Suppose $N = 2$ and conditions (7)-(8) hold. Platform 1 always attracts both agents and agent 1 is weakly better off than agent 2.*

1. *If $b_2(1, 1) > u(0, 0)$, both agents are better off with the platforms.*
2. *If $b_2(1, 1) < u(0, 0)$ and $b_1(1, 1) > \max\{u(0, 0), b_2(0, 2)\}$, both agents are worse off with the platforms.*
3. *If $b_2(1, 1) < u(0, 0)$ and $b_2(0, 2) > \max\{u(0, 0), b_1(1, 1)\}$, agent 1 is better off and agent 2 is worse off with the platforms.*

In case 1, platform 2 provides more value when agents split between platforms than the outside option does when no agents join. Thus, platform 2 exerts enough competitive pressure on platform 1 that both agents are better off with the platforms. In case 2, platform 1 offers more value when agents split than both the outside option and platform 2 when it attracts both agents. Here, neither platform 2 nor the outside option provides enough pressure, so both agents are worse off with the platforms. Finally, if platform 2 exerts enough competitive pressure on platform 1 for the first agent, but not after the first agent joins platform 1 or stays out, the first agent is better off with the platforms, while the second agent is worse off.

Similar results can be obtained for $N > 2$, but a full characterization becomes very messy. The intuition is the same: for at least one agent to be better off with competing platforms, the less preferred platform must be sufficiently competitive relative to the initial outside option. The results also hold when the agents' payoffs on each platform are independent of the number of agents who join the other platform, i.e. $b_1(n_1, n_2) = b_1(n_1)$ and $b_2(n_1, n_2) = b_2(n_2)$ for all (n_1, n_2) .

At the symmetric boundary, relaxing the preceding strict ranking, suppose $b_2(n_1, n_2) = b_1(n_2, n_1)$ for every combination of $n_1 \in \{0, 1, 2\}$ and $n_2 \in \{0, 1, 2\}$, agent 1 is always weakly better off with the platforms. This is not surprising: competition for the first agent is most intense, so that agent benefits. However, agent 2 may still be worse off, since case 3 can still arise, namely we can still have $b_2(0, 2) > u(0, 0) > b_1(1, 1) = b_2(1, 1)$. Whether agent 2 benefits depends on how effective platform 2 is as a competitor once platform 1 has attracted agent 1. If the value platform 2 can offer agent

2 in that case is low, platform 1 can extract enough surplus to leave agent 2 worse off than without the platforms.

So far, agents can join only one platform. If agents can multihome and benefit functions are independent, so $b_1(n_1, n_2) = b_1(n_1)$ and $b_2(n_1, n_2) = b_2(n_2)$, then payoffs are $b_1(n_1) + b_2(n_2)$ when multihoming and $u(n_1, n_2) = u(n_1) + u(n_2)$ when staying out. We then return to the baseline analysis. There is no real competition between platforms in this case. For example, this arises when agents are sellers on two marketplaces serving completely distinct buyer segments, each of which decides whether to go to the outside option or the particular marketplace of interest.

8 Real-world examples and policy implications

In this section, we connect our model’s predictions to real-world applications and discuss some implications for policy. The first prediction is that platform prices increase as more agents enter the market and join the platform. This can occur because early participants are pivotal and receive better terms (Proposition 1), or, even without pivotal agents, because each additional agent who joins weakens the bargaining position of those who follow (Proposition 6). These predictions are consistent with the widespread practice among real-world platforms of offering highly favorable terms early on (e.g., low fees, subsidies, exclusivity, preferential treatment), then raising fees and standardizing contracts as they scale.

For example, Amazon has not only raised fees in some categories, but also introduced new ones, that are difficult for newer sellers to avoid. Mitchell (2023), in an issue brief for the Institute for Local Self-Reliance, calculates that Amazon captured 45 percent of U.S. sellers’ revenue in the first half of 2023 (up from 35 percent in 2020 and 19 percent in 2014) through referral, advertising, fulfillment, and other fees. Another example is Uber. When entering new markets (cities), it subsidized drivers with incentives such as hourly or monthly earnings guarantees, which were later removed (Hall and Krueger, 2018). Similarly, Spotify initially offered record labels equity stakes, minimum revenue guarantees, and favorable licensing terms to secure key music catalogs. Later participants joined without such concessions.

A second key prediction concerns the order in which platforms attract agents. The logic of pivotal agents would seem to imply that platforms should target large, high-impact participants first and offer them preferential terms. That would also be the general prediction based on classic models of platforms with network effects: first attract the participants that generate the largest network effect on the platform. Shop-

ping malls illustrate this logic: developers typically first secure anchor tenants such as department stores or supermarkets with generous deals (below-market or even zero rents, long-term leases, and prime locations), then fill in with smaller retailers that pay higher rents and face co-tenancy clauses that tie their viability to the continued presence of anchor stores (Pashigian and Gould, 1998).

However, this ordering need not be optimal when participation also worsens agents' outside options (the direct channel). As shown in Propositions 7 and 14 (the latter is in the Supplementary Appendix), when only the superstar is pivotal or when all agents are pivotal and the outside option deteriorates rapidly at first and then more slowly, platforms may optimally reverse the order: attract "regular" (small) agents first to erode the outside option of "superstar" (large) ones, who are then compelled to join on less favorable terms.

The evolution of YouTube illustrates this logic. Early on (2005-2008), it focused on attracting individuals and incentivizing them to upload user-generated content. Only later (from around 2010) did it pivot to professional content and large media partners such as TV networks and movie studios. By contrast, Brightcove, one of YouTube's early competitors in Internet video, initially targeted major media companies (e.g., New York Times, Discovery Channel, MTV, CBS) with enterprise video infrastructure. Its aim was to build critical mass via professional content, before turning to user-generated content. This strategy backfired: large media partners imposed constraints that limited Brightcove's ability to scale a consumer platform, while YouTube faced no such restrictions. Ultimately, those same media companies were compelled to join YouTube, despite their initial reluctance.¹⁹

Third, from a policy perspective, the possible presence of a platform trap is not sufficient to warrant drastic interventions (e.g., banning platforms or imposing broad behavioral restrictions). As we showed, such traps can coexist with higher total surplus, especially in marketplace settings where buyer-side gains may compensate for seller-side surplus losses. However, we do think platform traps merit attention when evaluating the impact of platforms and related policies. Some platforms highlight that their suppliers enjoy higher revenues. For example, a report commissioned by DoorDash²⁰ cites a survey of individual restaurants, in which 85% of respondents said they would have lower overall revenue if they were to stop using DoorDash. But this does not imply that restaurants are collectively better off because of food delivery platforms. The key issue is that outside options are endogenous: they may be weak precisely because the

¹⁹See "Brightcove, Inc. (TN)," Harvard Business School teaching note 714-441, 2013.

²⁰See <https://doordash2024.publicfirst.co/>

platform is widely adopted. As a result, comparing a restaurant’s current payoff on the platform to its current outside option can be misleading, even accounting for diverted sales. A restaurant may be worse off leaving unilaterally because customers remain on the platform, whereas the outside option could be much stronger if many restaurants exited or if the platform were less prevalent.

Rather than banning platforms or imposing broad behavioral restrictions (e.g., limits on pricing), policymakers could mitigate platform traps by allowing affected agents to coordinate. When these agents are competing sellers, such coordination is usually treated as collusion under competition law. Our results suggest there may be value in carving out exceptions for sellers negotiating with dominant marketplaces or publishers negotiating with Google and Meta, provided coordination is limited to their role as platform customers. There are precedents. Australia’s Treasury Laws Amendment (News Media and Digital Platforms Mandatory Bargaining Code) Act 2021 requires designated platforms (e.g., Facebook and Google) to negotiate with news publishers or face arbitration, and explicitly allows publishers to bargain collectively — conduct that would otherwise risk being deemed collusive.

Another policy intervention would be to restrict platform practices that deliberately weaken agents’ outside options. For instance, some food delivery platforms (like DoorDash) withhold customer data (e.g., customer details, order histories, and detailed feedback) from participating restaurants. This makes it harder for restaurants to serve these customers well via direct channels, thereby pushing more consumers to use the platform. Another example is Apple’s iMessage, which shows texts coming from Android devices as green bubbles — Bursztyn et al. (2025b) find that this stigmatizes Android users and reduces the perceived quality of Android phones.

9 Concluding remarks

Commentators have noted that far from being a boon for participants, some platforms may end up hurting them. In online marketplaces, for example, this could involve saddling sellers with new fees to reach the same customers they previously served directly. Our model demonstrates how a platform trap can emerge despite agents being rational and forward-looking. Furthermore, the model explains increasing price dynamics: prices may rise because early agents are pivotal to attracting later agents, or because each joining agent endogenously weakens the bargaining position of those who follow. We also determine the optimal order for a platform to attract different (heterogeneous) agents in order to best engineer a platform trap.

Other forms of hold-up may also complement our theory in multi-period settings. For example, agents could be locked in by irreversible platform-specific investments, or, in the case of sellers joining a marketplace, by the buyers they bring to the platform becoming loyal to it (Karle et al., 2026). Unlike our mechanism, where an agent’s decision reduces the outside option for all agents, these hold-up mechanisms affect only the individual agent’s outside option. Hence, they do not directly produce a platform trap, but it would nevertheless be interesting to integrate these hold-up mechanisms into our framework.

It would also be interesting to explore multi-period versions of our model; we considered one such extension, as discussed in Section 3.1. With multiple agents making participation decisions each period and externalities across them, equilibrium selection plays a role in determining the outcome. In such a setting, equilibrium selection may be influenced by participants’ past choices, following the ideas in Halaburda and Yehezkel (2019) and Halaburda et al. (2020). In particular, equilibrium selection may be unfavorable for the platform at first, but if it can attract a critical number of agents to join, the platform may be able to enjoy favorable equilibrium selection after some point. This suggests a further mechanism by which a dynamic platform trap may arise.

10 Appendix

In this appendix we provide the proofs of propositions not proven in the text.

10.1 Proof of Propositions 1 and 3

Denote by $\Gamma(n, b(\cdot), u(\cdot))$ the generalized game from Section 5, with $N = n$ agents and payoff functions $b(\cdot)$ and $u(\cdot)$ satisfying Assumption A5. This formulation encompasses the microfounded version of the game from the baseline model, where agents are sellers on a marketplace. Thus, the proof that follows applies to both Proposition 1 and Proposition 3.

We start with two lemmas, which will be useful later.

Lemma 1: *If the platform can profitably attract all agents in the game $\Gamma(n, b(\cdot), u(\cdot))$, then it profitably attracts all agents in the game $\Gamma(n + 1, b(\cdot), u(\cdot))$ with optimal prices $P^k = \Delta(n + 1)$.*

Proof: The proof is by induction over games with an increasing number of agents. If the platform can profitably attract the sole agent in the game $\Gamma(1, b(\cdot), u(\cdot))$, then $b(1) > u(0)$, i.e. $\Delta(1) > 0$. Thus, in the game $\Gamma(2, b(\cdot), u(\cdot))$, if the first agent

doesn't join, the platform can still attract the second one. And $\Delta(2) \geq \Delta(1) > 0$, so if the first agent does join, the platform also profitably attracts the second one. Thus, it profitably attracts both agents in $\Gamma(2, b(\cdot), u(\cdot))$ by charging each $\Delta(2) > 0$.

Suppose the statement in the lemma is true for $N = n - 1 \geq 1$ and *any* payoff functions $b(\cdot)$ and $u(\cdot)$ such that $u(\cdot)$ is weakly decreasing and $\Delta(k) = b(k) - u(k - 1)$ is weakly increasing. We show it is also true for $N = n$. Namely, suppose the platform profitably attracts all agents in the game $\Gamma(n, b(\cdot), u(\cdot))$, so $\Delta(n) > 0$. We want to show the platform profitably attracts all agents in the game $\Gamma(n + 1, b(\cdot), u(\cdot))$ at prices $P^k = \Delta(n + 1)$ for $k = 1, \dots, n + 1$.

If the first of $n + 1$ agents doesn't join, when facing the second agent, the platform is in the same position as at the start of $\Gamma(n, b(\cdot), u(\cdot))$, so it still attracts the last n out of $n + 1$ agents. If the first agent joins and the second does not, then when facing agents $\{3, \dots, n + 1\}$, the platform is in the same position as when facing the second agent in $\Gamma(n, b(\cdot), u(\cdot))$ after the first agent has joined. By assumption, the platform attracts all agents in this situation, so the platform profitably attracts all agents $\{3, \dots, n + 1\}$ in $\Gamma(n + 1, b(\cdot), u(\cdot))$ after the first agent joins and the second does not. This is equivalent to the platform profitably attracting all agents in $\Gamma(n - 1, \tilde{b}(\cdot), \tilde{u}(\cdot))$, where $\tilde{b}(k) \equiv b(k + 1)$ and $\tilde{u}(k) \equiv u(k + 1)$ (note $\tilde{u}(\cdot)$ is weakly decreasing and $\tilde{\Delta}(k) = \tilde{b}(k) - \tilde{u}(k - 1) = b(k + 1) - u(k)$ is weakly increasing). The induction hypothesis ($N = n - 1$) then implies the platform profitably attracts all agents in $\Gamma(n, \tilde{b}(\cdot), \tilde{u}(\cdot))$. And this is equivalent to saying that in $\Gamma(n + 1, b(\cdot), u(\cdot))$, after the first agent joins, the platform profitably attracts all remaining n agents.

Thus, the platform profitably attracts the last n agents in $\Gamma(n + 1, b(\cdot), u(\cdot))$ regardless of whether the first agent joins or not, so it can attract the first agent with $P^1 = \Delta(n + 1) \geq \Delta(n) > 0$. Using a similar logic, the platform attracts all agents in $\Gamma(n + 1, b(\cdot), u(\cdot))$ with the same price $\Delta(n + 1)$, which maximizes profits. ■

Lemma 2: *In the game $\Gamma(n, b(\cdot), u(\cdot))$, either the platform optimally attracts all agents, or optimally attracts none of them.*

Proof: Suppose to the contrary, the platform finds it optimal to only attract $0 < n_0 < n$ agents in game $\Gamma(n, b(\cdot), u(\cdot))$. In this case, $\Delta(n_0) > 0$ because this is the highest price the platform can charge to participating agents. If the last agent that does not join is the last one overall, then the platform can attract it with a price $\Delta(n_0 + 1) \geq \Delta(n_0) > 0$, which is a contradiction. Suppose instead the last agent that does not join is the $(n - k_0)$ -th agent, so the last $k_0 \geq 1$ agents are all among the n_0 agents that join. Consider any one of these final k_0 agents. If that agent rejects, at most $n_0 - 1$ agents can ultimately join, so the price that can be charged to such an

agent can be no higher than

$$b(n_0) - u(n_0 - 1) = \Delta(n_0).$$

Consequently, the platform's continuation profit from the final k_0 participating agents is at most

$$k_0 \Delta(n_0).$$

The platform profitably attracts all agents in the game $\Gamma(k_0, \tilde{b}(\cdot), \tilde{u}(\cdot))$, where $\tilde{b}(k) = b(n_0 - k_0 + k)$ and $\tilde{u}(k) = u(n_0 - k_0 + k)$. Applying Lemma 1, the platform can also profitably attract all agents in the game $\Gamma(k_0 + 1, \tilde{b}(\cdot), \tilde{u}(\cdot))$ by charging each of them

$$\tilde{\Delta}(k_0 + 1) = \Delta(n_0 + 1) > 0.$$

This alternative continuation therefore earns

$$(k_0 + 1) \Delta(n_0 + 1) \geq (k_0 + 1) \Delta(n_0) > k_0 \Delta(n_0).$$

Revenues collected before this continuation are unchanged. The platform can therefore strictly increase its total profit, contradicting the supposed optimality of attracting only n_0 agents. ■

The proof of Proposition 1 also proceeds by induction over games with an increasing number of agents. Start with $\Gamma(N = 2, b(\cdot), u(\cdot))$. If $\Delta(1) > 0$, then neither agent is pivotal, so the platform maximizes profits by attracting both agents with prices $P^1 = P^2 = \Delta(2) \geq \Delta(1) > 0$. If instead $\Delta(1) \leq 0$, then the first agent is pivotal, so the profit-maximizing prices that attract both agents are $P^1 = b(2) - u(0)$ and $P^2 = b(2) - u(1)$. From Lemma 2, if the platform finds it optimal to attract any agents, it must attract both, so these prices are optimal if $2b(2) - u(1) - u(0) > 0$.

Suppose the following induction hypothesis holds for $N = n \geq 2$ (we have just shown it holds for $N = 2$) and *any* payoff functions $b(\cdot)$ and $u(\cdot)$ such that $u(\cdot)$ is weakly decreasing and $\Delta(k) = b(k) - u(k - 1)$ is weakly increasing in k :

- If the platform can profitably attract all agents in the game $\Gamma(n - 1, b(\cdot), u(\cdot))$, then it profitably attracts all agents in the game $\Gamma(n, b(\cdot), u(\cdot))$ with optimal prices $P^k = \Delta(n)$.
- If the platform cannot profitably attract agents in the game $\Gamma(n - 1, b(\cdot), u(\cdot))$, there exists $k_0 \in \{1, \dots, n - 1\}$ such that the platform's profit-maximizing prices

that induce agents to join in the game $\Gamma(n, b(\cdot), u(\cdot))$ are

$$P^k = \begin{cases} b(n) - u(k-1) & \text{if } 1 \leq k \leq k_0 \\ b(n) - u(n-1) & \text{if } k_0 + 1 \leq k \leq n \end{cases}.$$

If $\sum_{k=1}^n P^k > 0$, then these are the platform's optimal prices in $\Gamma(n, b(\cdot), u(\cdot))$ and all agents join. If instead $\sum_{k=1}^n P^k \leq 0$, then it is optimal to not attract any agents.

We now show this also holds for $N = n+1$. If the platform can profitably attract all agents in the game $\Gamma(n, b(\cdot), u(\cdot))$, then Lemma 1 directly implies it profitably attracts all agents in the game $\Gamma(n+1, b(\cdot), u(\cdot))$ with optimal prices $P^k = \Delta(n+1)$. Suppose instead the platform cannot profitably attract any agents in the game $\Gamma(n, b(\cdot), u(\cdot))$. In this case, the highest price at which the platform can attract the first agent in $\Gamma(n+1, b(\cdot), u(\cdot))$ is $P^1 = b(n+1) - u(0)$. Once the first agent participates, then when facing the remaining n agents, the platform is in a position equivalent to that at the start of game $\Gamma(n, \tilde{b}(\cdot), \tilde{u}(\cdot))$, where $\tilde{b}(k) = b(k+1)$ and $\tilde{u}(k) = u(k+1)$. We can then apply the induction hypothesis to conclude:

- If the platform can attract all agents in $\Gamma(n-1, \tilde{b}(\cdot), \tilde{u}(\cdot))$, then it attracts the last n agents in $\Gamma(n+1, b(\cdot), u(\cdot))$ after the first agent has joined with optimal prices

$$P^k = \tilde{b}(n) - \tilde{u}(n-1) = \Delta(n+1) > 0$$

for all $k \in \{2, \dots, n+1\}$. So we have

$$P^k = \begin{cases} b(n+1) - u(0) & \text{if } k = 1 \\ b(n+1) - u(n) & \text{if } 2 \leq k \leq n+1 \end{cases}.$$

- If the platform cannot profitably attract agents in $\Gamma(n-1, \tilde{b}(\cdot), \tilde{u}(\cdot))$, there exists $k_0 \in \{2, \dots, n\}$ such that in $\Gamma(n+1, b(\cdot), u(\cdot))$ after the first agent has joined, the platform's profit-maximizing prices for agents $k \in \{2, \dots, n+1\}$ that induce them to join are

$$P^k = \begin{cases} \tilde{b}(n) - \tilde{u}(k-2) = b(n+1) - u(k-1) & \text{if } 2 \leq k \leq k_0 \\ \tilde{b}(n) - \tilde{u}(n-1) = b(n+1) - u(n) & \text{if } k_0 + 1 \leq k \leq n+1 \end{cases}.$$

If $\sum_{k=2}^{n+1} P^k > 0$, these prices are optimal for agents $k \in \{2, \dots, n+1\}$. If instead $\sum_{k=2}^{n+1} P^k \leq 0$, the platform cannot profitably attract agents $k \in \{2, \dots, n+1\}$

even after the first agent joins, so the platform sets sufficiently high nonnegative prices, or makes no offers, and attracts no agents. Thus, the platform's profit-maximizing prices in $\Gamma(n+1, b(\cdot), u(\cdot))$ that induce all agents to join (recall Lemma 2) are

$$P^k = \begin{cases} b(n+1) - u(k-1) & \text{if } 1 \leq k \leq k_0 \\ b(n+1) - u(n) & \text{if } k_0 + 1 \leq k \leq n+1 \end{cases}.$$

If $\sum_{k=1}^{n+1} P^k > 0$, these are the platform's optimal prices and all agents join (note $\sum_{k=1}^{n+1} P^k > 0$ implies $\sum_{k=2}^{n+1} P^k > 0$ because P^k is weakly increasing in k). If instead $\sum_{k=1}^{n+1} P^k \leq 0$, the platform finds it optimal to attract no agents in $\Gamma(n+1, b(\cdot), u(\cdot))$.

Thus, the induction hypothesis holds for $N = n+1$. By induction, it holds for all $N \geq 2$. Moreover, if (4) holds, then the platform optimally attracts all agents in $\Gamma(N, b(\cdot), u(\cdot))$ even in the worst case when all of them are pivotal, so it can always profitably attract N agents. Finally, agent k 's payoff in $\Gamma(N, b(\cdot), u(\cdot))$ when all join is $b(N) - P^k$, which is either $u(N-1)$ (if agent k is not pivotal) or $u(k-1)$ (if agent k is pivotal). These are both weakly lower than $u(0)$ (strictly for $k = N$), which is the payoff each agent would obtain without the platform.

10.2 Proof of Corollary 1

The case $\Delta(1) > 0$ is already explained in the text before Proposition 1. If $b(N-1) > \frac{1}{N-1} \sum_{k=1}^{N-1} u(k-1)$, Proposition 1 implies that the platform profitably attracts the last $N-1$ agents even if the first agent does not join. The platform can then attract the first agent at price $P^1 = b(N) - u(N-1)$. By a similar logic, it attracts all remaining agents $k = 2, \dots, N$ at prices $P^k = b(N) - u(N-1)$.

10.3 Proof of Corollary 2

If $b(N-1) - u(N-2) \leq 0$, then $b(k) - u(k-1) \leq 0$ for all $k \in \{1, \dots, N-1\}$, i.e., all agents are pivotal. The logic used in the proof of Proposition 1 implies the platform optimally sets $P^k = b(N) - u(k-1)$ to the k -th agent for $k = 1, \dots, N$ provided (4) holds. Otherwise, the platform finds it optimal to attract no agents.

10.4 Proof of Corollary 3

From Proposition 1, agent k 's net surplus in the game $\Gamma(N, b(\cdot), u(\cdot))$ is either $u(k-1)$ or $u(N-1)$, so is weakly decreasing in N . From Lemma 2 in the Proof of Proposition 1, there are two possibilities:

- If the platform finds it optimal to attract all agents in $\Gamma(N, b(\cdot), u(\cdot))$, its profits are at most $N\Delta(N)$ and repeated application of Lemma 1 implies that platform profits in $\Gamma(N', b(\cdot), u(\cdot))$ are $N'\Delta(N') > N\Delta(N)$.
- If the platform finds it optimal to attract no agents in $\Gamma(N, b(\cdot), u(\cdot))$, then its profits in $\Gamma(N', b(\cdot), u(\cdot))$ are weakly higher (strictly when it can profitably attract all agents in $\Gamma(N', b(\cdot), u(\cdot))$).

If the platform is profitable in $\Gamma(N, b(\cdot), u(\cdot))$, then Lemma 1 implies that for any $N' > N$, the platform attracts all agents with prices equal to $\Delta(N') > 0$ in $\Gamma(N', b(\cdot), u(\cdot))$, so no agent is pivotal. Thus, the number of pivotal agents in $\Gamma(N', b(\cdot), u(\cdot))$ is weakly lower than in $\Gamma(N, b(\cdot), u(\cdot))$ (strictly whenever there are any pivotal agents in $\Gamma(N, b(\cdot), u(\cdot))$).

10.5 Proof of Proposition 4

Here, we denote by $\Gamma(n, b(\cdot), u(\cdot), P)$ the platform adoption game that unfolds when there are $N = n$ total agents, payoff functions are $b(\cdot), u(\cdot)$, and the platform has set price P .

First, we prove that the only possible equilibria in $\Gamma(n, b(\cdot), u(\cdot), P)$ are all agents join or no agents join. Suppose instead $1 \leq m < N$ agents join in equilibrium, and the last agent does not join. This implies $P > \Delta(m+1)$. The last agent to join obtains $b(m) - P$ in equilibrium. If that agent deviates to not joining, then all subsequent agents will continue to not join (they are less likely to do so because fewer prior agents will have joined; this can be easily proven by induction), so the agent's deviation payoff is $u(m-1)$. Since $P > \Delta(m+1) \geq \Delta(m)$, the deviation must be profitable. Thus, in an equilibrium in which the last agent doesn't join, there can be no prior agent that joins, so no agents join.

Suppose instead the equilibrium with $1 \leq m < N$ agents joining is such that the last agent joins, so $P \leq \Delta(m)$. The last agent that does not join obtains $u(m)$. If that agent deviates to joining, then all subsequent agents will continue to join (they are more likely to do so because more prior agents will have joined), so the agent's

deviation payoff is $b(m+1) - P$. Since $\Delta(m+1) \geq \Delta(m) \geq P$, the deviation is profitable. Thus, in an equilibrium in which the last agent joins, there can be no prior agent that does not join, so all agents join.

Next, we show by induction that when $b(\cdot)$ is weakly increasing, all agents join in $\Gamma(N, b(\cdot), u(\cdot), P)$ if $P \leq b(N) - u(0)$ and no agents join otherwise. In $\Gamma(1, b(\cdot), u(\cdot), P)$, the agent joins if and only if (iff) the platform sets $P \leq b(1) - u(0)$. Suppose the result holds for $N = n$ and any payoff functions $b(\cdot)$ and $u(\cdot)$ such that $b(\cdot)$ is weakly increasing and $u(\cdot)$ is weakly decreasing (so $\Delta(k) = b(k) - u(k-1)$ is weakly increasing in k). Consider $\Gamma(N = n+1, b(\cdot), u(\cdot), P)$. If the platform sets $P \leq b(n+1) - u(0)$, then after the first agent joins, the game is equivalent to $\Gamma(n, \tilde{b}(\cdot), \tilde{u}(\cdot), P)$, where $\tilde{b}(n) = b(n+1)$, $\tilde{u}(n) = u(n+1)$, and

$$P \leq b(n+1) - u(0) \leq b(n+1) - u(1) = \tilde{b}(n) - \tilde{u}(0).$$

The induction hypothesis implies all n agents join in this game. Thus, if the first agent joins in $\Gamma(n+1, b(\cdot), u(\cdot), P)$, its payoff will be $b(n+1) - P$, whereas if the agent doesn't join, its payoff will be at most $u(0)$. So the first agent joins because $b(n+1) - P \geq u(0)$, and so will all remaining n agents.

If instead the platform sets $P > b(n+1) - u(0)$ in $\Gamma(n+1, b(\cdot), u(\cdot), P)$, the first agent's payoff from joining will be less than $u(0)$. If the first agent doesn't join, then the remaining subgame is equivalent to $\Gamma(n, b(\cdot), u(\cdot), P)$ with

$$P > b(n+1) - u(0) \geq b(n) - u(0).$$

The induction hypothesis then implies no agent joins, so the first agent's payoff from not joining is $u(0)$, higher than its payoff from joining. The first agent doesn't join and neither will any of the remaining n agents. The result thus holds for $\Gamma(n+1, b(\cdot), u(\cdot), P)$.

By induction, the result holds for any $N \geq 2$. And we can then conclude that when $b(\cdot)$ is weakly increasing, the platform's optimal price is $P = b(N) - u(0)$ if $b(N) - u(0) > 0$, which attracts all agents, and any $P > 0$ if $b(N) - u(0) \leq 0$, which attracts no agents.

Finally, suppose $b(\cdot)$ is single-peaked and let m_0 be the lowest point at which $b(\cdot)$ reaches its maximum, i.e.

$$m_0 = \min \left\{ m \mid b(m) = \max_{1 \leq n \leq N} b(n) \right\}.$$

We prove that in the game $\Gamma(N, b(\cdot), u(\cdot), P)$, all agents join if $P \leq b(m_0) - u(0)$ and

no agents otherwise. Suppose $P \leq b(m_0) - u(0)$ and no agents join in equilibrium. The subgame starting with agent $N - m_0 + 1$ is equivalent to $\Gamma(m_0, b(\cdot), u(\cdot), P)$ (no prior agents have joined), where $b(\cdot)$ is weakly increasing from 1 to m_0 and $P \leq b(m_0) - u(0)$. The result above implies that all m_0 agents should join in the equilibrium of this game, which is a contradiction. Thus, all agents join in the equilibrium of $\Gamma(N, b(\cdot), u(\cdot), P)$ if $P \leq b(m_0) - u(0)$. If instead $P > b(m_0) - u(0)$, the only possible equilibrium is that no agents join. Indeed, if the first $N - 1$ agents do not join, the last agent does not join either because it obtains $u(0)$ by not joining and at most $b(m_0) - P$ by joining. Similarly, if the first $N - 2$ agents do not join, then neither do the last two. And so on, until we conclude that if the first agent does not join, then neither will the remaining $N - 1$ agents. So the first agent's payoff from not joining is $u(0)$, compared to at most $b(m_0) - P$ from joining. Thus, no agent joins in equilibrium if $P > b(m_0) - u(0)$. So the platform's optimal price in this case is $P = b(m_0) - u(0)$ if this is positive, and all agents join. Otherwise, no agents join.

10.6 Proof of Proposition 5

In the best possible equilibrium for the platform, the platform prices at $P = \Delta(N) > 0$ to all agents. Each agent joins because they expect all other agents to join. The platform attains its maximum feasible profits and all agents receive net payoff $u(N - 1)$, so they would be strictly better off without the platform.

Suppose $\Delta(1) > 0$. If the platform prices at $P = \Delta(1)$, each agent joins because they expect all other agents to join at this price, so they also join given they get $b(N) - P \geq b(N) - \Delta(1) \geq u(N - 1) > 0$ because $\Delta(\cdot)$ is weakly increasing. Note that, at this price, or any lower price, it is no longer an equilibrium for no agents to join, because even if an agent expects no other agent to join, it will want to join: i.e., $b(1) - P \geq u(0)$. So this is the worst the platform can obtain. And this is an equilibrium because if the platform deviates to $P > \Delta(1)$, then each agent expects no other agents to join, which is an equilibrium since under these expectations they would get $b(1) - P < u(0)$ from joining. In the equilibrium, agents get $b(N) - P = u(0) + (b(N) - b(1))$, so agents are strictly better off without the platform if and only if $b(1) > b(N)$.

If $\Delta(1) \leq 0$, the platform cannot profitably attract any agents in the worst equilibrium defined above.

11 References

- Armstrong, M. (2006) “Competition in Two-Sided Markets,” *RAND Journal of Economics*, 37(3), 668-691.
- Baye, M. R. and J. Morgan (2001) “Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets,” *American Economic Review*, 91, 454-474.
- Bergemann, D., A. Bonatti and T. Gan (2022) “The Economics of Social Data,” *RAND Journal of Economics*, 53(2), 263-296.
- Biglaiser, G., J. Crémer and A. Veiga (2022) “Should I Stay or Should I Go? Migrating away from an Incumbent Platform,” *RAND Journal of Economics*, 53(3), 453-483.
- Bursztyn, L., B.R. Handel, R. Jiménez-Durán and C. Roth (2025a) “When Product Markets Become Collective Traps: The Case of Social Media,” *American Economic Review*, 115(12), 4105-4136.
- Bursztyn, L., R. Jiménez-Durán, A. Leonard, F. Milojević and C. Roth (2025b) “Non-User Utility and Market Power: The Case of Smartphones,” NBER Working Paper 33642.
- Choi, J.P., D.-S. Jeon and B.-C. Kim (2019) “Privacy and Personal Data Collection with Information Externalities,” *Journal of Public Economics*, 173, 113-124.
- Farrell, J. and G. Saloner (1985) “Standardization, Compatibility, and Innovation,” *RAND Journal of Economics*, 16(1), 70-83.
- Farrell, J. and G. Saloner (1986) “Installed Base and Compatibility: Innovation, Product Preannouncements, and Predation,” *American Economic Review*, 76(5), 940-955.
- Gomes, R. and A. Mantovani (2025) “Regulating Platform Fees under Price Parity,” *Journal of the European Economic Association*, 23(1), 190-235.
- Halaburda, H. and Y. Yehezkel (2019) “Focality advantage in platform competition.” *Journal of Economics & Management Strategy*, 28, 49-59.
- Halaburda, H., B. Jullien and Y. Yehezkel (2020) “Dynamic competition with network externalities: Why history matters,” *RAND Journal of Economics*, 51(1), 3-31.

- Hall, J. V., and A. B. Krueger (2018). “An analysis of the labor market for Uber’s driver-partners in the United States,” *ILR Review*, 71(3), 705-732.
- Jullien, B., A. Pavan and M. Rysman (2021) “Two-sided markets, pricing, and network effects,” in K. Ho (Ed.), *Handbook of Industrial Organization* (Vol. 4, pp. 485-592). Elsevier B.V.
- Karle, H., M. Preuss and M. Reisinger (2026) “Selling on Recommender Platforms: Demand Boost versus Customer Migration,” Working paper.
- Katz, M.L. and C. Shapiro (1985) “Network Externalities, Competition, and Compatibility,” *American Economic Review*, 75(3), 424-440.
- Katz, M.L. and C. Shapiro (1986) “Technology Adoption in the Presence of Network Externalities,” *Journal of Political Economy*, 94, 822-841.
- Mitchell, S. (2023) “Amazon’s Monopoly Tollbooth in 2023,” Report of the Institute for Local Self-Reliance, Sep 21, 2023.
- Pashigian, B.P. and E.D. Gould (1998), “Internalizing Externalities: The Pricing of Space in Shopping Malls,” *Journal of Law and Economics*, 41(1), 115-142.
- Rey, P. and T. Vergé (2004) “Bilateral Control with Vertical Contracts,” *RAND Journal of Economics*, 35(4), 728-746.
- Rochet, J-C. and J. Tirole (2003) “Platform Competition in Two-Sided Markets,” *Journal of the European Economic Association*, 1(4), 990-1029.
- Segal, I. (1999) “Contracting with Externalities,” *Quarterly Journal of Economics*, 114(2), 337-388.
- Segal, I.R. and M.D. Whinston (2000) “Naked Exclusion: Comment,” *American Economic Review*, 90(1), 296-309.
- Sullivan, M. (2024) “Price Controls in a Multi-Sided Market,” Working paper.
- Wang, C. and J. Wright (2020) “Search Platforms: Showrooming and Price Parity Clauses,” *RAND Journal of Economics*, Vol. 51(1), 32-58.