Platform Traps*

Andrei Hagiu[†] and Julian Wright[‡]

April 9, 2025

Abstract

Rational agents can be induced to join a platform despite being better off without it, when each agent's outside option worsens as other agents participate on the platform. Such platform traps apply to sellers on dominant marketplaces and users on social networks, among other settings. We provide a dynamic theory of a platform trap in which the platform exploits a form of collective action problem through one (or more) of the following features: (i) the ability to make dynamic price adjustments; (ii) the combination of on-platform and off-platform negative externalities; (iii) favorable equilibrium selection when platforms make private offers.

1 Introduction

Many businesses are increasingly reliant on major platforms to reach customers, such as third-party sellers on Amazon, hotels on Booking.com and Expedia, and restaurants on DoorDash, Grubhub, and Uber Eats. As consumers shift to these platforms instead of shopping directly with sellers, concerns arise about sellers' increasing dependence on, and exploitation by, the platforms. This can involve saddling sellers with higher or new types of fees, as well as an increasing requirement to advertise if they want to be discovered on the platforms they depend on. But why then would businesses

*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, 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

join such platforms in the first place, knowing they will ultimately be held up in this way?

In this paper, we provide a general theory of why rational agents join a platform even when they are worse off as a result. We call this a platform trap. The theory applies to sellers participating on a marketplace that attracts consumers who previously searched directly for these sellers, but also to other settings: neighborhood shops opening outlets in a suburban shopping mall (drawing traffic away from their standalone outlets); brands supplying Walmart (diverting consumers from the brands' existing channels); individuals participating on an online social network (rather than relying on direct connections); news organizations signing deals with Facebook and Google to provide access to their content (resulting in these platforms becoming the primary gateway for accessing news). More generally, a platform trap can arise whenever novel technologies allow new platforms to emerge (e.g., Apple's Vision Pro), whose adoption inherently substitutes (at least partially) for transactions or interactions on legacy technologies (e.g., PCs, TVs).

The key ingredient in our theory is that each agent's outside option is not constant, but instead deteriorates as other agents join the platform, which then increases each agent's incentive to join. For example, when sellers list on a large marketplace like Amazon's, they divert more buyers to the platform, diminishing the profitability of sellers' outside option of selling directly (e.g., on their own websites). Similarly, a brand's sales via mom-and-pop stores may fall as more brands sell through Walmart if this redirects consumer traffic to Walmart from mom-and-pop stores.[1] As more individuals in a group join a dominant social network, those who remain outside (relying instead on direct connections) risk social exclusion and "fear of missing out". And when news outlets partner with Facebook and Google, others who don't may see less traffic on their own websites, since users will increasingly access news via these platforms.

Due to the negative externality each agent's platform participation decision imposes on other agents' outside options, agents face a collective action problem: each of them finds it individually attractive to join the platform but in doing so leads to a worsening of all agents' outside option. The platform exploits this by setting correspondingly higher prices, which can ultimately result in all agents being worse off. The problem is that while each individual agent takes into account the possibility that they may be better off not joining if that induces other agents to also not join, when they do join, they do not account for worsening the position of other agents who have not yet joined.

[1]The Walmart example shows our theory applies beyond standard platforms to other intermediaries (such as large retailers).

Our theory involves a finite number of agents who each receive an offer and decide sequentially whether to join the platform. Both the benefit from joining the platform and the benefit from the outside option can depend on how many agents join. We assume the benefit from the outside option is a weakly decreasing function of the number of agents participating on the platform and strictly decreasing at least at some point. We also assume the difference between the benefit of participating on the platform and the benefit of the outside option is weakly increasing as more agents join.

The sequential nature of the agents' participation decisions raises the possibility of an agent being pivotal — its participation decision affects whether subsequent agents can be profitably induced to participate. In this case, the platform must first attract a critical mass of participants before it can offer more value for each subsequent agent than the outside option. Once a platform achieves this critical mass, each subsequent agent knows, regardless of what it decides, the platform will be able to induce all remaining agents to join by appropriately adjusting its prices, and so it is willing to join even when it is charged a price based on all other agents joining.

As a baseline, we derive equilibrium prices and a sufficient condition for the platform trap to arise when the platform has full pricing power and can perfectly price discriminate among agents, and each agent can observe previous agents' decisions when making its own. When no agents are pivotal, the platform extracts the maximum surplus from each agent in equilibrium 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, provided the platform provides enough value, it attracts all agents with increasing prices, such that each agent is left with the outside option arising when only the agents joining before it participate. For intermediate cases, the platform sets increasing prices until the marginal agent joining is no longer pivotal, at which point it sets the same maximal price for all remaining agents as in the case without any pivotal agents.

In this baseline setting, whenever agents participate, all agents would be weakly better off and some strictly better off without the platform. In other words, they would be better off collectively committing to boycott the platform. When payoffs are such that all but a last set of agents are pivotal, the platform will sometimes need to subsidize the initial agents to induce them to participate. As a result, the platform might end up attracting agents to participate even when doing so is inefficient. Since the platform also attracts agents to participate whenever it is efficient, the problem of the platform trap is quite distinct from (and more pervasive than) traditional concerns about inefficiency.

Based on the logic of our baseline model, one might think that dynamic price adjustments are essential for a platform trap to emerge. Surprisingly, we show that even if the platform is only able to set a single price to all agents, a platform trap can arise if each agent's benefit from the platform is single-peaked and reaches its maximum before all agents have participated. The price is based on this maximum benefit, but provided agents' outside option decreases fast enough in the number of participating agents, all agents join and would be better off without the platform. Thus, a combination of negative externalities on and off the platform can generate a platform trap even if dynamic price adjustments are not possible.

We examine other factors — such as heterogeneous agents and competing platforms — that amplify or moderate the extent of the platform trap, as well as its boundaries: when would the platform make some or all agents better off? A key factor is imperfect information: agents do not observe others' offers or decisions. This gives rise to multiple equilibria. When equilibrium selection favors the platform, the trap emerges for a wider set of cases than in the full-information baseline, reflecting that agents can no longer be pivotal when their participation decision is not observed. On the other hand, with unfavorable equilibrium selection, a platform trap can still arise but only if externalities on the platform are negative.

2 Literature review

The concept of a platform trap is closely related to the idea of collective traps introduced and examined by Bursztyn et al. (2024). They provide experimental evidence showing that total consumer welfare from the availability of social media platforms Instagram and TikTok is negative among their student sample, even though these individuals are willing to pay to use them. They propose a mechanism to explain these results in which users impose a negative externality on non-users (e.g., via social exclusion or a fear of missing out). As more people join, marginal users may participate to avoid the worse outcomes of non-participation, even if they experience negative overall utility from the platform. This underlying mechanism is consistent with our own, in which the outside option weakens as platform participation increases. However, while Bursztyn et al. offer a conceptual framework to interpret their results, they do not formalize how (or under what conditions) a collective trap arises through individual choices, nor do they consider how platforms might induce such traps (e.g., via dynamic pricing). This reflects their different focus, which lies in empirically demonstrating and quantifying the collective trap rather than providing a theory for it.

On the theory side, our paper relates to settings in which a principal (the platform) contracts with multiple agents subject to cross-agent externalities. The most general formulation is that of Segal (1999), which encompasses many applications. One application is Katz and Shapiro (1986), in which a sponsored technology (the principal) that licenses its technology competes with an unsponsored technology (the outside option) for users, where both technologies are subject to network effects. Another application is Segal and Whinston (2000), in which an incumbent (the principal) signs exclusive contracts with agents, denying an entrant the scale needed to enter, thereby harming agents who do not contract with the principal. While each of these settings, including Segal's general framework, differs from ours in various modeling choices — for instance, they mostly focus on simultaneous decision-making by agents whereas we use a sequential baseline model — they also pursue a fundamentally different research question. They analyze the possible inefficiencies in the contracts offered by the principal, whereas we ask whether agents would be better off without the principal. As we will show, the determinants of inefficient outcomes are completely different from those of a platform trap.

A platform trap also differs fundamentally from the inefficiencies studied in the old literature on network externalities (e.g., Farrell and Saloner, 1985, and Katz and Shapiro, 1985), which shows that users can become locked into an inferior technology like the classic QWERTY keyboard example (excess inertia), or can inefficiently adopt a new technology due to excess momentum (Farrell and Saloner, 1986, and Katz and Shapiro, 1986). Even if only one of the technologies is sponsored (analogous to the platform in our setting, with the non-sponsored technology representing the outside option), again these papers focus on a completely different question — whether technology adoption can be inefficient, rather than whether agents would be better off without the sponsored technology altogether.

Our paper relates to previous literature on platforms in dynamic contexts (see Jullien et al., 2021). Arguably, closest to our approach from this literature is Biglaiser et al. (2022), in that their agents also make participation decisions sequentially. But their focus is quite different. They study the movement of agents from an incumbent platform to an entrant platform to understand the determinants of incumbency advantage.

Finally, there is a literature that captures the idea that platforms intensify competition between sellers (e.g., Baye and Morgan, 2001; Galeotti and Moraga-González, 2009; Wang and Wright, 2020, 2025; Gomes and Mantovani, 2025; and Hagiu and Wright, 2024) compared to a direct channel alternative. In such settings, sellers could be collectively better off without the platform because the platform lowers their margins due to intensifying competition between them. Our mechanism does not rely on the commodization of sellers that underpins the harm to sellers arising in these settings. Moreover, these works close down any kind of dynamic decision making to explain the emergence of an equilibrium where sellers decide to join despite ending up worse off, which is one of our key contributions.

3 Baseline model

There is a single platform and $N \geq 2$ agents. Each agent decides whether to join the platform or opt for a non-strategic alternative — the outside option (e.g., a direct sales channel in a marketplace, peer-to-peer connections in a social network, or a competitively provided product or network priced at cost). The game consists of a single period (i.e., payoffs are only realized once), with agents making sequential participation decisions as defined below, although as we will note, the results apply for any number of periods.

An agent's benefit from joining the platform when n total agents participate is $b(n) > 0$. The platform may exhibit either positive network effects (increasing $b(.)$) or negative externalities across participants (decreasing $b(.)$). For example, if agents are sellers on a marketplace, the negative effect of increased competition when more sellers participate may overwhelm the positive network effect of attracting more buyers. An agent that does not join the platform (i.e., chooses the outside option) earns payoff $u(n)$, where n is the number of agents who participate. We assume $u(.)$ is non-negative, weakly decreasing in its argument, and is such that $u(N-1) < u(0)$, implying an agent that remains outside is made strictly worse off if all other agents join the platform. Note these assumptions imply $u(0) > 0$.

Agents receive offers from the platform in a fixed sequence. Since agents are homogeneous in our baseline model, the order does not matter and can be exogenous or chosen by the platform.[2] The offer to agent $k \in \{1, .., N\}$ is a fixed price P^k to join. There is no discounting. If $n \leq N-1$ other agents participate, agent k's payoff is $b(n+1) - P^k$ if they join the platform and $u(n)$ if they do not. Each agent makes its decision upon receiving its offers, and all payoffs are realized after all N agents have made their decisions.

[2] In the extension with heterogeneous agents, the order of offers will matter, in which case we assume the platform chooses the order.

Define the surplus function

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

We assume $\Delta(.)$ is weakly increasing. This holds if $b(.)$ is weakly increasing or, if decreasing, does not decline as fast as $u(.)$. This ensures the marginal agent's incentive to join the platform weakly increases with the number of participating agents, implying (as we will show) that in equilibrium, either all agents join or none do. We also assume $\Delta(N) > 0$, so that participation on the platform yields positive surplus (before pricing) when all other agents join. Otherwise, the platform could not profitably attract any agents.[3] To avoid trivial equilibria, we assume that the platform must earn positive profits to operate.

We make two additional assumptions for our baseline model, both of which will be relaxed later: (i) each agent k can observe the decisions (join or not) of the previous $k-1$ agents in the sequence,[4] and (ii) the platform can set a different price for each agent. Formally, this is a dynamic game of complete and perfect information with N stages. The relevant solution concept is subgame perfect equilibrium, henceforth simply "equilibrium".

3.1 Discussion

Our baseline model assumes agents decide sequentially, which ensures a unique equilibrium. This rules out the possibility that a platform trap arises merely from "bad" equilibrium selection by agents. If instead agents decide simultaneously, multiple equilibria can emerge. We address this in Section 5.2, where we show that a platform trap can still arise under any equilibrium selection — though it requires an additional restriction on $b(.)$ when the worst equilibrium (from the platform's perspective) is selected.

There are two interpretations of our single-period model.

One is that it represents a single period of a repeated game (with finite or infinite periods), where the platform sets new prices in each period, agents can freely revise their participation decisions, and payoffs are discounted across periods. Because

[3]Even if the platform successfully attracted the first $N-1$ agents to join, it would not be able to charge a positive price to the last agent, given the last agent would get $b(N) - P$ from joining and $u(N-1)$ from not. And since $\Delta(n)$ is weakly increasing in n, repeating the same logic implies the platform will not attract any earlier agent either.

[4]As long as agent k can observe the decisions of the $k-1$ previous agents, it is irrelevant whether agent k also observes the prices they were offered. That being said, to make the baseline setting one of perfect information, we assume agents also observe these prices.

nothing links the decisions in one period to another, the equilibrium analysis can be conducted separately for any given period as we do in Section 4 below. Price dynamics across periods (in addition to the within-period dynamics we focus on) would only arise from exogenous changes in the number of agents (e.g., more agents become aware of the platform over time). Such changes, together with any exogenous changes in payoffs, can be handled in our single-period setting as straightforward comparative static exercises with respect to N (as we do at the end of Section 4). Similarly, extensions with additional stages in which agents can reverse earlier decisions can also be accommodated. We elaborate on these extensions in Online Appendix A.

Alternatively, the single-period model can be interpreted as capturing the full decision-making process, with agents making irreversible joining decisions. The payoffs $b(n)$ and $u(n)$ then represent the present discounted value of future benefits, over many (possibly infinite) periods. The absence of discounting across agents' decisions reflects the assumption that joining decisions are made over a short time frame, relative to the longer horizon over which benefits accrue.

In the baseline setting, the platform can set a different price for each agent, including negative prices if necessary. We relax this full price discrimination assumption in Section 5.1. Negative prices can be interpreted as subsidies or fixed participation benefits when the platform offers its services for free, while positive prices could represent exposure to advertising (e.g., in the case of social media platforms).

In the baseline model, we make the realistic assumption that when making an offer to a given agent, the platform cannot commit to the prices it will charge to subsequent agents. This shows that the platform trap does not depend on commitment power. In fact, as we show, the platform generally benefits from this lack of commitment. However, there are exceptions. For instance, if agents are pivotal and $b(.)$ is decreasing, committing to a sequence of decreasing prices can render earlier agents non-pivotal and increases the platform's profits and the scope for a platform trap. We prove this in Online Appendix B. A similar benefit arises if the platform can commit to prices contingent on the final number of participants, which is valuable when facing pivotal agents or when the platform is constrained to set the same price to all agents, as shown in Section 5.1.

Finally, in some interpretations of our model, agents who join the platform may not have to give up their outside option entirely. In such cases, $b(.)$ incorporates any remaining payoff they get from the outside option. For example, in a marketplace setting, sellers listing on the platform may still sell directly through their own websites. We illustrate this with a microfounded model in Online Appendix C, where buyers are non-strategic. This framework also allows us to assess whether buyers, like sellers, may be harmed by a platform trap. We also show how $b(.)$ can be decreasing (due to seller competition), while the surplus function $\Delta(.)$ remains increasing.

4 Analysis of the benchmark model

To build intuition for how the equilibrium is determined, consider two extreme cases: one where no agent is pivotal, and another where every agent (except the last) is pivotal. An agent is pivotal if the platform must attract the agent to profitably induce later agents to join. For example, with two agents ($N = 2$), suppose the platform needs both to join for it to offer positive value relative to the outside option; i.e., $\Delta(2) > 0 \geq \Delta(1)$. Then the first agent is pivotal: if it rejects the offer, the second agent cannot be profitably attracted. But if the first agent joins, the platform can profitably attract the second agent too. Now consider three agents ($N = 3$) under the same assumptions. The first agent is no longer necessarily pivotal, because even if it rejects the offer, the platform may still profitably attract the remaining two. However, if the first agent rejects its offer, the situation reduces to the $N = 2$ case, making the second agent pivotal.

For the first extreme where no agent is pivotal, assume $\Delta(1) > 0$, which implies $\Delta(n) > 0$ for all $n \geq 1$. This means the platform always creates more value for an individual agent than the outside option, regardless of how many other agents are expected to join. As a result, the platform can always profitably induce any agent to participate.

Since the last agent can be induced to join even if no one before them has joined, and since surplus $\Delta(n)$ is weakly increasing, each agent expects all subsequent agents to join regardless of their own decision. Thus, if $N' < N$ prior agents have joined, the k-th agent 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 would cause the agent to reject the offer, leading to zero revenue from them and lower revenue from all subsequent agents. By setting $P = P(N')$ at each stage, the platform ensures that all N agents 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 agents decided simultaneously and coordinated on the equilibrium most favorable for the platform; i.e., all agents join whenever that is an equilibrium. In that case, each agent expects the other $N-1$ agents 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 agent indifferent between joining and not, when they expect all others to participate.

Now consider the opposite extreme, where every agent (except the last) is pivotal. Specifically, suppose

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

which, given that $\Delta(.)$ is weakly increasing, implies $\Delta(k) \leq 0$ for all $k \leq N-1$. In this case, failure to attract even a single agent makes it impossible for the platform to profitably attract any of the remaining agents: the last agent would prefer the outside option when charged a positive price, so the platform has no incentive to attract them, so it cannot attract the second-to-last agent either, and so on. Thus, every agent except the last is pivotal. To induce participation, the platform must offer each agent at least their outside option assuming no subsequent agents join. Although all agents are induced to join in equilibrium, the price increases along the sequence, from $P^1 = b(N) - u(0)$ for the first agent to $P^N = b(N) - u(N-1) = \Delta(N)$ for the last agent. This is profitable for the platform provided $b(N) > \frac{1}{N}\sum_{k=1}^{N} u(k)$. Otherwise, the platform will choose not to induce any participation.

Between the two extremes, i.e. when $\Delta(1) \leq 0 < \Delta(N-1)$, some (but not all) agents may be pivotal. Specifically, we show in the proof of the next proposition that when some agents are pivotal, there is a threshold $k_0 \geq 1$, such that the first $k \leq k_0$ agents are pivotal and are charged prices $P^k = b(N) - u(k-1)$. Once these agents are on board, the remaining $N - k_0$ agents are no longer pivotal and are each charged $P^k = b(N) - u(N-1) = \Delta(N)$. These prices are consistent with those in the extreme cases discussed above. A scenario with $1 \leq k_0 \leq N-1$ can be interpreted as one where the platform must first attract a critical mass k_0 of agents to ensure it offers more value than the outside option for each subsequent agent regardless of whether those agents ultimately join.[5]

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

[5]The proof, as with all proofs that are not covered in the text, is relegated to the Appendix.

participation are

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

If $\sum_{k=1}^{N} P^k > 0$*, these are the platform's optimal prices and all* N *agents join, with the first* k_0 *agents being pivotal. If instead* $\sum_{k=1}^{N} P^k \le 0$*, then the platform cannot profitably attract any agents, so in equilibrium no one joins. In any equilibrium where agents join, they would all be weakly better off without the platform, and at least some would be strictly better off. A sufficient condition for all agents to join is*

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

We refer to the scenario in which agents join the platform despite being better off without it as a platform trap. This arises because when one agent joins, the value of the outside option falls for all others, increasing their incentive to participate. The platform thus exploits a collective action problem among agents. As shown in Proposition 1, whenever a platform trap occurs, at least some agents (always including the last to join) are strictly worse off than if the platform didn't exist, while the rest are weakly worse off. If $u(1) < u(0)$, the result is stronger: all agents are strictly worse off, except possibly the first if they are pivotal.

Note that (1) is only a sufficient condition for a platform trap to arise. A necessary condition is $\Delta(N) > 0$, which ensures the platform can attract any agents, and which we assumed holds.

For any given model primitives — the number of agents N, and their payoff functions $b(.)$ and $u(.)$) — one of two outcomes arises. Either the platform finds it optimal to attract all agents, yielding a unique equilibrium in which all join and prices satisfy the properties defined in Proposition 1; or the platform finds it optimal to attract no agents, in which case there is a continuum of equilibria (in prices) with no agent participation. When the platform attracts all agents, each agent $k \in \{1, .., N\}$ is either pivotal and charged $P^k = b(N) - u(k-1)$, or non-pivotal and charged $P^k = b(N) - u(N-1)$. This is because at any stage, the subgame has only two possible equilibria: either all remaining agents join or none do (as shown in the proof of Proposition 1). However, fully characterizing (in terms of model primitives) how many agents are pivotal becomes increasingly complex as N grows. To illustrate this and give a sense of how equilibrium prices relate to model primitives, in Online Appendix D, we provide the full characterization of equilibrium for the cases $N = 2$ and $N = 3$.

The characterization in Proposition 1 is sufficient for drawing conclusions about a platform trap. The two corollaries below fully spell out conditions under which the platform attracts agents, and thus when a platform trap emerges, in the two extreme cases discussed earlier.

Given the prices laid out in Proposition 1, an agent's expected payoff from joining the platform is either $u(N-1)$ if it is not pivotal, or $u(k-1)$ if it is pivotal. Since $u(k-1) \leq u(0)$ for all k and $u(N-1) < u(0)$, all agents would be weakly better off without the platform — i.e., if they collectively boycotted it — and at least the last agent would be strictly better off.[6] If at least one agent is pivotal, then the first agent to join earns $b(N) - (b(N) - u(0)) = u(0)$, making them indifferent with or without the platform.

Notably, the sufficient condition (1) in Proposition 1 may allow the platform to attract all agents even when $b(N) < u(0)$; i.e., even when the platform reduces welfare. In this case, both individual agents and total welfare (including the platform's profit) are worse off due to the platform's existence. This outcome hinges on the negative externality imposed on the outside option. Without it, a welfare-reducing platform could not profitably attract agents. However, the platform also attracts agents and creates a platform trap whenever it is efficient, so the problem of the platform trap is quite distinct from (and more pervasive than) traditional concerns about inefficiency.

We now fully characterize the conditions under which the platform profitably attracts all agents in the two extreme cases discussed above.

Corollary 1. *If* $\Delta(1) > 0$ *or* $b(N-1) > \frac{1}{N-1}\sum_{k=1}^{N-1} u(k-1)$*, there exists a unique equilibrium in which the platform attracts all* N *agents and charges* $P^k = \Delta(N) > 0$ *for all* $k \in \{1, .., N\}$*. 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 (1) holds. Under these conditions, there exists a unique equilibrium in which the platform charges* $P^k = b(N) - u(k-1)$ *to agent* $k \in \{1, .., N\}$*. All agents would be weakly better off without the platform, with at least some strictly better off.*

Corollary 1 characterizes equilibrium prices when no agent is pivotal. Whenever $N > 2$ and $b(.)$ is weakly increasing, the condition $b(N-1) > \frac{1}{N-1}\sum_{k=1}^{N-1} u(k-1)$ is sufficient, because $\Delta(1) > 0$ is a stronger condition in that case. Indeed, $b(N-1) >$

[6] We implicitly assume each agent's payoff without the platform equals their payoff $u(0)$ when no one joins. If agents' payoffs without the platform differ from $u(0)$, then our results should be interpreted as what happens if agents collectively boycott the platform.

$\frac{1}{N-1}\sum_{k=1}^{N-1} u(k-1)$ rules out the possibility of pivotal agents because even if one agent does not join, the platform can still profitably attract the remaining agents. On the other hand, Corollary 2 characterizes equilibrium prices when all agents are pivotal, and provides a sufficient condition for this case. Proposition 1 sometimes requires the platform to subsidize pivotal agents. This will happen when it is inefficient; i.e., when $b(N) < u(0)$.

We now return to the interpretation of our model (Section 3.1) as capturing one of many periods over which the game is repeated and over which payoffs are realized. A natural assumption is that the number of agents aware of the platform grows over time, so the number of available agents N_t in period t increases with t. This raises the question: how do prices, agents' net payoffs and platform profits evolve as N increases?

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 are weakly increasing in N, while prices for pivotal agents may be increasing or decreasing, depending on whether $b(.)$ is increasing or decreasing. However, net payoffs for all agents always weakly decline with N: when more agents are present, the total negative externality on the outside option is larger, making the platform trap more severe. Furthermore, it is intuitive that as the number of agents increases, each becomes less critical, making it easier for the platform to attract participation and reducing the likelihood of pivotal agents.

These results are consistent with a narrative according to which, as a platform becomes more well-known over time (so more agents enter the market), the trap deepens. The platform becomes less likely to rely on subsidies over time, its prices increase and agent net payoffs decline.

5 Additional platform trap mechanisms

In our baseline setting, we made certain assumptions on pricing and information to obtain our main platform trap result. These included that the platform can make a different take-it-or-leave it offer to every agent, and that agents have perfect information on the platform's and previous agents' decisions. In this section, we show that relaxing these assumptions sometimes limits and sometimes expands the scope for platform traps. It can also give rise to new platform trap mechanisms.

5.1 Different pricing assumptions

It is intuitive that the platform trap hinges on the platform having sufficient pricing power. Our baseline setting assumes it makes take-it-or-leave-it offers. To illustrate, in Online Appendix E we analyze the baseline setting with $\Delta(1) > 0$, where no agent is ever pivotal (the strongest version of the trap), and characterize equilibrium prices when the platform engages in Nash bargaining with each agent (sequentially). Not surprisingly, if agents have enough bargaining power, the platform trap may be avoided. Otherwise, some (the last agents to negotiate) or all agents would be better off without the platform.

We now turn to a more interesting question: what happens when the platform can no longer perfectly price discriminate across agents? One might think that dynamic price adjustments are essential for the platform trap to arise. Perhaps surprisingly, we show this is not the case. Even when the platform must offer a uniform price to all agents — i.e., the same price offered to the first agent must be offered to all others — a platform trap may still emerge. Since the platform must earn a profit, the uniform price must be positive, ruling out subsidies. Yet a platform trap can still occur when negative externalities dominate after enough agents have joined.

Proposition 2. *Suppose $b(.)$ is a single peaked function and the platform must charge the same price to all agents. There exists a unique equilibrium in which the platform attracts all N agents with the price $P = \max_{1 \leq n \leq N} \{b(n)\} - u(0)$, provided this is positive. In this case, agents are indifferent about the platform existing if $b(N) = \max_{1 \leq n \leq N} \{b(n)\}$ and would be strictly better off without the platform if $b(N) < \max_{1 \leq n \leq N} \{b(n)\}$. Alternatively, 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 a single fixed price to all agents and $b(.)$ is weakly increasing.[7] In this case, at the price stated in Proposition 2, agents are indifferent about the platform's existence because the platform cannot adjust prices to later agents based on the decisions of earlier agents, making each agent effectively pivotal (i.e., their decision could determine whether others join).

[7] As we show in the next section, a platform trap can re-emerge in this scenario when agents cannot observe anything but their own offers (or they decide simultaneously rather than sequentially), though it then depends on equilibrium selection.

However, when $b(.)$ is decreasing, or initially increasing and then decreasing, the platform trap reemerges despite the platform's inability to dynamically adjust prices. This occurs because agents keep joining the platform as their outside option worsens, even though each additional participant reduces the platform's benefit for others. The trap is now driven by both $b(n)$ and $u(n)$ declining in n, with $u(n)$ falling faster.

In fact, this new mechanism makes it possible that some agents are worse off under uniform pricing than under price discrimination. Recall that in the baseline model, if agent k was pivotal, the price it paid was $b(N) - u(k-1)$. With a single price, agent k is charged $b(n^*) - u(0)$, 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 uniform pricing.[8] This outcome requires that $b(n)$ declines sufficiently from its peak to $b(N)$, the point where all agents participate. Only early agents that were pivotal in the baseline may be worse off when the platform can no longer price discriminate. By contrast, non-pivotal agents (including the last one) are always weakly better off under uniform pricing. Indeed

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

Proposition 2 shows that when $b(.)$ is weakly increasing, a key driver of the platform trap is the platform's ability to adjust prices for future agents in response to deviations from the equilibrium path, in order to induce them to participate. However, if $b(N-1) > u(0)$, the platform trap only requires a very limited form of price discrimination. Specifically, suppose the platform can change its price only once: it sets P_1 for the first N_1 agents, and then P_2 for the remaining $N_2 = N - N_1$ agents. As in the baseline model, the platform cannot commit to future prices, so P_2 is chosen only when offered. The platform can also choose N_1, i.e. the point at which the price changes.

With this set-up, if $b(N-1) > u(0)$, the platform can replicate the maximal platform trap with full price discrimination from Corollary 1, despite having significantly less price flexibility. The platform offers the first agent $P_1 = \Delta(N) > 0$. From Proposition (2), the agent knows that rejecting this offer will lead the platform to set a uniform price $P_2 = b(N-1) - u(0)$ to all remaining agents, which ensures their participation. And if the first agent accepts, it becomes even easier to attract the others. Anticipating

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

that either way the remaining $N-1$ agents will join, the first agent accepts. The same logic applies for all subsequent agents: each of them accepts $P_1 = \Delta(N)$, so ultimately all agents join, just as in Corollary 1. Thus, we have proven:

Proposition 3. *If the platform can change its price only once, but chooses when to do so, and if $b(N-1) > u(0)$ and $b(.)$ is weakly increasing, there exists a unique equilibrium in which all agents join at a uniform price $P_1 = P_2 = \Delta(N) > 0$. All agents would be strictly better off without the platform.*

Comparing this result with Proposition 2 highlights that one of the key ways a platform trap can be generated is a platform's ability to adjust its prices to attract future agents. This ability can arise from the platform having some discretion to adjust its prices (Proposition 3) given we assumed commitment was not possible.

If instead, the platform can commit to an optimal mechanism, then the platform can always achieve maximal payoffs even without agent-specific pricing. It can do so by using second-degree price discrimination; i.e., by committing to a contingent price function $P(j)$, where all participating agents pay the same price based on the total number j of agents that join.

Proposition 4. *If the platform can commit to a single contingent pricing function offered to all agents, it maximizes profits by setting $P(j) = \Delta(j)$ for $j \in \{1, ..., N\}$. This induces all N agents to join and yields a profit of $N\Delta(N)$. All agents would be strictly better off without the platform.*

When some agents are pivotal under full price discrimination, the platform earns higher profits if it can commit to a single contingent pricing function.[9] However, this approach may require the platform to commit to selling at a loss off the equilibrium path; i.e., if only $n < N$ agents join and $\Delta(n) < 0$. Moreover, even if $\Delta(1) > 0$, implementing contingent pricing mechanisms is likely to be difficult in practice.

5.2 Private offers and simultaneous moves

In the baseline setting, we assumed agents had perfect information: each observed the entire history of offers and decisions before making their own decision. In this subsection we make the opposite assumption: agents only observe their own offers and

[9]Similarly, in Online Appendix F, we show that the platform can replicate the maximum platform trap from Proposition 4 if it can commit to run an auction in each stage for all unsigned agents, inducing them to compete for joining earlier.

nothing else (we call these private offers). Thus, agents do not know their position in the offer sequence or how many agents have previously accepted or rejected offers. As a result, multiple equilibria can arise depending on agents' off-equilibrium path beliefs, and the relevant equilibrium concept in this section is perfect Bayesian equilibrium (PBE).

In order to rule out PBE supported by unreasonable off-equilibrium path beliefs, we make the following refinement. Since agents do not observe whether others have already received offers, or the nature of any offers or decisions, there is no reason for the platform to change its offers to subsequent agents when a given agent changes their decision about whether to join or not. We therefore require that given any price it faces, each agent believes its decision to join or not doesn't affect the prices the platform will set to the subsequent agents. Since this refinement imposes a type of consistency on the agents' off-equilibrium beliefs, we refer to the resulting equilibria as strong PBE.[10] Given agents are homogeneous, there is no benefit to the platform from setting different prices, and the resulting strong PBE prices turn out to always be uniform.

Another reason for focusing only on the strong PBE we have defined is that they are also the equilibria that would be selected if we reorder the moves of the players so that all agents make their joining decisions (possibly simultaneously) after the platform has made its private offers to all agents. In that case, the platform cannot change its prices based on agents' decisions. Thus, the strong PBE we focus on do not depend on the exact ordering of the moves of players (provided, of course, the platform sets each agent's private offer before the agent decides).

The following proposition characterizes the best and the worst such equilibria for the platform.

Proposition 5. *Suppose agents observe only their own offers. There exists a continuum of strong perfect Bayesian equilibria. These are also the equilibria of the game in which the principal makes its private offers to all agents in the first stage and agents all decide simultaneously whether to accept or not in the second stage.*

1. *In the best equilibrium for the platform, each agent k is charged $P^k = \Delta(N)$, all N agents join, and platform profits are $N\Delta(N)$. All agents would be strictly better off without the platform.*

[10] Strong PBE only differ from weak PBE when $b(.)$ is decreasing, in which case there exist weak PBEs that involve even worse outcomes for the platform (the best equilibrium for the platform remains unchanged). However, such outcomes can only be supported by unreasonable beliefs: when an agent deviates from the equilibrium path by rejecting an offer, it must believe the platform will also deviate by increasing its prices to subsequent agents such that they will also reject its offers.

2. *If $\Delta(1) > 0$, in the worst equilibrium for the platform, each agent k is charged $P^k = \Delta(1)$, 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 would be strictly better off without the platform 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.*

The proposition characterizes the two most extreme equilibria from the platform's perspective when agents face private offers. Depending on equilibrium selection, the platform trap arises more broadly than in the baseline setting (case 1 in the proposition) or is more limited (case 2 when $b(.)$ is weakly increasing). More generally, the comparison depends on the equilibrium selection in this range, how many agents are pivotal in the baseline and the shape of the $b(.)$ function.[11]

The key reason the platform trap may apply more broadly is that agents are no longer pivotal: since individual participation decisions are unobservable, a rejection by one agent cannot influence subsequent agents. As a result, the best equilibrium for the platform may be even better with private offers than when agents' decisions are fully observed, extending the conditions under which the platform trap arises. Thus, even when the sufficient condition in Proposition 1 holds so the platform induces all agents to join under perfect information, agents are weakly worse off with private offers when the platform's best equilibrium is selected (agents that were previously pivotal are now strictly worse off).

Perhaps surprisingly, the full platform trap can still arise, even under the worst equilibrium for the platform. This happens when externalities on the platform are negative, i.e., $b(N) < b(1)$. Since the set of strong PBE in this private offers setting is equivalent to the set of equilibria when the platform sets a single price to all agents, and agents decide simultaneously, the results imply that a platform trap can arise even without any dynamics or price discrimination.

In the worst equilibrium for the platform, agents may still benefit from the platform, as they face a relatively low price ($P^k = \Delta(1)$) when on-platform network effects are positive. However, if $b(N) < b(1)$, all agents would be better off without the platform. The platform trap arises here because agents continue to join the platform as their outside option deteriorates, even though increased participation eventually

[11]For example, suppose $b(.)$ is decreasing and in the baseline all agents are pivotal. Then total agent payoffs 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.

reduces the platform's benefit for agents that join. Each agent is thus willing to pay $\Delta(1) = b(1) - u(0)$ to join, knowing they will end up getting $b(N)$ rather than $b(1)$, because rejecting the offer leaves them with $u(N-1)$, which is worse.[12]

Based on the above analysis, we can ask: what type of offers would the platform prefer if it could commit to a specific structure: public or private? The possible payoffs are given by Proposition 1 (public offers) and Proposition 5 (private offers). If the platform's best equilibrium is selected whenever multiple equilibria exist, it can't do better than with private offers (though public offers yield the same profit when no agents are pivotal). If instead the platform's worst equilibrium is selected, it does best with public offers when $b(.)$ is increasing. Otherwise, whether the platform does best with public offers or private offers depends on the number of pivotal agents under public offers and the shape of the $b(.)$ function.

6 Extensions

In this section we relax the assumption that agents are homogenous and that the platform is a monopoly, showing that platform traps can still arise in these more complicated environments. For brevity, the proofs for this section are contained in Online Appendix G.

6.1 Heterogeneous agents

So far we have assumed all agents are homogeneous. We now analyze a simple form of heterogeneity by adding one "superstar" agent alongside N regular (symmetric) agents. This raises the question: if the platform wants to engineer a platform trap, which type of agent should it attract first?

We assume the superstar is equivalent to $S > 1$ regular agents. Thus, when the superstar agent and $n \leq N$ agents join the platform, the payoff on the platform is $b(S+n)$ and the outside option payoff is $u(S+n)$. If the superstar does not join, payoffs are just the usual $b(n)$ and $u(n)$. We adjust our baseline assumption that every agent receives positive surplus from the platform when all others join by assuming $b(N+S) > u(N)$, which also implies $b(N+S) > u(N+S-1)$ since $S > 1$. To streamline the presentation of our results, we also assume $u(.)$ is strictly decreasing.

In the following proposition, we characterize the outcome for three parameter ranges that correspond to the most interesting configurations, depending on whether the su-

[12]Indeed, $u(N-1) \leq b(N) - \Delta(1)$ since $\Delta(1) \leq \Delta(N)$.

perstar is pivotal, and whether regular agents are pivotal.

Proposition 6. *Suppose there are N regular agents and one superstar agent, equivalent to $S > 1$ regular agents.*

1. *If* $\min\{b(N), b(N+S-1)\} > u(0)$*, the platform profitably attracts all agents, none is pivotal, and the order of the platform's offers is irrelevant.*
2. *If* $b(N+S-1) > u(0) > u(N-1) > b(N)$*, the platform profitably attracts all agents, only the superstar is pivotal, and the platform optimally offers to the superstar last.*
3. *If* $b(S+N-1) < u(S+N-2)$ *and* $b(N+S) > \frac{u(0)+\sum_{k=0}^{N-1} u(S+k)}{N+1}$*, the platform profitably attracts all agents, all are pivotal, and the platform optimally offers to the superstar first.*

In cases (1) and (2), there exists a unique equilibrium where the platform sets $P^s = b(N+S) - u(N)$ *for the superstar and* $P^k = b(N+S) - u(N+S-1)$ *for regular agent* $k \in \{1, ..., N\}$*. All agents would be strictly better off without the platform. In case (3), there exists a unique equilibrium where the platform sets* $P^s = b(N+S) - u(0)$ *for the superstar and* $P^k = b(N+S) - u(S+k-1)$ *for regular agent* $k \in \{1, ..., N\}$*. The regular agents would be strictly better off without the platform and the superstar is indifferent.*

When no agents are pivotal, each knows the platform can still attract all others regardless of their own decision. As a result, the order of offers doesn't matter and the platform can attain its maximum profit by leaving each agent with their lowest possible outside option — the one they would receive if they were the only agent to not join. This is a straightforward extension of the logic from the baseline model.

The platform can obtain the same maximum profit even if the superstar alone is pivotal, by approaching the superstar last, which removes its pivotal role. Doing so allows the platform to first extract maximum surplus from the regular (non-pivotal) agents, then charge the superstar the highest feasible price. In contrast, when all agents are pivotal, the platform can no longer achieve the same maximum profit. In this case it is optimal to attract the superstar first in order to maximize the decrease in the outside option for the remaining agents, which then allows the platform to extract a higher price from them. In all three cases, the superstar pays less than the regular agents, reflecting its higher outside option.

The more general logic uncovered by Proposition 6 is that the platform should first attract non-pivotal agents, as this minimizes the discounting of prices needed to attract pivotal agents. After attracting non-pivotal agents, the platform should first attract those pivotal agents whose participation generates the greatest negative externality on others' outside option. An implication is that agents with a constant outside option — i.e., whose outside option is unaffected by others joining — should be attracted first. This applies in particular to agents that have no outside option, who should be attracted before others since their presence on the platform helps attract others agents but the reverse is not true.

6.2 Platform competition

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

The agents get benefits $b_1(n_1, n_2)$ 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 $b_2(n_1, n_2)$, which means platforms create network effects. The outside option is $u(n_1, n_2)$, strictly decreasing in both its arguments, reflecting that more participation on either platform reduces the value of staying out.

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

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

$$b_1(2, 0) > u(0, 0) \quad (2)$$

$$u(0, 1) \geq u(1, 0). \quad (3)$$

Assumption (2) 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 (3) 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 two assumptions ensure that platform 1 always attracts both agents.

The following proposition focuses on the most salient cases.

Proposition 7. *Suppose $N = 2$ and conditions (2)-(3) hold. Platform 1 always attracts both agents and agent 1 is always weakly better off than agent 2.*

1. *If $b_2(1,1) > u(0,0)$, both agents are made better off by the presence of 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 made worse off by the presence of 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 made better off by the presence of the platforms and agent 2 is made worse off.*

In case 1, platform 2 provides more value when agents split between platforms than the outside option does when no agents join. In this case, platform 2 exerts sufficient competitive pressure on platform 1 that both agents are better off with the platforms present. 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 provide sufficient competitive pressure on platform 1, so both agents are worse off due to the presence of the platforms. Finally, if platform 2 exerts sufficient competitive pressure on platform 1 for the first agent, but not enough once the first agent joins platform 1 or stays out, then 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 providing the full characterization becomes very messy. Nevertheless, the intuition for our main result is the same: for at least one agent to be better off with competing platforms, the less desired platform needs to be sufficiently competitive relative to the initial outside option. Note that the results in the proposition continue to hold when the agent 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).

If the two platforms are sufficiently close to symmetric, i.e. if $b_2(n_1, n_2) \to b_1(n_2, n_1)$ for every combination of $n_1 \in \{0, 1, 2\}$ and $n_2 \in \{0, 1, 2\}$, then agent 1 is always weakly better off with the platforms. This is not surprising: if the two platforms are close competitors, then competition for the first agent is most intense, so that agent will be made better off. However, it is still possible for agent 2 to be worse off with approximately symmetric platforms, because 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 or not agent 2 is made better off depends on how effective platform 2 is as a competitor once platform 1 has attracted agent 1. If the maximum value platform 2 can offer agent 2 in that case is low enough, platform 1 can extract so much surplus from agent 2 so as to make it worse off than if the platforms didn't exist.

So far we have assumed that agents can only join one of the two competing platforms. If agents can multihome but the benefit functions are independent — i.e., $b_1(n_1, n_2) = b_1(n_1)$ and $b_2(n_1, n_2) = b_2(n_2)$ — so the agent payoff from multihoming is $b_1(n_1) + b_2(n_2)$, while payoff from the outside option is $u(n_1, n_2) = u(n_1) + u(n_2)$, then we are back to the analysis in the baseline model. There is no real competition between platforms in this case. For example, such a case would arise if agents are sellers on two marketplaces, and the marketplaces serve completely different buyer segments, each of which decides whether to go to the outside option or the particular marketplace of interest.

7 Conclusion

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 dynamic model demonstrates how a platform trap can emerge despite agents being rational and forward-looking. We also determined the optimal order for a platform to attract different (heterogeneous) agents in order to best engineer a platform trap.

While we do not advocate drastic policy interventions (such as banning platforms or imposing broad behavioral restrictions on them) just because a platform trap may be theoretically possible, we do think platform traps merit attention when evaluating the impact of platforms and policies targeted at them. Specifically, evaluating the real impact of a platform on surplus should recognize that the value of the alternative to the platform may be considerably higher if the platform didn't exist or were less widely adopted. Simply comparing *current* payoffs with the *current* outside option can be misleading.

Rather than banning a platform or imposing broad behavioral restrictions (e.g., restrictions on raising prices over time or across agents), an alternative way policymakers might reduce the adverse effects of a platform trap is to allow agents to coordinate. If the agents in question are competing sellers, such coordination would typically be considered collusion under competition law. Our results imply there could be value in making an exception for sellers negotiating with a dominant marketplace or publishers negotiating with Facebook and Google, provided such coordination is only in their role as platform customers. A challenge is that some future participants may not be present during any such negotiations, potentially leaving them disadvantaged or even excluded.

Another policy intervention would be to limit platform actions that undermine agents' outside options. For instance, some food delivery platforms (like DoorDash) withhold customer data (e.g., customer details and demographics tied to order histories, customer contact information, granular customer feedback) from participating restaurants when customers order through the platform. This makes it harder for restaurants to serve these customers well via direct channels, thereby pushing more consumers to use the platform, thereby undermining all restaurants' outside option.

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., 2024). 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.

Finally, a multi-period environment in which agents' decisions are linked across periods could be explored. Without introducing switching costs directly, one could model this as agents joining simultaneously each period, producing multiple participation equilibria each period. Over time, equilibrium selection may be influenced by participants' past choices, following the ideas in Halaburda and Yehezkel (2019) and Halaburda et al. (2020).

8 Appendix

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

8.1 Proof of Proposition 1

Denote by $\Gamma(n, b(.), u(.))$ the game from the baseline model with $N = n$ agents and payoff functions $b(.)$ and $u(.)$. We always assume these functions satisfy the assumptions from the baseline model, i.e. $u(.)$ is weakly decreasing and $\Delta(k) =$

$b(k) - u(k-1)$ is weakly increasing in k for $k = 1, .., n$.

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

Lemma 1: *If the platform can profitably attract all agents in the game* $\Gamma(n, b(.), u(.))$*, then it profitably attracts all agents in the game* $\Gamma(n+1, b(.), u(.))$ *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 $\Gamma(1, b(.), u(.))$, then $b(1) > u(0)$, i.e. $\Delta(1) > 0$. Thus, in the game $\Gamma(2, b(.), u(.))$, 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 $\Gamma(2, b(.), u(.))$ by charging each $\Delta(2) > 0$.

Suppose the statement in the lemma is true for $N = n \geq 1$ and *any* payoff functions $b(.)$ and $u(.)$ such that $u(.)$ is weakly decreasing and $\Delta(k) = b(k) - u(k-1)$ is weakly increasing. We show it is also true for $N = n+1$. Namely, suppose the platform profitably attracts all agents in the game $\Gamma(n+1, b(.), u(.))$, so $\Delta(n+1) > 0$. We want to show the platform profitably attracts all agents in the game $\Gamma(n+2, b(.), u(.))$ at prices $P^k = \Delta(n+2)$ for $k = 1, .., n+2$.

If the first of $n+2$ agents doesn't join, when facing the second agent, the platform is in the same position as at the start of $\Gamma(n+1, b(.), u(.))$, so it still attracts the last $n+1$ out of $n+2$ agents. If the first agent joins and the second does not, then when facing agents $\{3, .., n+2\}$, the platform is in the same position as when facing the second agent in $\Gamma(n+1, b(.), u(.))$ after the first agent has joined. By assumption, the platform attracts all agents in this situation, so the platform profitably attracts all agents $\{3, .., n+2\}$ in $\Gamma(n+2, b(.), u(.))$ after the first agent joins and the second does not. This is equivalent to the platform profitably attracting all agents in $\Gamma\left(n, \widetilde{b}(.), \widetilde{u}(.)\right)$, where $\widetilde{b}(k) \equiv b(k+1)$ and $\widetilde{u}(k) \equiv u(k+1)$ (note $\widetilde{u}(.)$ is weakly decreasing and $\widetilde{\Delta}(k) = \widetilde{b}(k) - \widetilde{u}(k-1) = b(k+1) - u(k)$ is weakly increasing). The induction hypothesis ($N = n$) then implies the platform profitably attracts all agents in $\Gamma\left(n+1, \widetilde{b}(.), \widetilde{u}(.)\right)$. And this is equivalent to saying that in $\Gamma(n+2, b(.), u(.))$, after the first agent joins, the platform profitably attracts all remaining $n+1$ agents.

Thus, the platform profitably attracts the last $n+1$ agents in $\Gamma(n+2, b(.), u(.))$ regardless of whether the first agent joins or not, so it can attract the first agent with $P^1 = \Delta(n+2) \geq \Delta(n+1) > 0$. Using a similar logic, the platform attracts all agents in $\Gamma(n+2, b(.), u(.))$ with the same price $\Delta(n+2)$, which maximizes profits. ■

Lemma 2: *In the game* $\Gamma(n, b(.), u(.))$*, 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 $\Gamma(n, b(.), u(.))$. 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 doesn't 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. This means the platform profitably attracts all agents in game $\Gamma\left(k_0, \widetilde{b}(.), \widetilde{u}(.)\right)$, where $\widetilde{b}(k) = b(n_0 - k_0 + k)$ and $\widetilde{u}(k) = b(n_0 - k_0 + k)$. Applying Lemma 1, this implies the platform profitably attracts all agents in game $\Gamma\left(k_0 + 1, \widetilde{b}(.), \widetilde{u}(.)\right)$ with prices equal to $\widetilde{\Delta}(k_0 + 1) = \Delta(n_0 + 1) > 0$. In other words, the platform can profitably attract the last $k_0 + 1$ agents with prices equal to $\Delta(n_0 + 1)$, strictly increasing profits. So it couldn't have been optimal to only attract n_0 agents to join.■

The proof of Proposition 1 also proceeds by induction over games with an increasing number of agents. Start with $\Gamma(N = 2, b(.), u(.))$. 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 $b(.)$ and $u(.)$ such that $u(.)$ 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 $\Gamma(n-1, b(.), u(.))$, then it profitably attracts all agents in the game $\Gamma(n, b(.), u(.))$ with optimal prices $P^k = \Delta(n)$.

- If the platform cannot profitably attract agents in the game $\Gamma(n-1, b(.), u(.))$, there exists $k_0 \in \{1, .., n-1\}$ such that the platform's profit-maximizing prices that induce agents to join in the game $\Gamma(n, b(.), u(.))$ are

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

If $\sum_{k=1}^{n} P^k > 0$, then these are the platform's optimal prices in $\Gamma(n, b(.), u(.))$

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 $\Gamma(n, b(.), u(.))$, then Lemma 1 directly implies it profitably attracts all agents in the game $\Gamma(n+1, b(.), u(.))$ with optimal prices $P^k = \Delta(n)$. Suppose instead the platform cannot profitably attract any agents in the game $\Gamma(n, b(.), u(.))$. In this case, the highest price at which the platform can attract the first agent in $\Gamma(n+1, b(.), u(.))$ 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 $\Gamma\left(n, \widetilde{b}(.), \widetilde{u}(.)\right)$, where $\widetilde{b}(k) = b(k+1)$ and $\widetilde{u}(k) = u(k+1)$. We can then apply the induction hypothesis to conclude:

- If the platform can attract all agents in $\Gamma\left(n-1, \widetilde{b}(.), \widetilde{u}(.)\right)$, then it attracts the last n agents in $\Gamma(n+1, b(.), u(.))$ after the first agent has joined with optimal prices
$$P^k = \widetilde{b}(n) - \widetilde{u}(n-1) = \Delta(n+1) > 0$$
for all $k \in \{2, ..., n+1\}$. So we have
$$P^k = \begin{cases} b(n+1) - u(0) & \text{if} \quad k = 1 \\ b(n+1) - u(n) & \text{if} \quad 2 \leq k \leq n+1 \end{cases}.$$

- If the platform cannot profitably attract agents in $\Gamma\left(n-1, \widetilde{b}(.), \widetilde{u}(.)\right)$, there exists $k_0 \in \{2, .., n\}$ such that in $\Gamma(n+1, b(.), u(.))$ after the first agent has joined, the platform's profit-maximizing prices for agents $k \in \{2, ..., n+1\}$ that induce them to join are
$$P^k = \begin{cases} \widetilde{b}(n) - \widetilde{u}(k-2) = b(n+1) - u(k-1) & \text{if} \quad 2 \leq k \leq k_0 \\ \widetilde{b}(n) - \widetilde{u}(n-1) = b(n+1) - u(n) & \text{if} \quad k_0 + 1 \leq k \leq n+1 \end{cases}.$$
If $\sum_{k=2}^{n+1} P^k > 0$, these prices are optimal for agents $k \in \{2, ..., n+1\}$. If instead $\sum_{k=2}^{n+1} P^k \leq 0$, the platform cannot profitably attract agents $k \in \{2, ..., n+1\}$ even after the first agent joins, so the platform sets any non-negative prices and attracts no agents. Thus, the platform's profit-maximizing prices in $\Gamma(n+1, b(.), u(.))$ that induce all agents to join (recall Lemma 2) are
$$P^k = \begin{cases} b(n+1) - u(k-1) & \text{if} \quad 1 \leq k \leq k_0 \\ b(n+1) - u(n) & \text{if} \quad k_0 + 1 \leq k \leq n+1 \end{cases}.$$

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 $\Gamma(n+1, b(.), u(.))$.

Thus, the induction hypothesis holds for $N = n + 1$. By induction, it holds for all $N \geq 2$. Moreover, if (1) holds, then the platform optimally attracts all agents in $\Gamma(N, b(.), u(.))$ 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 $\Gamma(N, b(.), u(.))$ 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 lower than $u(0)$ (strictly for $k = N$), which is the payoff each agent would obtain without the platform.

8.2 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 $k = 2, .., N$ at prices $P^k = b(N) - u(N-1)$.

8.3 Proof of Corollary 2

If $b(N-1) - u(N-2) \leq 0$, then $b(k) - u(k-1) \leq 0$ for all $k \in \{1, .., N-1\}$, i.e., all agents except the last one 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 all $k = (1, ..., N)$ provided (1) holds. Otherwise, the platform finds it optimal to attract no agents.

8.4 Proof of Corollary 3

From Proposition 1, agent k's net surplus in the game $\Gamma(N, b(.), u(.))$ is either $u(k)$ 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 $\Gamma(N, b(.), u(.))$, its profits are at most $N\Delta(N)$ and repeated application of Lemma 1 implies that platform profits in $\Gamma(N', b(.), u(.))$ are $N'\Delta(N') > N\Delta(N)$.

- If the platform finds it optimal to attract no agents in $\Gamma(N, b(.), u(.))$, then its profits in $\Gamma(N', b(.), u(.))$ are weakly higher (strictly when it can profitably attract all agents in $\Gamma(N', b(.), u(.))$).

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

8.5 Proof of Proposition 2

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

First, we prove that the only possible equilibria in $\Gamma(n, b(.), u(.), P)$ 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 they deviate 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 they deviate 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 $b(.)$ is weakly increasing, all agents join in $\Gamma(N, b(.), u(.), P)$ if $P \leq b(N) - u(0)$ and no agents join otherwise. In $\Gamma(1, b(.), u(.), P)$, the agent joins if and only if the platform sets $P \leq b(1) - u(0)$. Suppose the result holds for $N = n$ and any payoff functions $b(.)$ and $u(.)$ such that $b(.)$ is weakly increasing and $u(.)$ is weakly decreasing (so $\Delta(k) = b(k) - u(k-1)$ is weakly increasing in k). Consider $\Gamma(N = n+1, b(.), u(.), P)$. If the platform sets $P \leq b(n+1) - u(0)$, then after the first agent joins, the game is equivalent to $\Gamma\left(n, \widetilde{b}(.), \widetilde{u}(.), P\right)$, where $\widetilde{b}(n) = b(n+1)$, $\widetilde{u}(n) = u(n+1)$, and

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

The induction hypothesis implies all n agents join in this game. Thus, if the first agent joins in $\widetilde{\Gamma}(n+1, b(.), u(.), P)$, their payoff will be $b(n+1) - P$, whereas if they don't join, their 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 $\Gamma(n+1, b(.), u(.), P)$, 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 $\Gamma(n, b(.), u(.), P)$ 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 $\Gamma(n+1, b(.), u(.), P)$.

By induction, the result holds for any $N \geq 2$. And we can then conclude that when $b(.)$ 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 $b(.)$ is single-peaked and let m_0 be the lowest point at which $b(.)$ reaches its maximum, i.e.

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

We prove that in the game $\Gamma(N, b(.), u(.), P)$, 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 $\Gamma(m_0, b(.), u(.), P)$ (no prior agents have joined), where $b(.)$ 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 $\Gamma(N, b(.), u(.), P)$ 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 $P = b(m_0) - u(0)$ and all agents join.

8.6 Proof of Proposition 4

Suppose the platform offers the pricing function $P(j) = \Delta(j)$ for all $j = 1, .., N$. The last agent will pay $\Delta(N'+1)$ if they join and N' prior agents have joined, so they join (they are indifferent between joining and not). Knowing the last agent will join no matter what, the second-to-last agent will pay $\Delta(N'+2)$ if they join and N' prior agents have joined, which yields the same payoff $u(N'+1)$ as not joining. So the second-to-last agent joins too. Thus, by backwards induction, all agents join and along the equilibrium path they are all charged $\Delta(N)$. This is the maximum price that the platform can charge any agent while getting all agents to join.

8.7 Proof of Proposition 5

In the best possible equilibrium for the platform, the platform prices at $P = \Delta(N)$ to all agents, and each agent joins because they believe every other agent faces the same price and also joins. 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 $\Delta(1) > 0$. The worst equilibrium for the platform is supported by the following strategies and beliefs:

- The platform charges $P^k = \Delta(1)$ to every agent k in the sequence and all agents choose to join
- Provided $P^k \leq \Delta(1)$, each agent believes every other agent is charged $P^k = \Delta(1)$ and they will join, so they also join (indeed, joining yields $b(N) - P \geq b(N) - \Delta(1) \geq u(N-1)$ because $\Delta(.)$ is weakly increasing)
- If an agent is charged a price $P > \Delta(1)$, they believe every other agent is also charged the same P and does not join because they believe no other agent will join either.

Given the agents' strategy, the platform's pricing is optimal to induce participation whenever it is profitable. Conversely, each agent's strategy is optimal given their beliefs, which are fulfilled along the equilibrium path. The resulting platform profit, $N\Delta(1)$, is the minimum profit the platform can obtain given $\Delta(1) > 0$. To see this suppose there is an even worse equilibrium in which the platform sets the price $P^k < \Delta(1)$ to at least one agent. The platform can profitably increase its price P^k to $\Delta(1)$ for any such agent. Then regardless of the number n of other agents this agent expects to join after the price change, it will want to join. Indeed, joining yields payoff $b(n+1) - \Delta(1)$, not joining yields $u(n)$ (recall our PBE refinement implies n does not depend on the agent's own decision), and $b(n+1) - u(n) = \Delta(n+1) \geq \Delta(1)$. Similarly, the platform could change the price of any agent that doesn't join in the proposed equilibrium to $\Delta(1)$ and also profitably induce them to join. Thus, the platform can profitably deviate, so we can rule out any such worse equilibrium.

If $\Delta(1) \leq 0$, the platform cannot profitably attract any agents in the worst equilibrium defined above. Moreover, if it tries to charge a price above $\Delta(1)$ to any agent, then following the beliefs in the worst equilibrium above, such an agent will not join.

References

Baye, M. R. and Morgan, J. (2001). "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets," *American Economic Review*, 91, 454-474.

Biglaiser, G., J. Crémer, 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. Jimenez, and C. Roth (2024) "When Product Markets Become Collective Traps: The Case of Social Media," Working Paper 31771.

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.

Galeotti, A. and Moraga-González, J. L. (2009) "Platform Intermediation in a Market for Differentiated Products," *European Economic Review*, 53(4), 417-428.

Gomes, R. and A. Mantovani (2025) "Regulating Platform Fees under Price Parity," *Journal of the European Economic Association*, 23(1), 190-235.

Hagiu, A. and J. Wright (2024) "Optimal Discoverability on Platforms," *Management Science*, 70(11), 7770–7790.

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.

Jullien, B., A. Pavan and M. Rysman (2021) "Two-sided markets, pricing, and network effects," in In K. Ho (Ed.), *Handbook of Industrial Organization* (Vol. 4, pp. 485-592). Elsevier B.V.

Karle, H., M. Preuss, and M. Reisinger (2024) "Selling on 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.

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.

Wang, C. and Wright, J. (2020) "Search Platforms: Showrooming and Price Parity Clauses," *RAND Journal of Economics*, Vol. 51(1), 32-58.

Wang, C. and Wright, J. (2025) "Regulating Platform Fees," *Journal of the European Economic Association*, Vol. 23(2), 746-783.

Online Appendix: Platform Traps

Andrei Hagiu[1] and Julian Wright[2]

A Alternative timing

Consider first the case with finite $T > 1$ periods, possibly with a different number of agents N_t deciding in each period t, and possibly with discounting of payoffs between periods. The game in each period corresponds exactly to that specified for a single period in the baseline model. While some agents may expect to be active in multiple periods, they will make their participation decision in each period separately from their decision in any other period. This is because there is no linkage across the periods — an agent deciding to participate (or not) in period t has no implications for its payoffs or options in any future period $t' > t$. As shown in Corollary 3 in the main text, if N_t is increasing in t, then the prices charged to agents increase over time (strictly when $\Delta(N)$ is strictly increasing) and there are fewer pivotal agents (once N_t exceeds a certain threshold, there will no longer be any pivotal agents).

In case $T = \infty$, backwards induction no longer pins down a unique equilibrium, and the infinite horizon creates the potential for long-run punishments/rewards between agents that sustain cooperation (e.g., not participating on the platform). However, the finite-period equilibrium will remain an equilibrium in this case, and it would be selected by a Markov Perfect equilibrium refinement.

An alternative extension of our single period sequential model would be to add additional stages, in which agents can reverse their earlier decisions (exit if they have joined, or join if they passed previously). The simplest version is to just add another N stages, where each agent can make another decision. Provided the platform can adjust its price to each agent in these additional N stages, then the analysis is the same as in the second period of the multi-period version discussed above with $T = 2$ (which is also the same as our baseline model). The only difference is that payoffs are only realized once, at the end of the two periods. However, this doesn't affect the analysis for the last N stages of the game because the choices of the agents in the first N stages are irrelevant to their subsequent decisions in the last N stages. Indeed, the decisions made in the first N stages do not affect the options available to each agent in the last N stages and the payoffs that will be obtained at the end. Starting in stage $N + 1$, the platform and the N agents would solve the remaining N-stage game exactly as in the baseline analysis in the paper, ignoring the payoff-irrelevant choices in the earlier N stages.

[1] Boston University
[2] Department of Economics, National University of Singapore

B Commitment to prices

In this section, we analyze the scenario in which the platform commits upfront to the prices that will be charged to the N agents, who still arrive sequentially.

Proposition 8. *If $b(.)$ is weakly increasing, then the platform's profit-maximizing prices are $P_k = b(N) - u(k-1)$ for $k = 1, .., N$, and it profitably attracts all agents iff*

$$Nb(N) - \sum_{k=0}^{N-1} u(k) > 0.$$

If $b(.)$ is single-peaked such that $\max_{k \leq N} \{b(k)\} = b(n_0)$, with $1 \leq n_0 \leq N$, then the platform's profit-maximizing prices are

$$P_k = \begin{cases} b(N+1-k) - u(N-k) & \text{if} \quad 1 \leq k \leq N - n_0 \\ b(n_0) - u(k - (N - n_0 + 1)) & \text{if} \quad N - n_0 + 1 \leq k \leq N \end{cases},$$

and it profitably attracts all agents iff

$$\sum_{k=0}^{N-n_0-1} b(N-k) + n_0 b(n_0) - \sum_{k=0}^{N-1} u(k) > 0.$$

Proof: We start by proving Lemma 2 for the case when the platform commits to its prices, namely that it is optimal to either attract all agents or none of them. Suppose to the contrary, the platform finds it optimal to only attract $0 < n < N$ agents. In this case, the highest price that the platform can charge to any participating agent is $\Delta(n)$ and so we must have $\Delta(n) > 0$, otherwise the platform would not find it optimal to attract n agents. Now consider any agent that does not join at its optimal prices (it must exist) and suppose it is the k-th agent. If the platform keeps all prices unchanged except the price to agent k, which is changed to $P_k = \Delta(n+1) \geq \Delta(n) > 0$, then agent k will participate and all agents that were joining previously, when agent k did not join, will continue to join. So the platform has strictly increased profits, which means it couldn't have been optimal to only induce $n < N$ agents to participate.

We now proceed by induction for $N \geq 2$. Start with $N = 2$. If the platform finds it optimal to attract any agents, it must attract both. And there are two possibilities for the platform's profit-maximizing prices that attract both agents If the first agent is pivotal given the platform's prices, then we must have $P_1 = b(2) - u(0)$. This implies the profit-maximizing price for the second agent is $P_2 = b(2) - u(1)$. If on the other hand the first agent is not pivotal given the platform's prices, then we must have $P_1 = b(2) - u(1)$. In this case, the second agent must join even if the first agent does not join, so the profit-maximizing price for the second agent is $P_2 = b(1) - u(0)$.

Thus, if $b(2) \geq b(1)$, then the platform's profit-maximizing prices are $P_1 = b(2) - u(0)$ and $P_2 = b(2) - u(1)$, resulting in total profits $2b(2) - u(0) - u(1)$. If on the other hand $b(2) < b(1)$, then the platform's profit-maximizing prices are $P_1 = b(2) - u(1)$ and $P_2 = b(1) - u(0)$, resulting in total profits $b(1) + b(2) - u(0) - u(1)$. So the result holds for $N = 2$.

Now suppose the result holds for $N = n \geq 2$, and suppose $N = n+1$ and $b(.)$ is single-peaked such that $\max_{1 \leq k \leq n+1} \{b(k)\} = b(n_0)$, with $1 \leq n_0 \leq n+1$. There are two possibilities for the platform's profit-maximizing prices that attract all agents. If the first agent is pivotal given these prices, then we must have $P_1 = b(n+1) - u(0)$. This ensures the first agent joins, so we can apply the induction hypothesis to the remaining n agents and their prices $P_2, .., P_{n+1}$, with payoff functions $\widetilde{b}(k) = b(k+1)$ and $\widetilde{u}(k) = u(k+1)$, so that $\widetilde{b}(.)$ is single-peaked and

$$\max_{1 \leq k \leq n} \left\{\widetilde{b}(k)\right\} = \max_{2 \leq k \leq n+1} \{b(k)\} = b(\max\{n_0, 2\}) = \widetilde{b}(\max\{n_0, 2\} - 1),$$

where we have used that if $n_0 = 1$, then $b(n)$ must be decreasing, and so $\max_{2 \leq k \leq n+1} \{b(k)\} = b(2)$. We therefore must have

$$\begin{aligned} P_k &= \begin{cases} \widetilde{b}(n+2-k) - \widetilde{u}(n+1-k) & \text{if} \quad 2 \leq k \leq n+2-\max\{n_0,2\} \\ \widetilde{b}(\max\{n_0,2\}-1) - \widetilde{u}(k-(n+3-\max\{n_0,2\})) & \text{if} \quad n+3-\max\{n_0,2\} \leq k \leq n+1 \end{cases} \\ &= \begin{cases} b(n+3-k) - u(n+2-k) & \text{if} \quad 2 \leq k \leq n+2-\max\{n_0,2\} \\ b(\max\{n_0,2\}) - u(k-(n+2-\max\{n_0,2\})) & \text{if} \quad n+3-\max\{n_0,2\} \leq k \leq n+1 \end{cases}. \end{aligned} \tag{4}$$

So in this case total profits are

$$\begin{aligned} & b(n+1) + \sum_{k=2}^{n+2-\max\{n_0,2\}} b(n+3-k) + (\max\{n_0,2\}-1)\, b(\max\{n_0,2\}) - \sum_{k=0}^{n} u(k) \\ = \;& b(n+1) + \sum_{k=\max\{n_0,2\}+1}^{n+1} b(k) + (\max\{n_0,2\}-1)\, b(\max\{n_0,2\}) - \sum_{k=0}^{n} u(k). \end{aligned} \tag{5}$$

Now suppose the first agent is not pivotal given the platform's profit-maximizing prices. In this case, we must have $P_1 = b(n+1) - u(n)$. And the remaining n agents must join given the platform's prices even if the first agent does not join. Thus, applying the induction hypothesis to agents $k = 2, .., n+1$ and noting that $\max_{1 \leq k \leq n} \{b(k)\} = b(\min\{n_0, n\})$ (recall $b(.)$ is single-peaked), we must have

$$P_k = \begin{cases} b(n+2-k) - u(n+1-k) & \text{if} \quad 2 \leq k \leq n+1-\min\{n_0,n\} \\ b(\min\{n_0,n\}) - u(k-(n-\min\{n_0,n\}+2)) & \text{if} \quad n-\min\{n_0,n\}+2 \leq k \leq n+1 \end{cases}.$$

So in this case total profits are

$$
\begin{aligned}
& b(n+1) + \sum_{k=2}^{n+1-\min\{n_0,n\}} b(n+2-k) + \min\{n_0,n\}\, b(\min\{n_0,n\}) - \sum_{k=0}^{n} u(k) \\
= & \sum_{k=\min\{n_0,n\}+1}^{n+1} b(k) + \min\{n_0,n\}\, b(\min\{n_0,n\}) - \sum_{k=0}^{n} u(k).
\end{aligned}
$$

Suppose $b(.)$ is weakly increasing, which means $n_0 = n+1$. If the first agent is pivotal at the platform's profit-maximizing prices, we have from the first case above

$$P_k = b(n+1) - u(k-1)$$

for $1 \leq k \leq n+1$ and total profits

$$(n+1)\, b(n+1) - \sum_{k=0}^{n} u(k).$$

If the first agent is not pivotal at the platform's profit-maximizing prices, we have from the second case above

$$
\begin{aligned}
P_1 &= b(n+1) - u(n) \\
P_k &= b(n) - u(k-2) \text{ for } 2 \leq k \leq n+1
\end{aligned}
$$

and total profits

$$b(n+1) + nb(n) - \sum_{k=0}^{n} u(k).$$

Since $b(.)$ is increasing, we have

$$b(n+1) + nb(n) - \sum_{k=0}^{n} u(k) \leq (n+1)\, b(n+1) - \sum_{k=0}^{n} u(k),$$

which implies the profit-maximizing prices must be $P_k = b(n+1) - u(k-1)$ for $1 \leq k \leq n+1$, so every agent (including the first one) is pivotal.

Now suppose $b(.)$ is single-peaked and $\max_{1 \leq k \leq n+1} \{b(k)\} = b(n_0)$, with $1 \leq n_0 < n+1$. In this case, if the first agent is pivotal at the platform's profit-maximizing prices, then profits are given by expression (5) above. If the first agent is not pivotal at the platform's profit maximizing prices, then we have from the second case above

$$P_k = \begin{cases} b(n+2-k) - u(n+1-k) & \text{if} \quad 1 \leq k \leq n+1-n_0 \\ b(n_0) - u(k-(n-n_0+2)) & \text{if} \quad n-n_0+2 \leq k \leq n+1 \end{cases} . \tag{6}$$

and total profits are

$$\sum_{k=n_0+1}^{n+1} b(k) + n_0 b(n_0) - \sum_{k=0}^{n} u(k).$$

Comparing this expression with (5), we have

$$\begin{aligned} & \sum_{k=n_0+1}^{n+1} b(k) + n_0 b(n_0) - \sum_{k=0}^{n} u(k) \\ \geq \quad & b(n+1) + \sum_{k=\max\{n_0,2\}+1}^{n+1} b(k) + (\max\{n_0,2\}-1)\, b(\max\{n_0,2\}) - \sum_{k=0}^{n} u(k) \end{aligned} \tag{7}$$

for any n_0 such that $1 \leq n_0 \leq n$. Indeed, if $n_0 = 1$, so $b(.)$ is decreasing, then the inequality (7) reduces to

$$b(1) \geq b(n+1),$$

which is clearly true. If $2 \leq n_0 \leq n$, then the inequality (7) reduces to

$$b(n_0) \geq b(n+1),$$

which is also true. Thus, when $b(.)$ is single-peaked such that $\max_{1 \leq k \leq n+1} \{b(k)\} = b(n_0)$, with $n_0 \leq n$, the profit-maxizing platform prices are given by (6) above.

Thus, the induction hypothesis holds for $N = n+1$ when $b(.)$ is weakly increasing or single-peaked, so it holds for all $N \geq 2$.■

A few observations follow from this proposition.

- When $b(.)$ is weakly increasing, all agents are pivotal when the platform commits to prices (Proposition 8), so the platform's profit-maximizing profits are weakly lower than in the case without commitment (Proposition 1 in the baseline), strictly so if not all agents are pivotal in the baseline.

- When $b(.)$ is decreasing, we have

$$P_k = b(n+2-k) - u(n+1-k)$$

for $1 \leq k \leq n+1$, so prices are weakly decreasing in k because $\Delta(.)$ is weakly increasing. So with commitment, the first agent gets charged the highest price. Moreover, profit can now be higher than in the baseline case without commitment, depending on how many agents are pivotal in the baseline. E.g. profit is now strictly higher in the case that all agents are pivotal in the baseline. The logic behind this is that the platform commits to attract later agents (with a lower price) even if earlier agents don't join, so it takes away the pivotal role of earlier agents. The result is prices are decreasing rather than increasing. This strategy is profitable when attracting later agents in case earlier agents don't join is not too costly, which arises when $b(.)$ is decreasing, indeed meaning the platform can extract more in total by committing to these prices.

C Microfounded model

In this section, we provide a fully worked out microfoundation for our general model based on the case of a two-sided marketplace. We consider three versions of this microfoundation. In Section C.1 sellers are independent (monopolists) and the platform's pricing follows our assumptions in the general model. Section C.2 extends the setting with independent sellers to show how a two-sided platform trap can arise (i.e., buyers can also be made worse off by the presence of the platform) — in part, this is due to the fact that the platform charges sellers ad-valorem rather than fixed prices. Finally, Section C.3 allows for competing sellers, so that $b(.)$ is now decreasing rather than increasing, but shows that all the properties of our general model still hold (e.g. that $\Delta(.)$ is weakly increasing).

C.1 Independent sellers

Suppose there are $N \geq 2$ independent sellers (local monopolists). There is a continuum of measure one of buyers. Each buyer wants to buy from one of the independent sellers, with the particular seller a buyer is matched to being randomly drawn with equal probability across all sellers. Buyers have an outside option valued at zero.

We assume initially sellers always sell in the direct channel. For instance, this may be because there is no costs of doing so and there is no way for them to commit not to do so (we will relax this below). In addition, sellers can sell on a platform if they decide to list on the platform. In either channel, sellers face a standard downward sloping demand from each buyer, denoted $q_D(p)$ in the direct channel and $q_P(p)$ on the platform, where p is the seller's price which can be different on each channel. Sellers are assumed to have a constant marginal cost of c.

Each buyer faces a search cost of s of using the direct channel (which allows them to see all

N sellers in the direct channel) and so find their preferred seller for sure, with s drawn from some strictly increasing distribution function $G(s)$ with full support over $[0, S]$, where $S > 0$. On the other hand, buyers face no search cost of searching on the platform (which allows them to see all sellers, if any, listed on the platform). However, buyers can only choose one channel.[3]

The platform sets a fee to each seller to participate, denoted P^k for seller k. After each of the sellers has decided whether to participate on the platform or not, buyers observe how many sellers have decided to participate. Knowing their particular draw s, buyers then each decide (simultaneously) which channel to use (the platform or the direct channel). Since buyers do not know if the sellers they observe are a match until after they go to a channel and discover the sellers there, buyers care about how many sellers they can reach on the platform. The more sellers 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.

If n sellers participate on the platform, let $m(n)$ be the measure of buyers that choose to go to the platform, so $1 - m(n)$ is the measure of buyers that choose to go to the direct channel. A seller's per-buyer profit in the direct channel is then $\max_p \{(p - c)\, q_D(p)\}$, where each buyer that goes to the direct channel chooses how many transactions to make based on the price p in the direct channel (i.e., $q_D(p)$). Let $p_D = \arg\max_p \{(p - c)\, q_D(p)\}$. Then $\pi_D = (p_D - c)\, q_D(p_D)$ is a seller's profit per buyer on the direct channel and $\frac{1}{N}(1 - m(n))\, \pi_D$ is a seller's expected profit in the direct channel. Moreover, if the indirect utility function of buyers associated with the demand q_D is denoted $v_D(p)$, a buyer's indirect utility from purchasing from her matched seller in the direct channel is $v_D = v_D(p_D)$.

Similarly, a seller's per-buyer profit from sales on the platform (before the platform's fee) is $\max_p \{(p - c)\, q_P(p)\}$. Let $p_P = \arg\max_p \{(p - c)\, q_P(p)\}$. Then $\pi_P = (p_P - c)\, q_P(p_P)$ is a seller's profit per buyer on the platform and $\frac{1}{N} m(n)\, \pi_P$ is a seller's expected profit from sales on the platform. And given a buyer's indirect utility function $v_P(p)$, a buyer's indirect utility from purchasing from her matched seller on the platform is $v_P = v_P(p_P)$.

We assume $q_P(p) > q_D(p)$, so the platform generates higher demand for sellers from each buyer, reflecting various transactional benefits the platform adds. This implies $\pi_P > \pi_D$ and $v_P > v_D$.

If n sellers participate on the platform, a buyer will go to the platform if their expected utility $\frac{n}{N} v_P$ is greater than the expected utility of going to the direct channel, which is $v_D - s$. Thus, the measure of buyers going to the platform is

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

Clearly, $m(n)$ is weakly increasing in n. Note $m(N) = 1$ and $m(0) < 1$, which implies $m(n)$ is

[3]Equivalently, after searching on one channel, their search costs of searching in another channel becomes prohibitive.

strictly increasing at some point. We impose the weak assumption that

$$S + \frac{N-1}{N} v_P > v_D,$$

which says that if all but one seller is on the platform, the platform will attract some buyers (i.e. $m(N-1) > 0$). This ensures our assumption $u(N-1) < u(0)$ holds.

Now we translate this microfounded model to our general setting. Each of the N sellers is one of the agents. An agent's payoff from participating on the platform when n agents in total participate is $b(n) = \frac{1}{N}(m(n)\pi_P + (1-m(n))\pi_D)$. Since $\pi_P > \pi_D$, b is weakly increasing in n (strictly increasing at some point). An agent's payoff if it does not participate on the platform but n agents do is $u(n) = \frac{1}{N}(1-m(n))\pi_D$, which is weakly decreasing in n (strictly decreasing at some point). Note

$$\Delta(n) = b(n) - u(n-1) = \frac{1}{N}(m(n)(\pi_P - \pi_D) + m(n-1)\pi_D),$$

so given $\pi_P > \pi_D$, in this model $\Delta(1) > 0$ and $\Delta'(n) \geq 0$ for all n, with the inequality strict at some point. Finally,we note that (1) holds given this just requires $\pi_P > \sum_{k=1}^{N} \frac{1}{N}(1-m(k-1))\pi_D$, which is true given $\pi_P > \pi_D$ and $m(k-1) \geq 0$ for all k. Thus, all the conditions are satisfied for Proposition 1 to hold.

Proposition 1 implies all N sellers participate and face the platform fees characterized in the proposition. Since all sellers join, they get a payoff of $\frac{1}{N}m(N)\pi_P + \frac{1}{N}(1-m(N))\pi_D - P^k$. For non-pivotal sellers, this is $\frac{1}{N}(1-m(N-1))\pi_D$ with the platform, which is strictly lower than the payoff of $\frac{1}{N}\pi_D$ sellers would get in the absence of the platform. For a pivotal seller k, this is a payoff of $\frac{1}{N}(1-m(k-1))\pi_D$ with the platform, which is weakly lower than the payoff of $\frac{1}{N}\pi_D$ sellers would get in the absence of the platform, strictly so for seller $k = N$.

C.2 Two-sided platform trap

Up till now, we have focused on the payoffs of the sellers, which we interpreted as the relevant agents (the sellers are the agents that face fees and make strategic decisions). Clearly, buyers are always better off in the above microfounded model due to the existence of the platform. With the platform, buyers have the additional option to go to the platform, which is desirable if they draw a high search cost s for using the direct channel. Moreover, the value of that option increases as more sellers participate on the platform, reflecting that the platform fee does not feed back into buyer surplus (there is no pass-through in this model).

This raises the question of whether it is ever possible that the platform can induce sellers to participate even though both buyers and sellers could end up worse off? In order to consider this possibility, we modify the above model in two ways. First, sellers are assumed to be limited to one channel. If they participate on the platform, they no longer keep their direct channel, perhaps reflecting some coordination cost of trying to manage two different channels. Second, we modify the model so there is pass-through from the fee a platform charges a seller to the seller's price. This provides a mechanism by which buyers could ultimately end up being worse off because of the platform. To do so we assume the platform charges ad-valorem fees (a percentage of a seller's revenue) instead of a fixed fee per seller. When deciding whether to search on the platform or in the direct channel, we continue to assume buyers just observe how many sellers have listed on the platform, but they do not observe the fees the sellers are charged or their prices.

Let the ad-valorem fee charged to a seller be denoted τ, where $0 \leq \tau \leq 1$. Facing such a fee, a seller's per-buyer profit from sales on the platform is $\max_p \{((1-\tau)p - c)\, q_P(p)\}$. Let $p(\tau) = \arg\max_p \{((1-\tau)p - c)\, q_P(p)\}$, which is increasing in τ. Then $\pi_P(\tau) = ((1-\tau)p(\tau) - c)\, q_P(p(\tau))$ is a seller's profit per buyer, which is decreasing in τ and $\frac{1}{N} m(n)\, \pi_P(\tau)$ is a seller's expected profit from sales on the platform. And given a buyer's indirect utility function $v_P(p)$, a buyer's indirect utility from purchasing from her matched seller on the platform is $v_P(\tau) = v_P(p(\tau))$, which is decreasing in τ. Consistent with the case without any ad-valorem fee, we assume $q_P(p) > q_D(p)$, which implies $\pi_P(0) > \pi_D$ and $v_P(0) > v_D$.

When buyers decide which channel to go to, they only observe how many sellers are on each channel. The N sellers will be split between the platform and the direct channel. Buyers do not observe the actual fees charged to different sellers on the platform. Suppose buyers expect a single equilibrium fee τ^*. If n sellers participate on the platform, a buyer will go to the platform if their expected utility $\frac{n}{N} v_P(\tau^*)$ is greater than the expected utility of going to the direct channel, which is $\frac{N-n}{N} v_D - s$. Thus, the measure of buyers going to the platform is

$$m(n) = 1 - G\left(\frac{N-n}{N} v_D - \frac{n}{N} v_P(\tau^*)\right). \qquad (8)$$

Again, $m(n)$ is weakly increasing in n, with $m(N) = 1$ and $m(0) < 1$, implying $m(n)$ is strictly increasing at some point. We modify our previous assumption to this new setting. We assume

$$S + \frac{N-1}{N} v_P(\tau^*) > \frac{1}{N} v_D,$$

so that $m(N-1) > 0$ and therefore $u(N-1) < u(0)$ holds. Furthermore, to focus on the simplest case without any pivotal sellers, suppose $m(1)\, \pi_P(0) > \pi_D$, which says a seller obtains higher expected profit on the platform than the direct channel if it was charged no fee even if it is the only seller.

Consider the last seller deciding and facing the fee τ^N assuming N' other sellers have already joined the platform. If the seller joins, it will get

$$\frac{m\left(N'+1\right)\pi_P\left(\tau^N\right)}{N}$$

while if it does not join, it will get

$$\frac{\left(1-m\left(N'\right)\right)\pi_D}{N}.$$

The seller will participate provided τ^N is no more than the fee τ satisfying

$$\pi_P\left(\tau\right)=\frac{1-m\left(N'\right)}{m\left(N'+1\right)}\pi_D. \tag{9}$$

Suppose $m\left(N'\right)<1$. Since $m\left(1\right)\pi_P\left(0\right)>\pi_D$ we have $m\left(N'+1\right)\pi_P\left(0\right)>\left(1-m\left(N'\right)\right)\pi_D$. We also have $m\left(N'+1\right)\pi_P\left(1\right)=0<\left(1-m\left(N'\right)\right)\pi_D$. So there exists a unique value of $\tau>0$ solving (9), and the platform can always induce this last seller to join (even if no other sellers have joined by setting this fee). Similarly, if $m\left(N'\right)=1$, we have $m\left(N'+1\right)\pi_P\left(0\right)>\left(1-m\left(N'\right)\right)\pi_D$ and $m\left(N'+1\right)\pi_P\left(1\right)=0\leq\left(1-m\left(N'\right)\right)\pi_D$. Again, there exists a unique value of $\tau>0$ solving (9), and the platform can always induce this last seller to join (even if no other sellers have joined by setting this fee). Based on the same logic as in our baseline analysis, the platform can therefore induce all sellers to join by charging appropriate fees to each seller. The equilibrium fee is the unique $\tau>0$ solving (9) when $N'=N-1$, given each seller expects all other sellers to join in equilibrium.

In summary, the equilibrium τ^* is the $\tau>0$ solving (9) when $N'=N-1$ and $m\left(N-1\right)$ and $m\left(N\right)$ are determined by (8).

Sellers end up with expected payoffs $\frac{1}{N}m\left(N\right)\pi_P\left(\tau^*\right)$, which after using the equilibrium characterization of τ^* equals $\frac{1}{N}\left(1-m\left(N-1\right)\right)\pi_D$. This compares to getting $\frac{1}{N}\pi_D$ in the absence of the platform. So clearly all sellers are strictly worse off given $m\left(N-1\right)>0$.

Now consider buyers' expected utility. If the platform didn't exist, they would get v_D-s, with s drawn from distribution G. With the platform they get $v_P\left(\tau^*\right)$. Note that some buyers are strictly worse off with the platform provided $v_P\left(\tau^*\right)<v_D$. Indeed, our assumption that $S+\frac{N-1}{N}v_P\left(\tau^*\right)>\frac{1}{N}v_D$ does not rule out the possibility that all buyers are strictly worse off with the platform, which arises if $S+v_P\left(\tau^*\right)<v_D$. Otherwise, some buyers that use the platform will be worse off and others will be better off compared to the case without the platform, depending on their draw of s.

C.3 Competing sellers

We return to the setting in Section C.1 in which the platform sets fixed prices P^k for each seller k, and sellers do not give up their direct channels when they list on the platform. We make the following changes to the setting. Each buyer wants to buy one unit from one of the sellers, and obtains net utility from buying from seller i in the direct channel which is $v_D^i = v - p_D^i + \beta\varepsilon^i$, where ε^i is drawn i.i.d. from the standard (unit) Gumbel distribution over $(-\infty, \infty)$ across buyers and sellers, p_D^i is seller i's price in the direct channel, and $\beta > 0$ measures the importance of the match value. Suppose buyers randomly sample from two sellers in the direct channel. In contrast, if buyers go to the platform, they can choose from all sellers listed on the platform, and obtain a net utility from the seller i they purchase from of $v_P^i = v - p_P^i + \beta\varepsilon^i + s$, where s is distributed across buyers according to the uniform distribution on $[\underline{s}, \overline{s}]$, and could be positive (a benefit of transacting on the platform) or negative (a cost of doing so). We normalize $\overline{s} - \underline{s} = 1$, with $\overline{s} > 0$. Buyers must purchase from one of the sellers on one of the channels (they have no outside option). The timing follows the setting in Section C.1, with their draw of s being observed before they choose a channel, and buyers' sample of sellers in the direct channel and the draws of their match values in either channel only being observed after choosing a channel.

Provided $n \geq 2$, standard analysis from extreme value theory implies the expected match value of a buyer on the platform that draws s can be written as $v + s + \beta(\gamma + \ln(n))$, while in the direct channel it will be $v + \beta(\gamma + \ln(2))$, where γ is the Euler-Mascheroni constant. The corresponding equilibrium prices[4] are $\frac{\beta n}{n-1}$ and 2β. Note no buyers will go to the platform if there are no sellers listed, since they get no surplus on the platform vs. a positive surplus in the direct channel, i.e. $m(0) = 0$. In case only one seller lists on the platform, given the lack of an outside option, all buyers would have to purchase even facing an arbitrarily high price, so they would never go to the platform knowing this, i.e., $m(1) = 0$. For $n \geq 2$, comparing the expected match values less equilibrium prices across channels, and solving for the share of buyers that draw s high enough to go to the platform, we get

$$m(n) = \overline{s} + \beta\left(2 - \ln(2) + \ln(n) - \frac{n}{n-1}\right).$$

We assume $\overline{s}$, β and N such that $m(N) < 1$, which ensures $m(n)$ is strictly increasing in $n \geq 2$.

Now $u(n) = 2\beta(1 - m(n))$, so $u(1) = u(0) = 2\beta$, $u(n) > 0$ given $m(N) < 1$, and $u(n)$ is strictly decreasing in n for $n \geq 2$, so that $u(N-1) < u(0)$. Also, $b(n) = \frac{n}{n-1}\beta m(n) + 2\beta(1 - m(n))$, so $b(1) = 2\beta$, $b(n) > 0$ given $0 \leq m(N) < 1$, and $b(n)$ is strictly decreasing in n for $n \geq 2$ (given $0 < \frac{n}{n-1} < 2$ and is decreasing towards 1 as n increases). Finally, $\Delta(n) =$

[4] See S. Anderson and A. de Palma (1992) "The Logit as a Model of Product Differentiation," Oxford Economic Papers, 44(1), 51-67.

$b(n) - u(n-1)$, so that $\Delta(1) = \Delta(2) = 0$, and $\Delta(n) = \beta\left(2m(n-1) - \left(\frac{n-2}{n-1}\right)m(n)\right)$.

We want to show that $\Delta(n)$ is increasing for a wide range of parameters $\overline{s}$, β and N such that $m(N) < 1$. After substituting in the expression for $m(n)$, we have

$$\begin{aligned}\frac{\Delta(n+1) - \Delta(n)}{\beta} &= 2(\ln(n) - \ln(n-1)) - \frac{n-1}{n}(\ln(n+1) - \ln(n)) \\ &\quad + \frac{2n^3 - 5n + 2}{n^2(n-1)^2(n-2)} - \frac{1}{n(n-1)}\left(2 - \ln(2) + \ln(n) + \frac{\overline{s}}{\beta}\right).\end{aligned} \tag{10}$$

It can be verified that (10) is always positive for n large enough, and tends to zero as $n \to \infty$ for any fixed value of $\frac{\overline{s}}{\beta}$. However, for small values of n, (10) can become negative if $\frac{\overline{s}}{\beta}$ is too high. A sufficient condition for (10) to be positive for all $n \geq 2$ is $\frac{\overline{s}}{\beta} \leq \frac{61}{9}$. Thus, we can conclude that $\Delta(n)$ is increasing for $n \geq 2$ if $\overline{s}$, β and N are such that $\frac{\overline{s}}{\beta} \leq \frac{61}{9}$ and $m(N) < 1$.

D Full equilibrium characterization for $N = 2$ and $N = 3$

Suppose first $N = 2$. If $\Delta(1) \leq 0$, then the first agent is pivotal, so the highest prices the platform can charge to attract both agents are $P^1 = b(2) - u(0)$ to the first agent and $P^2 = b(2) - u(1)$ to the second agent. This is profitable if and only if (iff) $2b(2) - u(0) - u(1) > 0$. If $\Delta(1) > 0$, then no agent is pivotal, so the platform optimally charges $P^1 = P^2 = \Delta(2) > 0$ to both agents and makes positive profits. To sum up, platform profits are

$$\begin{cases} 0 & \text{if} \quad \max\left\{b(2) - \frac{u(0)+u(1)}{2}, b(1) - u(0)\right\} \leq 0 \\ 2b(2) - u(1) - u(0) & \text{if} \quad b(1) - u(0) \leq 0 < b(2) - \frac{u(0)+u(1)}{2} \\ 2(b(2) - u(1)) & \text{if} \quad b(1) - u(0) > 0 \end{cases}.$$

Now suppose $N = 3$. If the first agent doesn't join, from above, the platform attracts the two remaining agents if

$$b(2) - \frac{u(0) + u(1)}{2} > 0$$

or

$$b(1) - u(0) > 0.$$

If the first agent does join, the platform attracts the two remaining agents if

$$b(3) - \frac{u(1) + u(2)}{2} > 0$$

or

$$b(2) - u(1) > 0.$$

It is easily seen that

$$b(2) - u(1) \geq \max\left\{b(1) - u(0), b(2) - \frac{u(0)+u(1)}{2}\right\}$$

because $\Delta(2) \geq \Delta(1)$ and $u(.)$ is weakly decreasing. Thus, if the platform attracts the last two agents when the first agent does not join, it necessarily attracts the last two agents after the first agent joins. There are therefore three relevant cases:

1. $\max\left\{b(1) - u(0), b(2) - \frac{u(0)+u(1)}{2}\right\} > 0$, in which case the platform attracts the last two agents regardless of whether the first agent has joined or not
2. $\max\left\{b(1) - u(0), b(2) - \frac{u(0)+u(1)}{2}\right\} \leq 0 < \max\left\{b(2) - u(1), b(3) - \frac{u(1)+u(2)}{2}\right\}$, in which case the platform attracts the last two agents iff the first agent has joined
3. $\max\left\{b(2) - u(1), b(3) - \frac{u(1)+u(2)}{2}\right\} \leq 0$, in which case the platform cannot profitably attract any agent.

In the first case, the platform optimally charges $P^1 = P^2 = P^3 = \Delta(3) > 0$ and makes profits $3\Delta(3) > 0$.

In the second case, the first agent is pivotal so the platform optimally charges $P^1 = b(3) - u(0)$ to attract the first agent. Its maximum profits from the last two agents are then (applying the case $N = 2$ and taking into account the first agent has joined)

$$\begin{cases} 2b(3) - u(2) - u(1) & \text{if} \quad b(2) - u(1) \leq 0 < b(3) - \frac{u(1)+u(2)}{2} \\ 2(b(3) - u(2)) & \text{if} \quad b(2) - u(1) > 0 \end{cases},$$

so total profits are

$$\begin{cases} 3b(3) - u(2) - u(1) - u(0) & \text{if} \quad b(2) - u(1) \leq 0 < b(3) - \frac{u(1)+u(2)}{2} \\ 3b(3) - 2u(2) - u(0) & \text{if} \quad b(2) - u(1) > 0 \end{cases}.$$

And in the third case, the platform makes zero profits.

Combining all cases above, we have four final cases:

- if $\max\left\{b(1) - u(0), b(2) - \frac{u(0)+u(1)}{2}\right\} > 0$, the platform attracts all three agents, none are pivotal and the platform's profits are $3\Delta(3)$
- if $\max\left\{b(1) - u(0), b(2) - \frac{u(0)+u(1)}{2}\right\} \leq 0 < b(2) - u(1)$ and $b(3) > \frac{2u(2)+u(0)}{3}$, the platform attracts all three agents, only the first agent is pivotal, and the platform's profits are $3b(3) - 2u(2) - u(0) > 0$

- if $b(2) - u(1) \leq 0 < b(3) - \frac{u(2)+u(1)+u(0)}{3}$, the platform attracts all three agents, the first two agents are pivotal, and the platform's profits are $3b(3) - u(2) - u(1) - u(0) > 0$.
- otherwise the platform attracts no agents.

E Negotiated prices

Consider the negotiation of prices with each agent when the platform no longer has all the bargaining power. Specifically, we assume that the platform engages in Nash bargaining with each agent (sequentially). As before, the platform can set a different price for each agent, and each agent can observe the decisions of past agents. We focus on the case that no agents can be pivotal, so $\Delta(1) > 0$, which ensures the platform can always profitably induce all agents to participate in any equilibrium. This scenario leads to the strongest version of the platform trap, and so it serves as a useful benchmark to show that even in this case, when agents can negotiate, some may no longer suffer from the trap. The assumption also ensures we can obtain closed form solutions for the equilibrium prices despite the recursive nature of the problem.

To get a sense for how this works formally, consider the last agent to decide whether to participate and suppose N' other agents have already decided to participate. When the last agent decides whether or not to participate, it expects to get $b(N'+1) - P^N$ if it participates and $u(N')$ if it doesn't. Meanwhile, the platform obtains an additional profit of P^N if the agent participates and no additional profit if it doesn't. Assuming each agent's bargaining power relative to the platform is measured by the parameter α, the Nash bargaining solution is the price P^N that maximizes

$$\left(\Delta(N'+1) - P^N\right)^{\alpha} \left(P^N\right)^{1-\alpha},$$

which implies

$$P^N(N') = (1-\alpha)\,\Delta(N'+1).$$

Since $\Delta(N'+1) > 0$, the platform will induce the last agent to participate because it can extract a positive price from doing so.

Things are more complicated for agents that decide earlier because whether an agent participates or not affects the platform's negotiation position vis-a-vis all subsequent agents. Nevertheless, we are able to establish the following result (since it is long, we place the proof of the proposition at the end of this section).

Proposition 9. *Suppose $\Delta(1) > 0$ and the platform engages in Nash bargaining with each of the N agents in sequence, with $0 \leq \alpha < 1$ being each agent's bargaining power and $1-\alpha$ being the platform's bargaining power. There exists a unique equilibrium in which the platform induces all*

N agents to participate and the price charged to the k-th agent is

$$P^k = (1-\alpha)\left(\sum_{j=0}^{N-k} C_j^{N-k}(1-\alpha)^j \alpha^{N-k-j}\Delta(k+j)\right) \tag{11}$$

for all $k \in \{1,..,N\}$, where C_j^{N-k} is the binomial coefficient (the number of ways to choose j items from a set of $N-k$ items).

The price P^k is independent of k in the special case of $\alpha = 0$, so the platform sets the same price to all agents in equilibrium. When $\alpha = 0$, the platform has all the bargaining power, and we have $P^k = \Delta(N)$, so this is identical to Corollary 1 in which the platform makes a take-it-or-leave-it offer to each agent. Also, as $\alpha \to 1$, agents have all the bargaining power, so we have $P^k \to 0$ for all k. Whether agents would be better off without the platform in that case just depends on whether the platform is inefficient or not, i.e. whether $b(N) < u(0)$.

The more interesting case arises when $\alpha > 0$. Since $\sum_{j=0}^{N-k} C_j^{N-k}(1-\alpha)^j \alpha^{N-k-j} = 1$ for any k, the platform's optimal price to the k-th agent is $(1-\alpha)$ multiplied by a weighted average of the surplus obtained from joining the platform for each of the subsequent $N-k+1$ agents (including the k-th agent) deciding. That is, it is a weighted average of the $\Delta(n)$ terms, with n ranging from k to N. We wish to determine whether the platform's price is increasing across agents that sign up and whether agents are better or worse off due to the platform.

Since no agent is pivotal, by assumption, an agent expects all other agents to participate regardless of its decision. Thus, an agent's gross surplus $\Delta(N)$ from joining versus not doesn't change with k. Instead, any change in P^k over k happens due to changes in the additional payoff that the platform stands to gain on subsequent agents joining when agent k joins. Specifically, when agent k joins rather than doesn't, the additional profit that the platform stands to gain beyond the current price can be measured by the extent to which its negotiated price is expected to increase over subsequent agents. As k increases, with fewer potential agents left to join, the impact on the platform's profit from not signing an agent is felt over fewer subsequent agents signing up. This suggests the platform has less to lose from not attracting an agent as k increases, leading to P^k being increasing in k. But the additional profit that can be extracted from subsequent agents also depends on $\Delta(.)$. Given $\Delta(.)$ is an increasing function, as k increases, the average of the additional profit that can be extracted from the remaining agents increases, meaning the platform could actually have more to lose. The overall net effect on how much the platform stands to lose by not attracting an agent as k increases is ambiguous, and depends on the particular shape of $\Delta(.)$.

To proceed, we assume linear externalities so $b(n) = b_0 + b_1 n$ with $b_1 > 0$ and $u(n) = u_0 - u_1 n$ with $u_1 > 0$, which implies $\Delta(n) = \beta_0 + \beta_1 n$ where $\beta_0 = b_0 - u_0 - u_1$ and $\beta_1 = b_1 + u_1 > 0$.

Substituting these expressions into (11) implies

$$P^k = (1-\alpha)\left(\beta_0 + \beta_1\left(N - \alpha\left(N-k\right)\right)\right) \tag{12}$$

for all $k \in \{1,..,N\}$. The monotonicity of P^k in k allows us to characterize when the platform trap arises. Specifically, define the additional payoff of agent k due to the existence of the platform as $\Omega = b(N) - P^k - u(0)$, where P^k is defined in (12). The assumptions that $\Delta(1) \geq 0$ and $N > 1$ together with $b(n)$ being strictly increasing imply the platform is efficient, so $b(N) > u(0)$. This implies $\Omega > 0$ if $\alpha = 1$. Moreover, from Corollary 1 we have $\Omega < 0$ if $\alpha = 0$. Now consider $0 < \alpha < 1$. Since Ω is a quadratic function of α, with $\Omega'' < 0$, there exists a unique $\alpha(k)$ on $(0,1)$ such that $\Omega = 0$ for $\alpha = \alpha(k)$, $\Omega < 0$ for all $\alpha < \alpha(k)$ and $\Omega > 0$ for all $\alpha > \alpha(k)$. And given P^k is clearly increasing in k, $\alpha(k)$ must be increasing in k. This establishes the following proposition.

Proposition 10. *Suppose $\Delta(1) > 0$ and the platform engages in Nash bargaining with each of $N \geq 2$ agents in sequence. If $\alpha = 0$, all agents would be strictly better off without the platform. Now consider $0 < \alpha < 1$ and suppose externalities are linear as defined above. The price charged to the k-th agent, which is given in (12), is increasing in k. There exists $0 < \alpha_1 < \alpha_N < 1$ such that if $\alpha < \alpha_1$, all agents would be strictly better off without the platform, if $\alpha > \alpha_N$, all agents would be strictly worse off without the platform, while for $\alpha_1 \leq \alpha \leq \alpha_N$, agents with $k < \overline{k}$ would be strictly worse off without the platform and agents with $k > \overline{k}$ would be strictly better off without the platform, where $\overline{k}$ is defined as the unique level of k that makes $P^k = b(N) - u(0)$ for any given α in this range.*

The proposition shows that under linear externalities, provided the platform's bargaining power is strong enough, all agents would be strictly better off without the platform, while provided the agents' bargaining power is strong enough, the reverse is true. For some intermediate range of α, agents that receive early offers would be strictly worse off without the platform while those receiving later offers would be strictly better off without the platform.

E.1 Proof of Proposition 9

We first show by backwards induction that the price charged to the k-th agent to decide is

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

and all subsequent agents (including the k-th agent) will participate, where N' is the number of agents that have decided to participate before k.

We have already shown in the main text that this holds for the last agent to decide ($k = N$).

Suppose the result holds for all $k \geq K$ for some $K \leq N$. Let us now show it also holds for $k = K - 1$. To determine P^{K-1}, suppose N' other agents have already decided to participate. The $(K-1)$-th agent knows regardless of whether it decides to participate, the next $N - K + 1$ agents will participate (by the induction hypothesis). Thus, the agent's increase in surplus from participating relative to not participating is

$$\Delta\left(N' + N - K + 2\right) - P^{K-1}.$$

If the agent participates, the platform's profit from this and the next $N-K+1$ agents participating is

$$\begin{aligned} & P^{K-1} + \sum_{m=0}^{N-K} P^{K+m} \\ = \;& P^{K-1} + (1-\alpha)\sum_{m=0}^{N-K}\left(\sum_{j=0}^{N-K-m} C_j^{N-K-m}\left(1-\alpha\right)^j \alpha^{N-K-m-j}\Delta\left(N' + m + j + 2\right)\right). \end{aligned}$$

If the $(K-1)$-th agent does not participate, then the platform's profit from the subsequent agents is

$$(1-\alpha)\sum_{m=0}^{N-K}\left(\sum_{j=0}^{N-K-m} C_j^{N-K-m}\left(1-\alpha\right)^j \alpha^{N-K-m-j}\Delta\left(N' + m + 1 + j\right)\right).$$

Thus, the Nash bargaining solution is

$$\begin{aligned} P_t^{K-1} \;=\; & (1-\alpha)\,\Delta\left(N' + N - K + 2\right) \\ & -\alpha\left(1-\alpha\right)\sum_{m=0}^{N-K}\left(\sum_{j=0}^{N-K-m} C_j^{N-K-m}\left(1-\alpha\right)^j \alpha^{N-K-m-j}\left(\begin{array}{c}\Delta\left(N' + m + j + 2\right) \\ -\Delta\left(N' + m + 1 + j\right)\end{array}\right)\right) \end{aligned}$$

After some re-arrangement of the various terms, this can be re-written

$$P_t^{K-1} = (1-\alpha)\left(\begin{array}{c} (1-\alpha)^{N-K+1}\Delta\left(N' + N - K + 2\right) + \alpha^{N-K+1}\Delta\left(N' + 1\right) \\ +\sum_{s=1}^{N-K}\alpha^{N-K+1-s}\Delta\left(N' + s + 1\right)\left(\begin{array}{c}\sum_{m=0}^{s} C_{s-m}^{N-K-m}\left(1-\alpha\right)^{s-m} \\ -\sum_{m=0}^{s-1} C_{s-1-m}^{N-K-m}\alpha\left(1-\alpha\right)^{s-1-m}\end{array}\right)\end{array}\right)$$

We wish to show this can be rewritten as

$$P_t^{K-1} = (1-\alpha)\left(\sum_{s=0}^{N-K+1} C_s^{N-K+1}\left(1-\alpha\right)^s \alpha^{N-K+1-s}\Delta\left(N' + s + 1\right)\right),$$

so we need to prove

$$\begin{aligned}
&\sum_{m=0}^{s}\frac{(N-K-m)!}{(N-K-s)!\,(s-m)!}\,(1-\alpha)^{s-m}-\sum_{m=0}^{s-1}\frac{(N-K-m)!}{(N-K-s+1)!\,(s-1-m)!}\alpha\,(1-\alpha)^{s-1-m}\\
=\;& C_s^{N-K+1}\,(1-\alpha)^s
\end{aligned}$$

for all $s\in\{1,..,N-K\}$ and $\alpha\in[0,1]$, which can be rewritten

$$\sum_{m=0}^{s}\frac{(N-K-s+m)!}{(N-K-s)!m!}\,(1-\alpha)^m-\sum_{m=0}^{s-1}\frac{(N-K-s+1+m)!}{(N-K-s+1)!m!}\alpha\,(1-\alpha)^m=C_s^{N-K+1}\,(1-\alpha)^s\,.$$

With the change of variable $x=1-\alpha$, this is equivalent to proving

$$\begin{aligned}
&\frac{1}{(N-K-s)!}\sum_{m=0}^{s}(N-K-s+m)\times...\times(1+m)\,x^m\\
&-\frac{1}{(N-K-s+1)!}\sum_{m=0}^{s-1}(N-K-s+1+m)\times...\times(1+m)\,(1-x)\,x^m\\
=\;&\frac{(N-K+1)!}{(N-K+1-s)!s!}x^s
\end{aligned}$$

for all $s\in\{1,..,N-K\}$ and $x\in[0,1]$.

To do this, let

$$\begin{aligned}
f_s(x) &= \frac{1}{(N-K-s)!}\sum_{m=0}^{s}(N-K-s+m)\times...\times(1+m)\,x^m\\
&\quad-\frac{1}{(N-K-s+1)!}\sum_{m=0}^{s-1}(N-K-s+1+m)\times...\times(1+m)\left(x^m-x^{m+1}\right)\\
g_s(x) &= \frac{(N-K+1)!}{(N-K+1-s)!s!}x^s.
\end{aligned}$$

Note that both $f_s(x)$ and $g_s(x)$ are polynomials in x with the highest power of x being s. We have

$$f_s(0)=0=g_s(0)\,.$$

For all $k \in \{1,..,s-1\}$ we have

$$\begin{aligned}
\frac{d^k f_s}{dx^k}(x=0) &= \frac{(N-K-s+k)!}{(N-K-s)!} - \frac{((N-K-s+1+k)! - (N-K-s+k)!k)}{(N-K-s+1)!} \\
&= \frac{(N-K-s+k)!}{(N-K-s)!} - \frac{(N-K-s+k)!}{(N-K-s)!} \\
&= 0 = \frac{d^k g_s}{dx^k}(x=0)
\end{aligned}$$

Finally,

$$\begin{aligned}
\frac{d^s f_s}{dx^s}(x=0) &= \frac{(N-K)!}{(N-K-s)!} + \frac{(N-K)!s}{(N-K-s+1)!} \\
&= \frac{(N-K+1)!}{(N-K-s+1)!} = \frac{d^s g_s}{dx^s}(x=0).
\end{aligned}$$

Thus, we can conclude $f_s(x) = g_s(x)$ for all $x \in [0,1]$ and any $s \in \{1,..,N-K\}$.

So by induction we have proven

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

for any $k \leq N$, when N' agents have decided to participate prior to agent k.

Since all agents participate along the equilibrium path, we can conclude that in equilibrium the price charged to agent k is

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

F Repeated auctions

Instead of setting a price to each agent, the platform can achieve the maximum platform trap (i.e., the same outcome as in Proposition 4) by committing to run an auction in each of the N stages in which there are multiple agents left to sign. Specifically, the auction would invite all agents that haven't yet joined to bid to join the platform (given the number of agents that have already joined), where the agent bidding the highest gets to join at that price and if there are multiple highest bids, then the platform randomly picks one of them to join. We assume all invited agents enter their bids simultaneously, which could involve not entering a bid if they so choose. If there is only a single agent left to join, which can only happen in the last stage, the platform just charges the agent $\Delta(N)$. We show that the only equilibrium resulting from these repeated auctions involves unsigned agents bidding $\Delta(N)$ in each stage, and the last remaining agent joining at the price $P^N = \Delta(N)$ in the last stage. Note, like the contingent pricing function of Proposition 4, such a mechanism relies on commitment (should the highest bid turn out to be less than $\Delta(N)$, the platform still honors the auction, including the case the highest bid is negative).

Proof: Suppose there is one stage left, $n \geq 1$ agents left that haven't joined and $n' \geq 0$ agents have already joined. If $n = 1$, then that last agent is charged $\Delta(n'+1)$ and is left with surplus $u(n')$. If $n' > 1$, then in equilibrium all of these agents bid $b(n'+1) - u(n'+1)$: the winner and all losers obtain $u(n'+1)$ net payoff, with no agent having an incentive to deviate.

Suppose there are $k = 2$ stages left, $n \geq 2$ agents left that haven't joined and $n' \geq 0$ agents have already joined. If $n = 2$, then the equilibrium is for both agents to bid $\Delta(n'+2)$. That is an equilibrium because in the last stage the platform will attract the loser from the second-to-last stage with a price $\Delta(n'+2)$ (based on analysis with one stage left above), so both the winner and the loser obtain net payoff $u(n'+1)$. The only other candidate equilibrium is nobody bids in the second-to-last stage. Then in the last stage there are two agents left and n' agents have joined, so we know from the previous case that both agents obtain net payoff $u(n'+1)$. But then one agent can deviate by submitting a bid slightly below $\Delta(n'+2)$ in the second-to-last stage, obtaining a final net payoff slightly above $u(n'+1)$. If $n > 2$, then the equilibrium is for everyone to bid $b(n'+2) - u(n'+2)$. Indeed, after the winner joins, the losers go to the last stage, all of them bid $b(n'+2) - u(n'+2)$ and one of them joins. So the winner in the second-to-last stage obtains net payoff $u(n'+2)$, and so do all the losers. Again, the only other candidate equilibrium is nobody bids in the second-to-last stage, but this cannot be an equilibrium by the same logic as above.

Now suppose the following induction hypothesis holds. With $k \geq 2$ stages left, $n \geq k$ agents left and $n' \geq 0$ agents having joined, the equilibrium in case $n = k$ is that everyone bids $\Delta(n'+k)$ and all agents obtain net payoff $u(n'+k-1)$, while if $n > k$, the equilibrium is everyone bids $b(n'+k) - u(n'+k)$ and all agents obtain net payoff $u(n'+k)$.

Consider now the case with $k+1 > 2$ stages left, $n \geq k+1$ agents left, and $n' \geq 0$ agents having joined. If $n = k+1$, then the equilibrium is for all agents to bid $\Delta(n'+k+1)$ in the current stage. To see this, note that with these bids, the winner joins and the $n-1 = k$ losers go on to bid $\Delta(n'+k+1)$ in the next stage and everyone obtains net payoff $u(n'+k)$, so no one wants to deviate (at the end there will be $n'+k+1$ agents that join by induction). The only other candidate equilibrium is nobody bids, which is not an equilibrium because in this case all agents would move on to the next stage and obtain net payoff $u(n'+k)$. But then one agent could deviate by bidding slightly below $\Delta(n'+k+1)$ in the current period and obtain net payoff slightly above $u(n'+k)$.

If $n > k+1$, then the equilibrium is for all agents to bid $b(n'+k+1) - u(n'+k+1)$ in the current stage. To see this, note that with these bids, the winner joins and the $n-1 > k$

losers go on to bid $b(n'+k+1)-u(n'+k+1)$ in the next stage and everyone obtains net payoff $u(n'+k+1)$, so no one wants to deviate (at the end there will be $n'+k+1$ agents that join by induction). The only other candidate equilibrium is nobody bids, which once again is not an equilibrium for the same reason as above.

Thus, by induction, the result holds for all N stages. Working backwards to stage 1, we conclude that in stage $k \geq 2$ there are $N-k+1$ agents left and all agents bid $\Delta(N)$, and the last agent also pays $\Delta(N)$.

G Proofs of remaining propositions

This section contains the proofs of the two propositions in Section 6.

G.1 Heterogenous agents

In this section we prove Proposition 6.

Suppose first $\min\{b(N), b(N+S-1)\} > u(0)$. We show that in this case the platform can profitably attract all agents and extract maximum profits regardless of the order of offers. Suppose the platform has attracted the first N agents. If the last agent is the superstar, the platform attracts it with a price of $b(N+S)-u(N)$ and if it is a regular agent the platform attracts it with a price of $b(N+S)-u(N+S-1)$, both of which are positive by assumption.

Now suppose the platform has attracted the first $N-1$ agents it has made offers to and is now facing the second-to-last agent. We know from the previous case that if this agent participates, then so will the last agent. Now suppose the second-to-last agent does not participate. Then, if the last agent is the superstar, the platform attracts it by charging $b(N+S-1)-u(N-1)>0$, and if it is a regular agent, the platform attracts it by charging $b(N)-u(N-1)>0$ (if the second-to-last agent was the superstar) or $b(N+S-1)-u(N+S-2)>0$ (if the second-to-last agent was not the superstar). Thus, the last agent will be attracted regardless of what the second-to-last agent does, so the platform can attract it by charging $b(N+S)-u(N)$ if it is the superstar, or $b(N+S)-u(N+S-1)$ if it is a regular agent.

We can use the same logic all the way back to the first agent, concluding that the platform profitably attracts all agents, charging $b(N+S)-u(N)$ to the superstar and $b(N+S)-u(N+S-1)$ to regular agents, regardless of the order of offers. The platform cannot do any better because it extracts the maximum surplus from each agent given the participation of the other agents.

Next suppose $b(N+S-1) > u(0) \geq u(N-1) > b(N)$, so the superstar is pivotal but none of the regular agents are. In this case, the only way for the platform to achieve maximum profits is by approaching all regular agents first and the superstar last. All regular agents are offered a price $b(N+S) - u(N+S-1)$ and the superstar is offered a price $b(N+S) - u(N)$, resulting in total profits

$$(N+1)\,b(N+S) - u(N) - Nu(N+S-1) > 0.$$

This can be easily verified as an equilibrium outcome given the same backwards induction logic as above. Once again, the platform attains its absolute maximum profit, except that here the order matters. If the platform approaches the superstar earlier than in the last slot, say in slot $k+1 < N+1$, then it would only be able to charge it $b(N+S) - u(k)$ (if the superstar does not participate, the platform will not be able to attract any of the remaining agents), which is strictly less than $b(N+S) - u(N)$.

Finally, suppose $u(N+S-2) > b(N+S-1)$, so all agents are pivotal. Note this implies $u(N-1) > b(N)$ because

$$b(N) - u(N-1) = \Delta(N) \leq \Delta(N+S-1) = b(N+S-1) - u(N+S-2) < 0.$$

In this case, we show that it is optimal for the platform to approach the superstar agent first if it wants to attract any agents. Indeed, by doing so, it can charge $b(N+S) - u(0)$ to the superstar agent and $P^k = b(N+S) - u(S+k-1)$ to agent $k \in \{1, .., N\}$ in the sequence of regular agents. Total profits are

$$(N+1)\,b(N+S) - u(0) - u(S) - u(S+1) - ... - u(S+N-1)\,,$$

which is positive if

$$b(N+S) > \frac{u(0) + \sum_{k=0}^{N-1} u(S+k)}{N+1}.$$

Now suppose the platform approaches the superstar second (after one regular agent). Then it charges $b(N+S) - u(0)$ to the first regular agent, $b(N+S) - u(1)$ to the superstar, $b(N+S) - u(S+1)$ to the second regular agent, and so on. Total profits are now

$$(N+1)\,b(N+S) - u(0) - u(1) - u(S+1) - ... - u(S+N-1)\,.$$

This is lower than the profit obtained by approaching the superstar first since $u(S) < u(1)$. And it is easily seen that the same will be true when the platform approaches the superstar in any but the first position.

G.2 Competing platforms

In this section we prove Proposition 7.

There are four cases to consider.

Case (i): $b_1(1,1) > \max\{u(0,0), b_2(0,2)\}$. In this case, platform 1 will profitably attract agent 2 regardless of what agent 1 does. Thus, if agent 1 joins platform 1, it obtains $b_1(2,0) - P_1^1$, if it joins platform 2 it obtains $b_2(1,1) - P_2^1$, and if it doesn't join either platform, it obtains $u(1,0)$. Platform 2 is not willing to subsidize agent 1 because it has no chance of attracting agent 2, therefore platform 1 will attract agent 1 because $b_1(2,0) > \max\{u(1,0), b_2(1,1)\}$. And platform 1's prices are

$$P_1^1 = P_1^2 = b_1(2,0) - \max\{u(1,0), b_2(1,1)\} > 0,$$

meaning each agent's net payoff is

$$\max\{u(1,0), b_2(1,1)\}.$$

Thus, both agents are worse off with the platforms if $b_2(1,1) < u(0,0)$ and they are better off with the platforms if $b_2(1,1) > u(0,0)$.

Case (ii): $u(0,0) > b_1(1,1) \geq b_2(0,2)$. In this case, if agent 1 joins platform 1, then platform 1 will also profitably attract agent 2 because $b_1(2,0) \geq b_1(1,1) > b_2(1,1)$ and $b_1(2,0) > u(1,0)$. If agent 1 joins platform 2, then agent 2 will still join platform 1 because $b_1(1,1) \geq b_2(0,2) > u(0,1)$. Finally, if agent 1 does not join either platform, then agent 2 joins platform 1 if $b_1(1,0) > u(0,0)$, and it joins neither platform if $u(0,0) \geq b_1(1,0)$. Thus, platform 2 is unwilling to subsidize agent 1 because it has no chance of attracting agent 2. Thus, if agent 1 joins platform 1, its payoff will be $b_1(2,0) - P_1^1$, if it joins platform 2 its payoff is $b_2(1,1) - P_2^1$ and if it joins neither platform, its payoff will be $u(0,0)$ (if $u(0,0) \geq b_1(1,0)$) or $u(1,0)$ (if $b_1(1,0) > u(0,0)$). Since platform 2 is unwilling to subsidize, platform 1 will always attract agent 1 by charging

$$P_1^1 = \begin{cases} b_1(2,0) - u(0,0) & \text{if } u(0,0) \geq b_1(1,0) \\ b_1(2,0) - \max\{u(1,0), b_2(1,1)\} & \text{if } u(0,0) < b_1(1,0) \end{cases}$$

Then platform 1 can profitably attract agent 2 by charging

$$P_1^2 = b_1(2,0) - \max\{b_2(1,1), u(1,0)\} > 0.$$

Platform 1 profits are

$$2b_1(2,0) - u(0,0) - \max\{b_2(1,1), u(1,0)\} \geq 0$$

when $u(0,0) \geq b_1(1,0)$, or

$$2b_1(2,0) - 2\max\{b_2(1,1), u(1,0)\} \geq 0$$

when $u(0,0) < b_1(1,0)$. Note in both cases profits are positive because

$$b_1(2,0) > \max\{b_2(1,1), u(1,0)\}.$$

Agent 1's net payoff is $u(0,0)$ and agent 2's net payoff is $\max\{b_2(1,1), u(1,0)\}$. Since $b_2(1,1) \leq b_2(0,2) < u(0,0)$ in this case, agent 2 is strictly worse off with the platforms. Agent 1 is indifferent.

Case (iii): $b_2(0,2) > b_1(1,1) > u(0,0)$. In this case, if agent 1 joins platform 2, then platform 2 also profitably attracts agent 2. If agent 1 joins platform 1 or neither platform, then platform 1 profitably attracts agent 2 because

$$b_1(2,0) \geq b_1(1,1) > \max\{b_2(1,1), u(0,0)\} \geq \max\{b_2(1,1), u(1,0)\}$$

and $b_1(1,0) > b_2(0,1)$ and $b_1(1,0) \geq b_1(1,1) > u(0,0)$ imply

$$b_1(1,0) > \max\{b_2(0,1), b_1(1,1)\} \geq \max\{b_2(0,1), u(0,0)\}.$$

Thus, since $b_2(0,2) > u(0,0) \geq u(1,0)$, the binding constraint on platform 1 for attracting agent 1 is platform 2. Namely, platform 2 is willing to set its price for agent 1 as low as

$$P_2^1 = -(b_2(0,2) - \max\{b_1(1,1), u(0,1)\}) = -(b_2(0,2) - b_1(1,1)),$$

i.e. subsidize agent 1 by the amount platform 2 would be able to extract from agent 2 were it able to attract agent 1. This means platform 1 must set

$$P_1^1 = b_1(2,0) - (2b_2(0,2) - b_1(1,1))$$

in order to attract agent 1, which then implies platform 1 attracts agent 2 by setting

$$P_1^2 = b_1(2,0) - \max\{b_2(1,1), u(1,0)\}.$$

Total profit for platform 1 is

$$2b_1(2,0) - 2b_2(0,2) + b_1(1,1) - \max\{b_2(1,1), u(1,0)\},$$

which is positive and confirms that platform 1 wins. Net payoffs are $2b_2(0,2) - b_1(1,1)$ for agent

1 and $\max\{b_2(1,1), u(1,0)\}$ for agent 2. Note

$$2b_2(0,2) - b_1(1,1) > u(0,0)$$

in this case, so agent 1 is strictly better off with the platforms. Agent 2 is better off if $b_2(1,1) > u(0,0)$, otherwise it is worse off.

Case (iv): $b_1(1,1) < \min\{b_2(0,2), u(0,0)\}$. In this case, if agent 1 joins either platform, then agent 2 will join the same platform because

$$\begin{aligned} b_2(0,2) &> \max\{b_1(1,1), u(0,1)\} \\ b_1(2,0) &> \max\{b_2(0,2), u(1,0)\} \geq \max\{b_2(1,1), u(1,0)\}. \end{aligned}$$

And if agent 1 joins neither platform, then agent 2 joins neither platform when $u(0,0) \geq b_1(1,0)$), or joins platform 1 when $u(0,0) < b_1(1,0)$. Thus, agent 1's payoffs are $b_1(2,0) - P_1^1$ from joining platform 1, $b_2(0,2) - P_2^1$ from joining platform 2, and $u(0,0)$ or $u(1,0)$ from joining neither platform. Platform 2 is willing to subsidize agent 1 by setting

$$P_2^1 = -(b_2(0,2) - \max\{b_1(1,1), u(0,1)\}),$$

so to win agent 1 platform 1 must set

$$P_1^1 = b_1(2,0) - \max\{u(0,0), 2b_2(0,2) - \max\{b_1(1,1), u(0,1)\}\}$$

if $u(0,0) \geq b_1(1,0)$ or

$$P_1^1 = b_1(2,0) - \max\{u(1,0), 2b_2(0,2) - \max\{b_1(1,1), u(0,1)\}\}$$

if $u(0,0) < b_1(1,0)$.

Platform 1 then attracts both agents and charges

$$P_1^2 = b_1(2,0) - \max\{b_2(1,1), u(1,0)\}$$

to agent 2. Total profits for platform 1 are

$$2b_1(2,0) - \max\{u(0,0), 2b_2(0,2) - \max\{b_1(1,1), u(0,1)\}\} - \max\{b_2(1,1), u(1,0)\}$$

if $u(0,0) \geq b_1(1,0)$ or

$$2b_1(2,0) - \max\{u(1,0), 2b_2(0,2) - \max\{b_1(1,1), u(0,1)\}\} - \max\{b_2(1,1), u(1,0)\}$$

if $u(0,0) < b_1(1,0)$. It is easily verified that under assumptions (2) and (3), these profits are always positive, so platform 1 does indeed profitably attract both agents. Agent 1's net payoffs are then

$$\max\{u(0,0), 2b_2(0,2) - \max\{b_1(1,1), u(0,1)\}\}$$

if $u(0,0) > b_1(1,0)$ or

$$\max\{u(1,0), 2b_2(0,2) - \max\{b_1(1,1), u(0,1)\}\}$$

if $u(0,0) \leq b_1(1,0)$. Agent 2's net payoffs are

$$\max\{b_2(1,1), u(1,0)\}.$$

So if $b_2(0,2) > u(0,0) > b_1(1,1)$, then agent 1 is strictly better off with the platforms and agent 2 is strictly worse off.

Based on the four cases laid out above, if $b_2(1,1) > u(0,0)$ (case 1 in the text of the proposition), then we are in cases (i) or (iii), and under this condition, in both cases both agents are better off with the platforms. If $b_1(1,1) > \max\{u(0,0), b_2(0,2)\}$ and $b_2(1,1) < u(0,0)$ (case 2 in the text of the proposition), then we are in case (i) and both agents are worse off with the platforms. Finally, if $b_2(0,2) > \max\{u(0,0), b_1(1,1)\}$ and $b_2(1,1) < u(0,0)$ (case 3 in the proposition), then we are in cases (iii) or (iv), and it is easily verified that agent 1 is better off with the platforms, while agent 2 is worse off.