Open Transaction Networks and Platform Competition*

Paul Belleflamme† UCLouvain

Sumit Shrivastav† UCLouvain

Kishen Shastry‡ University of Cambridge & GDN

Daniele Rinaldo§ University of Exeter

Jean-Louis Arcand¶ Global Development Network & Geneva Graduate Institute

August 28, 2026

Abstract

This paper studies Open Transaction Networks (OTNs), open protocol-based networks in which buyers and sellers connected through independent applications can transact with one another. India's Open Network for Digital Commerce (ONDC) is the leading real-world instance. We develop a two-sided Hotelling model in which two competing platforms choose whether to adopt the shared protocol on the buyer side, the seller side, both sides, or neither. Existing theory suggests that platforms typically underinvest in interoperability, the ability of users of different networks to interact with one another. We show that full interoperability between competing platforms can instead emerge as the equilibrium outcome of platform competition. Compliance with the protocol expands the trading opportunities of a platform's users, including access to participants beyond both platforms, and it softens competition for exclusive control over users. Much of the resulting profit gain is a transfer, since opening the network allows platforms to withdraw the cross-side subsidies they concede under exclusivity. When compliance is costless, total welfare is maximised by full compliance. Users gain relative to the closed benchmark only when the newly reachable pools are sufficiently valuable, and a planner who weighs users alone often prefers compliance on a single side, which expands reach without raising fees. Beyond arranging interoperability, OTNs redistribute network-generated rents, with implications for the design of inclusive digital public infrastructure.

*We thank Sabhya Rai for excellent research assistance. This research was supported by the Bill & Melinda Gates Foundation through the Global Development Network's Economics of Open Transaction Networks programme (INV-093368). We are particularly grateful to Tanner Lewis for his support and guidance throughout the project, and to Aarti Khanna for coordinating the programme.

†CORE/LIDAM, UCLouvain, Belgium.

‡Cambridge Centre for Alternative Finance, University of Cambridge, UK; Global Development Network, New Delhi, India. Corresponding author: kshastry@gdn.int.

§University of Exeter Business School, University of Exeter, UK.

¶Global Development Network, New Delhi, India; Geneva Graduate Institute, Geneva, Switzerland.

1 Introduction

Digital platform markets such as e-commerce, mobility, and food delivery are increasingly organised around large, vertically integrated intermediaries that control discovery, matching, and transaction flows between buyers and sellers. In these closed architectures, access to trading partners is proprietary: a seller must join a given platform to access its users, and cross-platform interoperability is limited or absent. The economics of these closed platforms is well understood: intermediaries create value by coordinating participation on the two sides of the market and set prices that internalise cross-group network effects (Caillaud and Jullien, 2003; Rochet and Tirole, 2003; Armstrong, 2006). Control over who can transact with whom is central to this business model, and policy analyses argue that it also raises barriers to entry and allows platforms to capture a substantial share of the value generated within their ecosystems (Crémer, de Montjoye, and Schweitzer, 2019; Furman, Coyle, Fletcher, Marsden, and McAuley, 2019).

Recent initiatives illustrate a different market architecture, which we refer to as an *Open Transaction Network* (OTN). In general terms, an OTN is a network in which transactions are coordinated by an open, shared protocol rather than by a single proprietary intermediary: any application that implements the protocol can participate in the network, and buyers and sellers connected through different applications can discover one another and transact. India's Open Network for Digital Commerce (ONDC) is the leading real-world instance and serves as our motivating example throughout the paper.[1] Conceptually, the protocol separates a shared transaction layer, which coordinates discovery, matching, and order routing, from a competitive access layer where applications compete through interface design, search quality, logistics, payments, credit, and other complementary services. As a result, access to potential counterparties is no longer solely controlled by individual platforms but can instead be coordinated through a shared layer of infrastructure. This coordination extends only to compliant interactions, and each user continues to reach the market through the application they have chosen.

The emergence of OTNs is part of a broader movement toward open digital infrastructures and protocol-based ecosystems. Beyond ONDC and Namma Yatri, similar principles underpin developments in open banking, digital identity systems, and interoperable payment networks, where common standards allow multiple providers to compete on top of shared infrastructure. A growing policy discussion on Digital Public Infrastructure (DPI) argues that such architectures can promote inclusion, innovation, and contestability by separating core transactional infrastructure from the services built upon it (D'Silva, Filková, Packer, and Tiwari, 2019; Alonso, Bhojwani, Hanedar, Prihardini, Uña, and Zhabska, 2023; World Bank, 2023). Yet despite growing policy interest, relatively little is known about how shared transaction infrastructures affect the incentives of competing platforms.

[1] ONDC applies this architecture to digital commerce, while Namma Yatri, an open mobility network launched in Bengaluru, applies it to urban ride-hailing. Both are built on the Beckn protocol, an open specification that enables transactions across independent buyer-facing and seller-facing applications.

Understanding these incentives is critical because the long-run sustainability of OTNs depends on whether participation remains privately attractive for market intermediaries as much as on the social benefits the network generates.

From an economic perspective, this architecture changes how network effects translate into market power. In conventional platform markets, firms compete partly by controlling access to users and transactions. In an OTN, by contrast, the transaction layer becomes interoperable while users may continue to singlehome, that is, to join at most one access application on each side. As a result, network effects are no longer fully internalised within proprietary ecosystems, departing from the canonical framework of two-sided platform competition in which platforms seek to internalise cross-group externalities (Rochet and Tirole, 2006; Belleflamme and Peitz, 2021). Instead, some of the value generated through participation becomes accessible across organisational boundaries.

This raises a set of economic questions. How does limiting proprietary control over trading opportunities reshape platform competition? Under what conditions does interoperability strengthen competition, and under what conditions might it instead soften rivalry among access providers? More fundamentally, what happens to the value generated by network effects when access to transactions becomes a shared resource? In particular, how does interoperability change both the amount of value that network effects create and the way that value is distributed across market participants?

A long literature on network compatibility[2] shows that compatibility can raise welfare by expanding the effective network available to users, but also that private and social incentives to adopt it need not coincide (Katz and Shapiro, 1985; Farrell and Saloner, 1985, 1992). In particular, firms with large installed bases may resist compatibility because it weakens exclusivity and erodes the rents generated by network effects (Katz and Shapiro, 1985). Although this literature identifies circumstances in which standardisation can also be excessive (Farrell and Saloner, 1985), the underprovision case has become one of the principal economic justifications for interoperability mandates in digital markets. Recent studies extend the analysis to two-sided platform settings in which the relevant externalities operate across sides (Rasch and Wenzel, 2014; Maruyama and Zennyo, 2015). There too, private compatibility choices need not align with social ones, and platforms may adopt less compatibility than is socially desirable.

OTNs differ, however, from the compatibility arrangements studied in this literature, in two respects. First, compatibility is usually modelled as a bilateral or industry-wide

[2]The earlier literature speaks of *compatibility* between products or networks, while the recent policy debate speaks of *interoperability* between platforms or services. Both denote arrangements that allow users of different systems to interact, so that network effects extend across firm boundaries. The terms differ mainly in connotation: compatibility typically refers to the one-off adoption of a common standard, while interoperability suggests ongoing interconnection between live services, which is the relevant case for OTNs. Throughout the paper, we use compatibility when discussing the earlier literature, interoperability for the ability of one platform's users to transact with the other's, and compliance for a platform's decision to adopt the OTN protocol.

arrangement that links the installed bases of existing networks. Adherence to an OTN protocol is instead unilateral and side-specific: a platform can open the buyer side, the seller side, or both, independently of what its rival does, and interoperability between rival platforms arises only when their compliance choices mesh. Second, the network extends beyond the competing platforms. Because any application that implements the protocol can join, an OTN also connects participants who are attached to neither platform and who reach the market through other protocol-compliant applications. Compliance therefore does more than reallocate trade between two installed bases. It expands the market that a platform's users can reach. Both features change the incentives to adopt interoperability and the way its benefits are distributed across market participants.

This paper shows that the underinvestment logic need not apply to OTNs. We develop a model of platform competition in which two two-sided platforms may adopt an OTN-style protocol on the buyer side, the seller side, both sides, or neither side, generating a rich set of interoperability configurations.

Our central result is that full compliance can be a strictly dominant strategy, so that full interoperability emerges without any mandate. Compliance does two things at once: it makes a platform more attractive to users by enlarging the pool of participants they can transact with, and, once platforms are compliant on opposite sides, it makes them interoperable, which reduces the intensity of competition for exclusive control over users. Together, these gains can outweigh the loss of exclusivity, reversing the standard underinvestment result: platforms may adopt interoperability more readily than their incumbent users would like.

However, the welfare implications depend on the welfare standard applied. Total welfare, the sum of platform profits and user surpluses, is maximised by full compliance everywhere in the costless model, whereas the assessment on behalf of users is more nuanced. Because full interoperability softens competition and leads to higher access fees, a planner who weighs only the surplus of users may prefer a lower degree of interoperability, such as compliance on a single side, which expands users' reach at unchanged fees. User welfare rises above the closed benchmark under full compliance only when the pools of previously unattached participants are sufficiently large in value terms. Where this condition holds, opening the network benefits users as well as platforms. Where it fails, the market still opens but the gains accrue mainly to the platforms.

The deeper contribution of the paper is to reinterpret OTNs through a transfer-of-value lens. Over and above whatever new value it creates, interoperability changes who captures the value generated by network effects. OTNs transfer a portion of network generated rents away from participants who benefit from exclusive access to established network effects, toward participants who were previously excluded from those effects. In the context of digital commerce, these participants may include small merchants, geographically dispersed businesses, informal-sector actors, or consumers who would otherwise remain outside established platform ecosystems. OTNs therefore serve as mechanisms for reallocating the benefits generated by network effects.

This interpretation has important implications for the design of digital public infrastructure. Much of the policy discussion treats interoperability as a public good that markets will systematically underprovide, thereby justifying regulatory mandates or public intervention. Our analysis suggests a different possibility. When interoperability expands participation and redistributes access to network effects, the resulting infrastructure can become both privately self-sustaining and socially valuable. OTNs should therefore be understood as instruments of market design that reshape how network-generated value is created and distributed, a role that goes beyond their usual billing as tools for increasing competition or reducing dependence on dominant intermediaries. The relevant policy question is then whether an architecture that intentionally redistributes network rents can remain self-sustaining while preserving incentives for participation, investment, and innovation.

Formally, our model is a Hotelling (Hotelling, 1929) framework with two competing platforms and two user sides that single-home, following the standard approach in the economics of platforms (Belleflamme and Peitz, 2021). The benchmark corresponds to the canonical two-sided market environment of Armstrong (2006), with cross-side network effects and access fees charged to both buyers and sellers. We then introduce an OTN-like protocol as a platform choice over interoperability. The resulting game has three stages: platforms first choose compliance strategies, then set access fees, and finally buyers and sellers select which platform to join. This structure allows us to compare a closed-platform benchmark with a wide range of OTN architectures, including full interoperability, partial interoperability, and asymmetric compliance decisions.

Related literature: This paper sits at the intersection of the literatures on two-sided platforms, interoperability, and digital public infrastructure.

The starting point is the economics of two-sided markets. Caillaud and Jullien (2003), Rochet and Tirole (2003), Rochet and Tirole (2006) and Armstrong (2006) established the foundational framework for analysing platforms that serve multiple user groups linked by indirect network effects. In these models, platforms create value by coordinating participation across different sides of the market and by internalising cross-group externalities through their pricing and governance decisions. Subsequent work by Evans (2003), Hagiu (2006), Belleflamme, Peitz, and Toulemonde (2022) and Belleflamme and Peitz (2021) further examined platform competition, governance, and value appropriation. A common feature of this literature is that platforms retain proprietary control over interactions between users. This paper departs from this framework by studying an environment in which the ability to connect buyers and sellers can be shared through an open transaction protocol instead of remaining under the exclusive control of individual platforms.

The paper is more directly related to the literature on compatibility and interoperability.

In the context of *direct network effects*, Katz and Shapiro (1985), Farrell and Saloner (1985) and Farrell and Saloner (1992) show that compatibility can increase welfare by expanding the effective size of networks and enabling users of different systems to interact. At the same time, these papers demonstrate that firms may have limited incentives to adopt compatibility because it weakens exclusivity and reduces the rents generated by network effects. This tension between private and social incentives remains central to contemporary debates on digital platform regulation. Our paper is more closely related to the literature on compatibility in two-sided platform environments characterised by *indirect network effects*. Rasch and Wenzel (2014) study compatibility in a platform market where content provision is endogenous, and show that its welfare effects are ambiguous: compatibility is harmful when it leads to less content and beneficial when content increases sufficiently. Closest to our setting, Maruyama and Zennyo (2015) consider platforms choosing compatibility unilaterally and show that symmetric compatibility cannot be an equilibrium unless the cost of achieving compatibility is zero. In recent work, Ekmekci, White, and Wu (2025) analyse platform competition when users respond to net fees and show that interoperability can reduce market dominance and lower prices. In our symmetric duopoly with membership fees, by contrast, interoperability softens price competition and raises fees, so the competitive effect of interoperability depends on the environment into which it is introduced.

Our paper departs from this stream of literature in two ways. First, compliance is side-specific: a platform can choose compliance on the buyer side, the seller side, or both. Second, while compatibility in the existing literature connects two existing networks and reallocates trade between them, OTN compliance also reaches users attached to neither platform, who access the network through other applications built on the protocol.

Another distinction concerns the nature of interoperability itself. OTNs create a shared transaction layer through which interactions can be coordinated across independent applications. In this respect, our paper is related to work on platform openness, modularity, and ecosystem design. West (2003), Boudreau (2010), and Baldwin and Woodard (2009) emphasise how common standards and modular architectures can support competition and innovation at one layer while preserving coordination at another. OTNs represent a particularly distinctive form of platform architecture because they separate the transaction layer from the access layer, allowing applications to compete for users while relying on a common protocol for discovery, matching, and transaction execution. Whereas this literature primarily focuses on innovation, governance, and ecosystem structure, our focus is on how such architectures reshape pricing incentives, interoperability choices, and the distribution of rents in two-sided markets. We therefore provide a bridge between the platform architecture literature and the economics of platform competition.

Finally, our paper is related to the following policy issues. Recent policy reports, including Crémer et al. (2019) and Furman et al. (2019), argue that interoperability can reduce lockin, lower barriers to entry, and enhance competition in digital markets. Our results differ from the conventional underinvestment logic that emerges from this literature. In our setting, interoperability may become privately profitable and emerge endogenously even in the absence of regulatory mandates.

The paper also relates to an emerging institutional and policy literature on Digital Public Infrastructure and open digital ecosystems (D'Silva et al., 2019; Alonso et al., 2023; World Bank, 2023). Initiatives such as India Stack, ONDC, Beckn, open banking frameworks, and interoperable payment systems are motivated by the idea that shared digital rails can broaden participation, foster innovation, and reduce dependence on dominant intermediaries. While much of this literature focuses on governance, implementation, and public policy objectives, less attention has been paid to the incentives of the private intermediaries expected to operate within such systems. Our paper contributes to this discussion by providing a formal economic model of an OTN. More broadly, we introduce a transfer-of-value perspective on interoperability, emphasising how open transaction architectures reshape the distribution of network-generated rents, a dimension that analyses of efficiency and competition tend to leave in the background. In doing so, we connect the economics of platform competition with emerging debates on digital inclusion, market access, and the design of interoperable digital ecosystems.

The remainder of the paper proceeds as follows. Section 2 introduces the two-sided platform model and the OTN compliance choices available to competing platforms. Section 3 derives users' participation decisions and the pricing equilibrium for all feasible compliance configurations. Section 4 analyses the platforms' compliance choices at Stage 1 and identifies the two channels, competitive advantage and interoperability, through which compliance shapes competition. Section 5 studies the Stage 1 equilibrium numerically and evaluates the resulting welfare and surplus implications. Section 6 discusses the broader implications for OTNs, digital public infrastructure, and the governance of interoperable digital markets. Section 7 discusses how the model can be brought to the data. Section 8 concludes.

2 Model

We present here our baseline model, which is a straightforward extension of the two-sided singlehoming model of Armstrong (2006).

Players and utilities: There are two proprietary platforms, 1 and 2, competing with each other for a unit mass of sellers and a unit mass of buyers. Platforms are located at the extreme ends of a unit Hotelling line, and sellers and buyers are uniformly distributed over the unit line. It is assumed that buyers and sellers *singlehome*, which means that they can only be active on a single platform over a given time period.[3] A buyer and a

[3]This assumption is natural in settings where users face switching or learning costs, contractual or exclusivity constraints, or strong preferences over interface and policies.

seller located at x on the unit line derives, respectively, the following utility:

$$u_b(x) = \begin{cases} r_b + \alpha_b \mathbb{S}_1 - b_1 - \tau_b x & \text{joining platform 1} \\ r_b + \alpha_b \mathbb{S}_2 - b_2 - \tau_b (1-x) & \text{joining platform 2} \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

$$u_s(x) = \begin{cases} r_s + \alpha_s \mathbb{B}_1 - s_1 - \tau_s x & \text{joining platform 1} \\ r_s + \alpha_s \mathbb{B}_2 - s_2 - \tau_s (1-x) & \text{joining platform 2} \\ 0 & \text{otherwise,} \end{cases} \tag{2}$$

with

$$\begin{aligned} \mathbb{B}_i &= B_i + \sigma_i B_u + \mu \sigma_i \beta_j B_j, \\ \mathbb{S}_i &= S_i + \beta_i S_u + \mu \beta_i \sigma_j S_j. \end{aligned}$$

The notation is as follows.

- r_b is the stand-alone benefit that a buyer gets when registering with any platform. The stand-alone benefit for a seller is r_s.
- α_b is the cross-side network effect that a buyer enjoys when interacting with an additional seller. The cross-side network effect that a seller enjoys when interacting with an additional buyer is α_s.
- τ_b is the mismatch (transportation) cost per unit of distance that a buyer incurs by registering with a platform that does not correspond to their ideal specification. The corresponding cost for a seller is τ_s.
- $\mathbb{S}_i$ is the "effective" mass of sellers that a buyer can interact with on platform i. Likewise, $\mathbb{B}_i$ is the effective mass of buyers that a seller can interact with on platform i.
- b_i and s_i are the access fee that platform i sets on the buyer and seller side, respectively.
- B_i and S_i are the masses of buyers and of sellers registered with platform i.
- B_u and S_u are the masses of unattached buyers and of unattached sellers.
- $\beta_i \in \{0, 1\}$ is the level of platform i's OTN-compliance on the buyer side, and $\sigma_i \in \{0, 1\}$ is the level on the seller side. A value of one means that the platform complies on that side and a value of zero that it does not.
- μ, with $0 \leq \mu \leq 1$, measures the degree of interoperability across proprietary platforms when they are OTN-compliant on opposite sides.

Interpretation: In a buyer–seller marketplace, each utility component has a natural reading. Beginning with buyers, the stand-alone benefit r_b can be seen as the intrinsic value of being able to browse and purchase through the marketplace interface, with its search tools, reviews, convenient payments, and returns, even before considering how many sellers are present. The cross-side network effect $\alpha_b \mathbb{S}_i$ measures the extra value from the breadth and depth of supply on platform i, since more sellers mean better assortment, lower prices, and a higher chance of finding a good match. The mismatch cost τ_b, multiplied by the distance on the Hotelling line, measures the loss from using a marketplace whose "style" or policies are further from the buyer's ideal, whether in user experience, curation level, language, data and privacy stance, or location and delivery options. The access fee b_i is the monetary or non-monetary price for buyers to participate, such as membership fees, higher shipping or payment fees, or the effective "cost" of being shown ads or giving up data on platform i. On the seller side, the stand-alone benefit r_s corresponds to the baseline value of being listed on the marketplace, from tools for listing and inventory to payment processing, the reputation system, and basic analytics, irrespective of the number of potential transactions. The cross-side network effect $\alpha_s \mathbb{B}_i$ is the benefit from more active buyers on platform i, via higher demand, more orders, and better matching to the seller's products. The mismatch cost τ_s measures the loss from selling on a marketplace whose rules and features are far from the seller's preferred terms, whether commission structure, return policies, branding control, or logistics integration, again proxied by distance. The access fee s_i encompasses the seller-side charges: listing fees, subscription fees for storefronts, or required spending on ads to access demand on platform i.[4]

Extension to OTNs: Alongside the unit masses of buyers and sellers who choose to register with either platform 1 or platform 2, there exist buyers and sellers who are *unattached*, that is, registered with neither platform. These users are not outside the market altogether: they can be reached through the open network, typically via lightweight, protocol-compliant applications that we do not model as strategic platforms. Although we treat these users as exogenous, we can think of them as users for whom the proprietary platforms' net value is negative, so they rationally choose the outside option rather than joining. Concretely, they may (i) face high fixed or variable fees that make participation unprofitable, which weighs especially on small or low-intensity users, (ii) have strong preference or regulatory mismatches with existing platforms' rules, quality, or data practices, or (iii) be constrained by digital exclusion, such as a lack of smartphones, payment instruments, or trust, so that the expected benefits from joining do not compensate these costs and risks. In this sense, unattached users represent latent demand on each

[4] Similar interpretations can be given in the context of ride-hailing platforms, for example. The stand-alone benefit is the basic value of the app: on-demand rides for riders, and dispatch and earnings opportunities for drivers. The cross-side effect is that more drivers improve riders' wait times and choice, and more riders improve drivers' utilisation. The mismatch cost captures how far the platform's pricing, quality, rules, and coverage are from a user's ideal. The access fee collects all charges that effectively make participation more expensive for each side, such as booking fees and surcharges for riders, and commissions and fixed fees for drivers.

side that can be activated if a new architecture, such as an OTN, offers access under more favourable terms, without assuming that all potential users are already captured by incumbent platforms.

Initially, buyers and sellers registered with platform i cannot transact with unattached users. To enable such transactions, platform i must adhere to the OTN protocol on one side and/or the other. For instance, if platform i is "OTN-compliant" on the buyer side, its registered buyers can engage with unattached sellers. Symmetrically, OTN compliance on the seller side allows registered sellers to interact with unattached buyers. Therefore, OTN-compliance is an investment in quality, enhancing the platform's value for specific user groups. We refer to this as the *external* effect of compliance. Additionally, compliance can also create an *internal* effect by changing the nature of competition between the two proprietary platforms. If, say, platform 1 is OTN-compliant on the buyer side, while platform 2 is compliant on the seller side, they become partially interoperable (at degree μ), enabling buyers from platform 1 to conduct transactions with sellers of platform 2.

We translate these two effects in the model as follows. Suppose, for instance, that platform 1 is OTN-compliant on the buyer side ($\beta_1 = 1$) and platform 2 is *not* compliant on the seller side ($\sigma_2 = 0$). Then, buyers on platform 1 have access not only to the products offered by the sellers on platform 1, but also to those offered by the unattached sellers. The mass of sellers they can interact with is therefore $\mathbb{S}_1 = S_1 + S_u$, where S_1 is the mass of "attached" users who have decided to register with platform 1, while S_u is the *effective* mass of unattached sellers. By "effective", we mean the number of unattached sellers weighted by the quality or intensity of the transactions that can be conducted with them: how frequently they are available, how well they are integrated into the protocol, and how well their offers match buyers' needs. A completely analogous interpretation applies to B_u on the buyer side: when a platform becomes OTN-compliant on the seller side, its attached sellers can reach both the buyers registered on that platform and a mass of unattached buyers. In both cases, B_u and S_u are thus reduced-form parameters that capture the additional trading opportunities created by the OTN beyond those provided by the proprietary platforms themselves. They allow us to distinguish between (i) the scale of demand and supply internal to each platform, B_i and S_i, and (ii) the extra scale made available through the network, B_u and S_u, which is only accessible if the platform chooses to comply with the OTN protocol.

Assume now that platform 2 is OTN-compliant on the seller side ($\sigma_2 = 1$). This enables transactions between platform 1's buyers and platform 2's sellers. From the perspective of buyers on platform 1, the mass of sellers they can access is now $\mathbb{S}_1 = S_1 + S_u + \mu S_2$, where $\mu \in [0, 1]$ measures the extent of interoperability between the two proprietary platforms. Interoperability may be imperfect in practice because technical integration can be partial, with only some product categories or features federated, because contractual frictions may limit cross-platform listing or visibility, and because users may not perceive or exploit all cross-platform matches that are theoretically feasible. For this reason, we allow μ to be strictly below one, even if in a fully idealised OTN one would expect $\mu = 1$. At the same time, we keep μ as a free parameter so that we can also clearly disentangle the internal and external effects of compliance.

Timing: We consider the following three-stage game. In stage 1, platforms choose their compliance strategy. They have the following four options: no compliance (N: $\beta_i = \sigma_i = 0$), compliance on the buyer side (B: $\beta_i = 1, \sigma_i = 0$), compliance on the seller side (S: $\beta_i = 0, \sigma_i = 1$), or compliance on both sides (A: $\beta_i = \sigma_i = 1$). In stage 2, platforms set access fees for buyers and sellers. In stage 3, attached buyers and sellers choose which platform to register with. We solve the game for its subgame-perfect equilibrium.

Limitations: One limitation of the framework of Armstrong (2006) is that it only considers access (membership) fees. It is thus implicitly assumed that a platform's pricing does not depend on how many interactions each user actually has with the other side, only on joining decisions. This means the model cannot capture margins like commissions per sale or per-ride, or how platforms might use per-transaction fees as an alternative instrument to manage the intensity of platform usage.[5] However, one way to accommodate transaction fees within this framework is to reinterpret (or decompose) the "access fee" as the expected total payment a user anticipates, which already bundles a fixed membership charge and a per-transaction fee times expected usage.

Another limitation is that the model abstracts from any form of same-side network effects, whereas in our two leading examples, such effects are clearly present. In marketplaces, additional sellers increase variety, a positive cross-side effect for buyers, but they also intensify competition among sellers themselves, a negative same-side network effect. In ride-hailing, more users on a given side can create congestion or crowding, with more drivers chasing the same pool of riders or more riders competing for a limited supply of drivers, so that same-side participation is not purely beneficial.[6] Our setting can nevertheless be viewed as a reasonable approximation for two reasons. First, to the extent that same-side competition mainly feeds through prices or effective quality, its impact is partially captured in reduced-form parameters, for instance in the stand-alone utilities or in the strength of the cross-side network effects, so the model still speaks to how interoperability shifts market power and participation, even if it does not trace the full price-formation process among same-side users. Second, explicitly modelling same-side network effects, especially when they are negative and possibly asymmetric, would substantially increase the analytical complexity and blur the central mechanisms we want to highlight, namely how making transactional reachability interoperable affects platform

[5] Armstrong (2006) points out that a continuum of equilibria exists when platforms compete with access and usage fees, which undermines the predictive power of the model. Reisinger (2014) shows that introducing heterogeneous trading behaviour selects a unique equilibrium with intuitive properties.

[6] Belleflamme and Toulemonde (2009) analyse platform competition when such negative intra-group externalities are present alongside positive cross-group effects. Belleflamme and Peitz (2021) discuss the broader role of within-group external effects in platform markets.

pricing incentives, the distribution of users across platforms, and the rents that can sustain an open transaction network.

3 Pricing and participation decisions

In this section, we analyse pricing and participation decisions of platforms and users, respectively, following platforms' compliance decisions in stage 1.

3.1 Users' participation decisions (Stage 3)

At **Stage 3**, the "attached" users choose between the two platforms. To determine the demands for participation, we identify on each side the user who is indifferent between joining platform 1 or platform 2. Using expressions (1) and (2), we see that the indifferent buyer is located at x_b and the indifferent seller at x_s with:

$$
\begin{aligned}
r_b + \alpha_b \mathbb{S}_1 - b_1 - \tau_b x_b &= r_b + \alpha_b \mathbb{S}_2 - b_2 - \tau_b (1 - x_b) \\
&\Leftrightarrow \quad x_b = \frac{1}{2} + \frac{b_2 - b_1 + \alpha_b (\mathbb{S}_1 - \mathbb{S}_2)}{2\tau_b}, \\
r_s + \alpha_s \mathbb{B}_1 - s_1 - \tau_s x_s &= r_s + \alpha_s \mathbb{B}_2 - s_2 - \tau_s (1 - x_s) \\
&\Leftrightarrow \quad x_s = \frac{1}{2} + \frac{s_2 - s_1 + \alpha_s (\mathbb{B}_1 - \mathbb{B}_2)}{2\tau_s}.
\end{aligned}
$$

We assume throughout the analysis that the stand-alone benefits r_b and r_s are sufficiently large to guarantee that all attached buyers and sellers participate. As a result, buyers located between 0 and x_b prefer to join platform 1, while buyers located between x_b and 1 prefer to join platform 2. It follows that $B_1 = x_b = 1 - B_2$. Symmetrically, we have that $S_1 = x_s = 1 - S_2$.

Using the definitions of $\mathbb{B}_i$ and $\mathbb{S}_i$, we can compute:

$$
\begin{aligned}
\mathbb{B}_1 - \mathbb{B}_2 &= (1 - \mu\beta_1\sigma_2) B_1 - (1 - \mu\sigma_1\beta_2)(1 - B_1) + (\sigma_1 - \sigma_2) B_u, \\
\mathbb{S}_1 - \mathbb{S}_2 &= (1 - \mu\sigma_1\beta_2) S_1 - (1 - \mu\beta_1\sigma_2)(1 - S_1) + (\beta_1 - \beta_2) S_u.
\end{aligned}
$$

It follows that:

$$
\begin{cases}
B_1 = \dfrac{1}{2} + \dfrac{(b_2 - b_1) + \alpha_b \left((1 - \mu\sigma_1\beta_2) S_1 - (1 - \mu\beta_1\sigma_2)(1 - S_1) + (\beta_1 - \beta_2) S_u\right)}{2\tau_b}, \\
S_1 = \dfrac{1}{2} + \dfrac{(s_2 - s_1) + \alpha_s \left((1 - \mu\beta_1\sigma_2) B_1 - (1 - \mu\sigma_1\beta_2)(1 - B_1) + (\sigma_1 - \sigma_2) B_u\right)}{2\tau_s}.
\end{cases}
$$

Solving the previous system of equations for B_1 and S_1, we get the demand functions for participation on the two sides (recalling that $B_2 = 1 - B_1$ and $S_2 = 1 - S_1$):

$$\begin{cases} B_1 = \frac{1}{2} - \frac{4\tau_s(b_1-b_2)+2\alpha_b\gamma_p(s_1-s_2)}{2\left(4\tau_s\tau_b-\gamma_p^2\alpha_b\alpha_s\right)} + K_b,\ B_2 = 1 - B_1, \\ S_1 = \frac{1}{2} - \frac{4\tau_b(s_1-s_2)+2\alpha_s\gamma_p(b_1-b_2)}{2\left(4\tau_s\tau_b-\gamma_p^2\alpha_b\alpha_s\right)} + K_s,\ S_2 = 1 - S_1, \end{cases} \tag{3}$$

where:

$$\begin{aligned} K_b &\equiv \frac{\alpha_b\gamma_m\left(\gamma_p\alpha_s - 2\tau_s\right)}{2\left(4\tau_s\tau_b - \gamma_p^2\alpha_b\alpha_s\right)} + \frac{2\alpha_b\alpha_s\gamma_p\delta_\sigma B_u + 4\alpha_b\tau_s\delta_\beta S_u}{2\left(4\tau_s\tau_b - \gamma_p^2\alpha_b\alpha_s\right)}, \\ K_s &= -\frac{\alpha_s\gamma_m\left(\gamma_p\alpha_b - 2\tau_b\right)}{2\left(4\tau_s\tau_b - \gamma_p^2\alpha_b\alpha_s\right)} + \frac{2\alpha_b\alpha_s\gamma_p\delta_\beta S_u + 4\alpha_s\tau_b\delta_\sigma B_u}{2\left(4\tau_s\tau_b - \gamma_p^2\alpha_b\alpha_s\right)}, \end{aligned}$$

$$\begin{aligned} \gamma_p &\equiv 2 - \mu\left(\sigma_1\beta_2 + \beta_1\sigma_2\right),\ \gamma_m \equiv \mu\left(\sigma_1\beta_2 - \beta_1\sigma_2\right), \\ \delta_\beta &\equiv \beta_1 - \beta_2 \text{ and } \delta_\sigma \equiv \sigma_1 - \sigma_2. \end{aligned}$$

To ensure that demands for participation on platform i decrease with platform i's fees (b_i and s_i), we assume that: $\tau_s\tau_b > \alpha_b\alpha_s$. For future reference, it is useful to note the following.

Finding 1. *The attached users' participation decisions are unaffected by the masses of unattached users if the two proprietary platforms make identical compliance decisions at the first stage of the game.*

If platforms make the same compliance decisions, then $\beta_1 = \beta_2$ and $\sigma_1 = \sigma_2$. It follows from the above expressions that $\delta_\beta = \delta_\sigma = 0$, so that (K_b, K_s), and thus (B_i, S_i), do not depend on (B_u, S_u). The intuition is straightforward: unattached users only matter if platforms are compliant, and when both platforms are compliant on the same side, they offer exactly the same additional benefit to users on that side, leaving the identity of the marginal (indifferent) user, and with it participation demands, unchanged. In short, by choosing the same side on which to be compliant, platforms effectively neutralise each other.

The user surpluses, written for the full mass on each side and integrating the linear utility across the Hotelling segment served by each platform, are

$$UB = \sum_{i=1}^{2} \int_{B_i\text{-segment}} \left(r_b + \alpha_b\mathbb{S}_i - b_i - \tau_b d(x,i)\right) dx, \tag{4}$$

$$US = \sum_{i=1}^{2} \int_{S_i\text{-segment}} \left(r_s + \alpha_s\mathbb{B}_i - s_i - \tau_s d(x,i)\right) dx, \tag{5}$$

and total welfare is $W = \pi_1 + \pi_2 + UB + US$.

3.2 Platforms' pricing decisions (Stage 2)

At **stage 2**, Platform i ($i = 1, 2$) chooses access fees b_i and s_i to maximise its profit $\pi_i = b_i B_i + s_i S_i$, given platforms' compliance choices. Solving the system of four first-order conditions, we find the following equilibrium fees:

$$
\begin{aligned}
b_1^* &= \tau_b - \frac{\gamma_p}{2}\alpha_s + \frac{6\tau_b\left(4\tau_b\tau_s - \gamma_p^2\alpha_b\alpha_s\right)K_b - \gamma_p\left(2\alpha_b + \alpha_s\right)\left(4\tau_b\tau_s - \gamma_p^2\alpha_b\alpha_s\right)K_s}{D}, \\
b_2^* &= \tau_b - \frac{\gamma_p}{2}\alpha_s - \frac{6\tau_b\left(4\tau_b\tau_s - \gamma_p^2\alpha_b\alpha_s\right)K_b - \gamma_p\left(2\alpha_b + \alpha_s\right)\left(4\tau_b\tau_s - \gamma_p^2\alpha_b\alpha_s\right)K_s}{D}, \\
s_1^* &= \tau_s - \frac{\gamma_p}{2}\alpha_b + \frac{6\tau_s\left(4\tau_b\tau_s - \gamma_p^2\alpha_b\alpha_s\right)K_s - \gamma_p\left(\alpha_b + 2\alpha_s\right)\left(4\tau_b\tau_s - \gamma_p^2\alpha_b\alpha_s\right)K_b}{D}, \\
s_2^* &= \tau_s - \frac{\gamma_p}{2}\alpha_b - \frac{6\tau_s\left(4\tau_b\tau_s - \gamma_p^2\alpha_b\alpha_s\right)K_s - \gamma_p\left(\alpha_b + 2\alpha_s\right)\left(4\tau_b\tau_s - \gamma_p^2\alpha_b\alpha_s\right)K_b}{D},
\end{aligned}
$$

where

$$D \equiv 36\tau_b\tau_s - \gamma_p^2\left(2\alpha_b + \alpha_s\right)\left(\alpha_b + 2\alpha_s\right).$$

The second-order conditions impose that $D > 0$. Recalling that $\gamma_p = 2 - \mu\left(\sigma_1\beta_2 + \beta_1\sigma_2\right)$, we see that the largest value is reached for $\beta_i = \sigma_i = \mu = 0$. So, we impose the most stringent condition, that is:

$$D_0 \equiv 9\tau_b\tau_s - \left(2\alpha_b + \alpha_s\right)\left(\alpha_b + 2\alpha_s\right) > 0 \tag{6}$$

We can now compute the equilibrium participation levels as:

$$
\begin{aligned}
B_1^* &= \frac{1}{2} + \frac{\left(12\tau_b\tau_s - \gamma_p^2\alpha_s\left(\alpha_b + 2\alpha_s\right)\right)K_b - 4\gamma_p\tau_s\left(\alpha_b - \alpha_s\right)K_s}{D}, \\
S_1^* &= \frac{1}{2} + \frac{4\tau_b\gamma_p\left(\alpha_b - \alpha_s\right)K_b + \left(12\tau_b\tau_s - \alpha_b\gamma_p^2\left(2\alpha_b + \alpha_s\right)\right)K_s}{D}, \\
B_2^* &= 1 - B_1^* \text{ and } S_2^* = 1 - S_1^*.
\end{aligned}
$$

To guarantee an interior solution, we impose:

$$
\begin{cases}
-D/2 < \left(12\tau_b\tau_s - \gamma_p^2\alpha_s\left(\alpha_b + 2\alpha_s\right)\right)K_b - 4\gamma_p\tau_s\left(\alpha_b - \alpha_s\right)K_s < D/2, \\
-D/2 < 4\tau_b\gamma_p\left(\alpha_b - \alpha_s\right)K_b + \left(12\tau_b\tau_s - \alpha_b\gamma_p^2\left(2\alpha_b + \alpha_s\right)\right)K_s < D/2,
\end{cases}
$$

which, substituting for K_b and K_s, is equivalent to:

$$
\begin{cases}
-D < 2\gamma_p\alpha_s\left(\alpha_b + 2\alpha_s\right)\delta_\sigma B_u + 12\alpha_b\tau_s\delta_\beta S_u - \gamma_m\left(6\alpha_b\tau_s - \gamma_p\alpha_s\left(\alpha_b + 2\alpha_s\right)\right) < D, \\
-D < 2\gamma_p\alpha_b\left(2\alpha_b + \alpha_s\right)\delta_\beta S_u + 12\alpha_s\tau_b\delta_\sigma B_u + \gamma_m\left(6\alpha_s\tau_b - \gamma_p\alpha_b\left(2\alpha_b + \alpha_s\right)\right) < D.
\end{cases}
$$

The latter conditions are the most stringent when one platform is compliant on both sides and the other is not compliant, that is, when either $\delta_\beta = \delta_\sigma = 1$ or $\delta_\beta = \delta_\sigma = -1$. In both cases, $\gamma_p = 2$ and $\gamma_m = 0$. Rearranging, we get the following sufficient conditions for an interior equilibrium in all subgames:

$$\begin{cases} \alpha_s (\alpha_b + 2\alpha_s) B_u + 3\alpha_b \tau_s S_u < D_0, \\ \alpha_b (2\alpha_b + \alpha_s) S_u + 3\alpha_s \tau_b B_u < D_0. \end{cases} \tag{7}$$

We can now express the platforms' equilibrium profits at stage 2 as:

$$\begin{cases} \pi_1^* (\beta_1, \beta_2, \sigma_1, \sigma_2; B_u, S_u) = b_1^* B_1^* + s_1^* S_1^*, \\ \pi_2^* (\beta_1, \beta_2, \sigma_1, \sigma_2; B_u, S_u) = b_2^* B_2^* + s_2^* S_2^*. \end{cases} \tag{8}$$

3.2.1 Welfare analysis

Finally, let us derive the users' surpluses at the pricing equilibrium. The buyers attached to platform 1 occupy the segment $[0, B_1^*]$ on the Hotelling line by the indifferent-buyer construction, and analogously for platform 2 on the segment $[B_1^*, 1]$. Integrating the buyer utility over each platform's segment gives a closed form. Proceeding analogously for sellers, we find:

$$UB = \sum_{i=1}^{2} B_i^* \left[r_b + \alpha_b \mathbb{S}_i^* - b_i^* - \tfrac{1}{2} \tau_b B_i^* \right], \tag{9}$$

$$US = \sum_{i=1}^{2} S_i^* \left[r_s + \alpha_s \mathbb{B}_i^* - s_i^* - \tfrac{1}{2} \tau_s S_i^* \right]. \tag{10}$$

For policy purposes, we will need three welfare measures derived from these objects. Total welfare is $W = \pi_1 + \pi_2 + UB + US$. User welfare is $UW = UB + US$. The weighted welfare $W_\theta = \theta(\pi_1 + \pi_2) + (1 - \theta)(UB + US)$ traces the trade-off between the producer-inclusive and the user-only conventions as the weight $\theta \in [0, 1]$ varies. The Bergson-Samuelson default corresponds to $\theta = 1/2$ and produces W_θ proportional to total welfare. The consumer-welfare standard typical of platform regulation corresponds to $\theta = 0$.

3.2.2 Pricing under symmetric compliance choices

We consider here the four subgames in which platforms make identical compliance choices: no compliance (N, N), full compliance (A, A), and the two subgames of same-side compliance (B, B) and (S, S). Let us start with our benchmark, (N, N), which corresponds exactly to the two-sided singlehoming situation analysed by Armstrong (2006). Here, $\mathbb{B}_i = B_i$ and $\mathbb{S}_i = S_i$. Applying the preceding analysis, we find the equilibrium prices as:

$$b_{NN} = \tau_b - \alpha_s \text{ and } s_{NN} = \tau_s - \alpha_b.$$

(Throughout, we use the subscript KL to refer to the subgame in which platform 1 chooses strategy K and platform 2, strategy L, with $K, L \in \{N, B, S, A\}$.) As Armstrong (2006) explains, equilibrium access fees on each side equal the marginal cost (here zero) plus the corresponding mismatch cost minus the cross-side network benefit that users of that side generate for the opposite side. As platforms set the same fees, buyers and sellers split equally: $B_{NN} = S_{NN} = \frac{1}{2}$. We can then compute the platforms' profits and the users' net utilities at equilibrium as:

$$\begin{aligned}
\pi_{NN} &= \tfrac{1}{2}\left(\tau_b + \tau_s - \alpha_s - \alpha_b\right), \\
UB_{NN} &= \int_0^{\frac{1}{2}} \left(r_b + \alpha_b \tfrac{1}{2} - (\tau_b - \alpha_s) - \tau_b x\right) dx = \tfrac{1}{4}\left(2r_b + \alpha_b + 2\alpha_s\right) - \tfrac{5}{8}\tau_b, \\
US_{NN} &= \int_0^{\frac{1}{2}} \left(r_s + \alpha_s \tfrac{1}{2} - (\tau_s - \alpha_b) - \tau_s x\right) dx = \tfrac{1}{4}\left(2r_s + \alpha_s + 2\alpha_b\right) - \tfrac{5}{8}\tau_s.
\end{aligned}$$

(By symmetry, we report user surpluses per platform in this subsection. The surpluses over the full mass on each side are twice these values.)

In subgame (B, B), both platforms are compliant on the buyer side. As a result, buyers can interact with the unattached sellers as well: $\mathbb{S}_1 = S_1 + S_u$ and $\mathbb{S}_2 = S_2 + S_u$. Buyers thus receive additional benefits on each platform, but the difference $\mathbb{S}_1 - \mathbb{S}_2$ is unchanged, since the extra access to S_u cancels out. On the other hand, platforms remain not compliant on the seller side, meaning that $\mathbb{B}_1$ and $\mathbb{B}_2$ do not change, and neither does their difference. Demands for participation therefore remain as in the benchmark, and the rest of the analysis carries through unchanged. The same reasoning applies to subgame (S, S), leading to the conclusion that same-side compliance does not affect profits relative to the benchmark:

$$\pi_{BB} = \pi_{SS} = \pi_{NN} = \tfrac{1}{2}\left(\tau_b + \tau_s - \alpha_s - \alpha_b\right).$$

Users' surplus, however, increases on the compliant side:

$$\begin{aligned}
UB_{BB} &= \tfrac{1}{4}\left(2r_b + \alpha_b + 2\alpha_s\right) - \tfrac{5}{8}\tau_b + \tfrac{1}{2}\alpha_b S_u > UB_{NN} = UB_{SS}, \\
US_{SS} &= \tfrac{1}{4}\left(2r_s + \alpha_s + 2\alpha_b\right) - \tfrac{5}{8}\tau_s + \tfrac{1}{2}\alpha_s B_u > US_{NN} = US_{BB}.
\end{aligned}$$

Subgame (A, A) both resembles and differs from the other three symmetric subgames. It resembles them in that it yields the same participation demands. It differs because compliance on both sides generates interoperability across the two platforms, which further expands users' benefits: $\mathbb{B}_i = B_i + B_u + \mu B_j$ and $\mathbb{S}_i = S_i + S_u + \mu S_j$. As a result, $\mathbb{B}_1 - \mathbb{B}_2 = (1-\mu)(B_1 - B_2)$ and $\mathbb{S}_1 - \mathbb{S}_2 = (1-\mu)(S_1 - S_2)$. This lowers demand elasticity with respect to relative prices and, thereby, relaxes price competition. This is because interoperability reduces users' marginal gains from switching. In the non-interoperable benchmark, participation on a given platform grants exclusive access to that platform's users, so small fee differentials translate into discrete changes in the set of potential trading partners, making cross-platform substitution highly responsive to price. With interoperability, users on one platform can interact with users on the rival platform as well, so a marginal change in fees alters the effective interaction set less. As a result, the competitive externality associated with undercutting a rival is weakened, best-response functions become less aggressive, and equilibrium prices are correspondingly higher than in the non-interoperable case.

Using the previous analysis, we compute the equilibrium fees and profits as:

$$b_{AA} = \tau_b - (1-\mu)\,\alpha_s \text{ and } s_{AA} = \tau_s - (1-\mu)\,\alpha_b,$$

$$\pi_{AA} = \tfrac{1}{2}\left(\tau_b + \tau_s - \alpha_s - \alpha_b\right) + \tfrac{1}{2}\mu\left(\alpha_b + \alpha_s\right) > \pi_{NN} = \pi_{BB} = \pi_{SS}.$$

Fees and profits grow larger as interoperability strengthens. At the limit ($\mu = 1$), they no longer depend on the cross-side network effects. Users are affected in contrasting ways: they enjoy larger network benefits but pay higher fees. The net impact is ambiguous, whether we compare full compliance with no compliance or with same-side compliance

$$UB_{AA} - UB_{NN} = \tfrac{1}{2}\alpha_b S_u + \tfrac{1}{4}\mu\left(\alpha_b - 2\alpha_s\right) \text{ and } US_{AA} - US_{NN} = \tfrac{1}{2}\alpha_s B_u + \tfrac{1}{4}\mu\left(\alpha_s - 2\alpha_b\right).$$

$$UB_{AA} - UB_{BB} = \tfrac{1}{4}\mu\left(\alpha_b - 2\alpha_s\right) \text{ and } US_{AA} - US_{SS} = \tfrac{1}{4}\mu\left(\alpha_s - 2\alpha_b\right).$$

We note that both the latter expressions can be negative if cross-side network effects are similar on the two sides ($\alpha_b \approx \alpha_s$).

Collecting the previous results, we can state the following.

Finding 2. *When both platforms are* compliant on the same side*, (i) they earn the same profits as in the benchmark, because competition fully dissipates the additional benefits they confer to users on the compliant side, and (ii) users are strictly better off on the compliant side and indifferent on the other side. When both platforms are* compliant on both sides*, (iii) they earn larger profits than in the benchmark, because the interoperability induced by full compliance relaxes price competition, and (iv) if cross-side network effects have similar intensities, extending compliance from one side to both reduces users' surplus.*

4 Platforms' compliance choices (Stage 1)

We now turn to platforms' compliance choices. Expressions (8) can be used to generate the second-stage equilibrium profits in the 16 subgames resulting from the platforms' first-stage decisions (recall that each platform has four compliance strategies in stage 1, $C_i \in \{N, B, S, A\}$). Because platforms are symmetric, we can focus our attention on platform 1 without loss of generality. That leaves us with ten subgames to analyse. Among them is subgame (N, N), which corresponds to the benchmark model of Armstrong (2006), as neither platform is OTN-compliant. We categorise the remaining nine subgames into four distinct classes, as depicted on Figure 1.

- **Single platform compliance.** If only platform 1 is compliant, then compliance generates purely external effects. The three subgames in this class are: (B, N), (S, N), and (A, N).

- **Same-side compliance.** If platforms are compliant on the same side and on this side only, they both generate the same extra benefits for a specific group of users and remain thus symmetric (as noted in Finding 1). The two subgames in this class are: (B, B) and (S, S).
- **Opposite-sides compliance.** If platforms are compliant on opposite sides, then they become interoperable, which means that compliance induces both external and internal effects. The three subgames in this class are: (S, B), (A, B), and (A, S).
- **Full compliance**. In subgame (A, A), both platforms are compliant on both sides, which combines same-side and opposite-sides compliance.

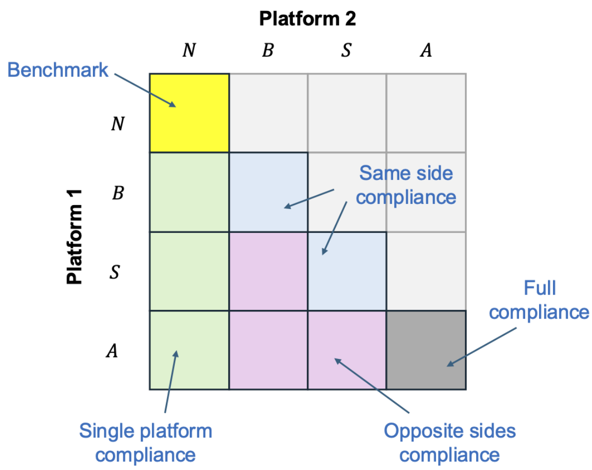

Figure 1: Classes of subgames in Stages 2 and 3. Each cell of the 4×4 compliance matrix is assigned to one of five categories: the closed benchmark (N, N), single platform compliance, same-side compliance, opposite-sides compliance, and full compliance.

Fully solving for the Stage-1 Nash equilibrium is, in principle, feasible but technically demanding, given the large number of parameters in the model. We therefore resort to numerical simulations in the next section and, in this section, focus on the two main channels through which compliance choices shape price competition between platforms. First, a platform that is more compliant than its rival gains a competitive advantage by becoming more attractive to attached users. Second, when platforms are compliant on different sides, they become interoperable, which softens price competition. From this analysis, we derive several general insights that will help interpret the numerical results presented below.

4.1 Compliance as competitive advantage

By becoming compliant on, say, the buyer side, a platform increases the benefits it offers to its buyers, as they can now interact with the pool of unattached sellers. This gives the platform a competitive advantage in attracting buyers *unless, as shown in Finding*

1, the rival platform is also compliant on the buyer side. In this sense, the presence of unattached users pushes platforms to be *compliant on one more side* than their rival. Conversely, when the rival is "more compliant", matching its compliance becomes necessary to avoid falling at a competitive disadvantage. To isolate the incentives generated by these mechanisms, we abstract from the possibility that compliance on opposite sides makes platforms interoperable and impose $\mu = 0$. The following two instructive results emerge from our analysis.

Finding 3. *Suppose $\mu = 0$ and consider a platform that is not compliant on any side. This platform's incentive to be compliant (on one side or on both sides) is the largest when the other platform is not compliant and the smallest when the other platform is compliant on both sides.*

Formally, we can establish the following inequalities (see Appendix A for the proof):

$$\pi_{KN} - \pi_{NN} > \begin{cases} \pi_{KB} - \pi_{NB} \\ \pi_{KS} - \pi_{NS} \end{cases} > \pi_{KA} - \pi_{NA}, \text{ with } K \in \{B, S, A\}.$$

It follows that if a platform prefers any form of compliance to non compliance when the rival is fully compliant ($\pi_{KA} - \pi_{NA} > 0$), then the same applies for any other rival's strategy, meaning that non compliance (strategy N) is a dominated strategy. Reversing the argument, if a platform prefers non compliance to any form of compliance when the rival is not compliant ($\pi_{KN} - \pi_{NN} < 0$), then the same applies for any other rival's strategy, meaning that full compliance (strategy A) is a dominated strategy.

So far, non compliance has been the common starting point. Now, we also consider a platform's incentives to become compliant on one additional side. We obtain the following result.

Finding 4. *Suppose $\mu = 0$. Becoming compliant on one side when the rival is not is more profitable than becoming compliant on that side when the rival already is (irrespective of the platforms' compliance on the other side).*

The formal relationships are (see Appendix A for the proof):

$$\pi_{KN} - \pi_{NN} = \pi_{AL} - \pi_{LL} > \pi_{KK} - \pi_{NK} = \pi_{AA} - \pi_{LA}, \text{ with } K, L \in \{B, S\}, K \neq L.$$

To understand what these relationships mean, set $K = B$ and $L = S$. On the left-hand side, we measure the platform's incentive to be compliant on the buyer side when the rival is not. Interestingly, this incentive is the same whether the platform is not compliant on the seller side ($\pi_{BN} - \pi_{NN}$) or already is ($\pi_{AS} - \pi_{SS}$). The right-hand side, which is smaller, captures the incentive to be compliant on the buyer side when the rival *is already compliant on that side.* Again, the magnitude of the incentive does not depend on whether the platform is compliant on the seller side ($\pi_{AA} - \pi_{SA}$) or not ($\pi_{BB} - \pi_{NB}$). Roughly speaking, these results say that *it is more profitable to move first than to catch up.*

It is important to recall that the previous results hold when assuming that one platform's buyers cannot interact with the other platform's sellers even when platforms are compliant on opposite sides ($\mu = 0$). By continuity, they also hold for very small values of μ. However, for larger values of μ, compliance may also affect profits by making the platforms interoperable. For instance, by moving from (N, S) to (B, S), platform 1 allows its buyers to interact not only with the unattached sellers but also with the rival platform's sellers. By setting $\mu = 0$, we considered only the first effect. As shown in Finding 2, interoperability tends to relax the intensity of price competition between the platforms, which suggests that the second effect further increases a platform's incentive to become compliant. However, the complexity of the model does not allow us to prove this conjecture in general.

4.2 Compliance and interoperability

To assess the internal impacts of compliance, that is, those that operate through interoperability, we mute the external impacts by setting $B_u = S_u = 0$, so that there are no unattached users. In this extreme case, compliance only affects profits in the subgames in which platforms are compliant on opposite sides, that is, in (B, S), (A, B), (A, S), their mirror images, and (A, A). All the other subgames lead to the same price equilibrium as in the benchmark (N, N). This implies that (N, N) *is always a Nash equilibrium in the absence of unattached users*, as no unilateral deviation is strictly profitable. Recall that the platforms' equilibrium profit in that subgame is $\pi_{NN} = \frac{1}{2}(\tau_b + \tau_s - \alpha_s - \alpha_b)$. We also showed in the previous section that when both platforms are fully compliant, they achieve a larger equilibrium profit (irrespective of the values of B_u and S_u):

$$\pi_{AA} = \pi_{NN} + \tfrac{1}{2}\mu(\alpha_b + \alpha_s).$$

In the opposite-sides subgames other than (A, A), we compute the equilibrium profits as:

$$\begin{aligned}
\pi_{BS} &= \pi_{AS} = \pi_{BA} = \pi_{NN} + \tfrac{1}{2}\mu Y_b, \\
\pi_{SB} &= \pi_{AB} = \pi_{SA} = \pi_{NN} + \tfrac{1}{2}\mu Y_s,
\end{aligned}$$

where Y_s and Y_b are defined in Appendix A. The first line collects the subgames in which platform 1's buyers reach platform 2's sellers through the cross-platform channel, and the second line the subgames in which platform 1's sellers reach platform 2's buyers.

The number of parameters in the model prevents us from ascertaining the sign of Y_b and Y_s, nor their ranking with respect to $(\alpha_b + \alpha_s)$. However, we can show that $Y_b + Y_s > 0$, so that at least one of the two variables is strictly positive. Suppose that $Y_s > 0$. A platform whose rival is compliant on the buyer side then earns strictly more by becoming compliant on the seller side, a deviation that makes the pair interoperable. This destabilises subgames (B, B), (N, B), and (B, N). The argument for $Y_b > 0$ is symmetric and destabilises (S,S), (N,S), and (S,N). Moreover, (N,A) and (A,N) are never equilibria for $\mu > 0$, because the non-compliant platform gains $\frac{1}{2}\mu(\alpha_b+\alpha_s) > 0$ by moving to full compliance. Throughout the parameter ranges covered by our numerical simulations, Y_b and Y_s are both strictly positive. When both are positive, no subgame in which a platform complies without generating interoperability survives as an equilibrium, while the benchmark (N,N) always remains one, since no unilateral deviation from it creates interoperability. The remaining situations can be part of a subgame-perfect equilibrium under the following conditions:

- (B,S) and (S,B) if $Y_b, Y_s \geq 0$;
- (A,B) and (B,A) if $Y_s \geq 0$ and $Y_b \geq \max\{\alpha_b+\alpha_s, Y_s\}$;
- (A,S) and (S,A) if $Y_b \geq 0$ and $Y_s \geq \max\{\alpha_b+\alpha_s, Y_b\}$;
- (A,A) if $\alpha_b+\alpha_s \geq \max\{Y_b, Y_s\}$.

In the numerical simulations we performed, we observe that $\alpha_b+\alpha_s \geq \max\{Y_b, Y_s\} \geq 0$. If so, the game has four coexisting equilibria: (N,N), (B,S), (S,B), and (A,A).

4.3 Summary

Taken together, our analysis shows that OTN compliance affects platform competition through two distinct channels. First, by expanding access to unattached users on a given side, compliance serves as a competitive advantage in attracting attached users, and it is generally more profitable to lead than to catch up in adopting additional compliance measures. Second, when platforms are compliant on different sides, compliance creates interoperability across proprietary platforms, which relaxes price competition and can raise platform profits relative to a closed benchmark. These mechanisms jointly determine how many sides platforms choose to make compliant, which compliance configurations can arise as Nash equilibria, and how interoperability interacts with the presence of unattached users.

In the next section, we use numerical simulations to explore how these forces manifest across different parameter constellations and assess their implications for platform profits, user surplus, and overall welfare.

5 Numerical analysis

The closed forms in Section 3 describe the equilibrium of each subgame in isolation. The Stage 1 game and the planner's problem both require comparisons across the sixteen profiles. Several of these comparisons admit closed-form resolution, and we report the most informative of them as findings below. The full comparison across the admissible parameter space is analytically unwieldy in the general case, and we evaluate it numerically. We work from the closed-form expressions of Section 3, which deliver the equilibrium of any profile upon substitution of that profile's compliance indicators. The procedure is described in Appendix B.

The benchmark calibration is:

$$\tau_b = \tau_s = 1, \quad \alpha_b = 0.3, \quad \alpha_s = 0.4, \quad r_b = r_s = 2, \quad B_u = S_u = 0.5, \quad \mu = 0.5.$$

The mismatch costs furnish the unit of account and are normalised to one. Fees, profits and surpluses are expressed in units of the per-unit mismatch cost. The stand-alone benefits are set to two, which is large enough to sustain full coverage under every profile. We check that these parameters satisfy conditions (6) and (7), securing an interior price equilibrium in every subgame. The inequality $\alpha_s > \alpha_b$ reflects an empirical regularity in retail marketplaces and ride-hailing, namely that an additional buyer raises a seller's transactional value by more than an additional seller raises a buyer's. The unattached masses are set to half the attached unit mass on each side, capturing a regime in which the protocol provides material but not overwhelming reach. The interoperability intensity is placed at the intermediate value $\mu = 0.5$, so that the comparative-statics exercises can move it in either direction.

5.1 Stage 1 equilibrium

Our simulations reveal that in the costless model, the Stage 1 game has a unique solution in dominant strategies.

Finding 5. *A is a dominant strategy for each platform across the admissible region of the parameter space we have tested numerically. The compliance profile (A, A) is the unique pure-strategy Nash equilibrium of the Stage 1 game.*

The finding is established by an exhaustive numerical search over a five-dimensional grid of the primitives $(\alpha_b, \alpha_s, B_u, S_u, \mu)$. At every admissible point on the grid, platform 1's profit at (A, L) exceeds its profit at (K, L) for every $K \in \{N, B, S\}$ and every choice L of the rival. The inequality is strict in the interior of the region and binds only at boundaries where a cross-side coefficient or an unattached mass degenerates to zero. The grid is dense enough that the dominance result holds robustly across the admissible region, and we verify each comparison by direct evaluation of the closed-form profit expressions (8). The numerical values that support the dominance ordering at the benchmark calibration are collected row by row in Appendix C.

The mechanism is the one that determines the equilibrium fees in Section 3. Opening a side of compliance enlarges the cross-side mass that the platform's users on that side can reach. The induced increase in those users' participation raises the access fees the platform can sustain in equilibrium on both sides through the cross-side externality. The gain is present whatever the rival plays, and it compounds when both sides of the platform are opened. Full compliance therefore dominates the alternative strategies row by row in the Stage 1 matrix. The origin of the dominance is already visible in the single-platform subgames analysed earlier. A platform that complies while its rival does not draws users toward itself on the compliant side and earns more than in the closed benchmark. Compliance is a profitable unilateral deviation, and the same logic applied to the rival's choice yields (A, A) as the dominant-strategy equilibrium.

The open architecture is therefore reached without coordination between the platforms and without regulatory mandate. Each platform finds full compliance privately optimal whatever its rival's choice. Since compliance occurs without prompting, the question the costless model leaves open is how the surplus it generates is divided between platforms and users, and the remainder of the section addresses this distributional question.

5.2 Welfare comparison

We evaluate the market equilibrium against the configuration that a planner would select, taking the welfare of users $UB + US$ as the planner's objective. This is the consumer-welfare standard most commonly applied in the two-sided markets literature. It is also the standard appropriate to the open-network setting, where the policy case for open architectures rests on the protection of buyers and small sellers, with the profitability of intermediaries a secondary concern.[7]

Under user welfare, the planner and the market part company across the greater part of the parameter space. Figure 2 reports the configuration that maximises user welfare over the plane of unattached masses, with the remaining parameters held at the benchmark. Full compliance is preferred only in the upper-right region, where both unattached pools are large. Across the rest of the plane, the planner prefers one of the same-side configurations. The planner chooses (B, B) where the unattached seller pool is large relative to the buyer pool, and (S, S) in the opposite case. A narrow band of asymmetric profiles separates the two regions.

A parallel pattern arises in the plane of cross-side coefficients. Figure 3 reports the user-welfare-maximising configuration as α_b and α_s vary at the benchmark unattached pools. The planner opens the buyer side when the buyer cross-side coefficient is the larger of the two, and the seller side in the opposite case. The diagonal $\alpha_b = \alpha_s$ separates the two same-side regions, with a narrow band of asymmetric profiles along it, and single-platform profiles appear where both coefficients are large. At the benchmark pools, full compliance is not user-optimal anywhere on this slice. The pattern reinforces the message of the previous figure. Same-side compliance dominates full compliance whenever cross-

[7]The unweighted total-welfare criterion $W = \pi_1 + \pi_2 + UB + US$ affords no discrimination among configurations in the costless model. Full compliance maximises it at every point of the admissible region. When compliance is free, opening a side is a pure enlargement of the trading opportunities available to users and can only raise aggregate surplus, whatever the division of that surplus between platforms and users. The total-welfare comparison acquires content only once compliance is costly, because the cost then enters as a deadweight loss that the planner must weigh against the surplus the protocol creates, a case we leave for companion work.

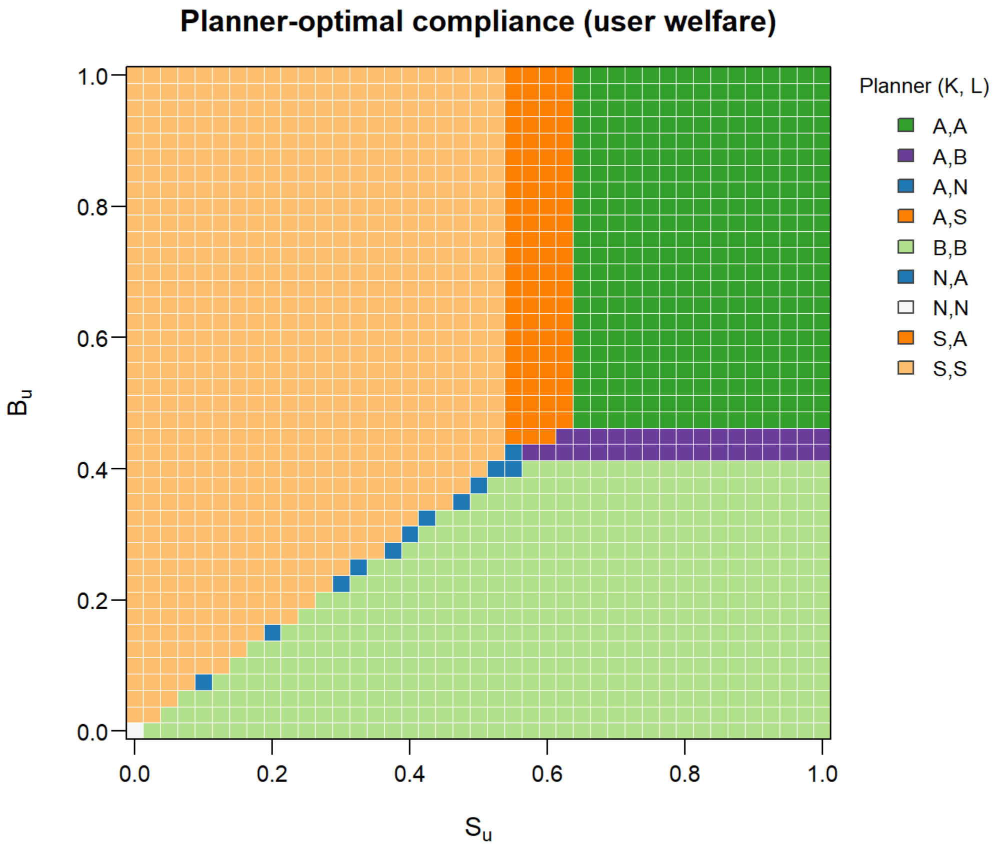

Figure 2: Compliance configuration that maximises user welfare over the plane of unattached masses, with the unattached seller pool S_u on the horizontal axis and the unattached buyer pool B_u on the vertical axis, and the remaining parameters at the benchmark. Full compliance (A, A) is preferred only where both pools are large. Same-side compliance, (B, B) where S_u is large relative to B_u and (S, S) in the opposite case, is preferred across the remainder of the plane, with a narrow band of asymmetric profiles separating the regions.

side externalities are sizeable relative to the protocol's reach.

The source of the divergence lies in the structure of equilibrium fees, as discussed in Section 3. At the closed benchmark, the fee on each side carries a discount equal to the cross-side benefit the side confers on the other, so that $b_{NN} = \tau_b - \alpha_s$ and $s_{NN} = \tau_s - \alpha_b$, the familiar property of two-sided pricing. Under full compliance with complete interoperability, $\mu = 1$, this discount vanishes and the fees rise to the Hotelling level $b_{AA} = \tau_b$ and $s_{AA} = \tau_s$. Once the two platforms are fully interoperable, the cross-side benefit becomes common to both platforms, whichever one a user happens to join, so competing for participation no longer requires conceding the cross-side subsidy. Users obtain wider reach but pay the higher fee, and the surplus the protocol creates is retained by the platforms in the form of higher margins.

Same-side compliance avoids this redistribution: in (B, B) and (S, S), no channel opens

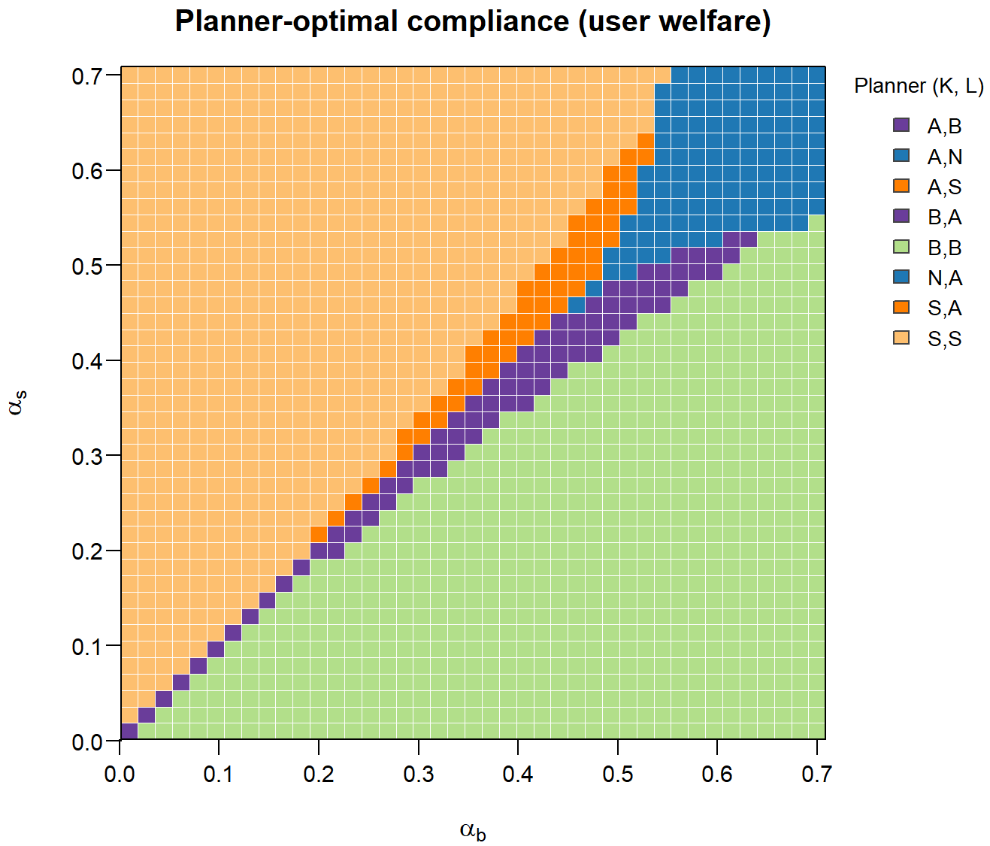

Figure 3: Compliance configuration that maximises user welfare over the plane of cross-side coefficients, with unattached pools at the benchmark $B_u = S_u = 0.5$. The planner picks (B, B) below the diagonal and (S, S) above it. A narrow band of asymmetric profiles runs along the diagonal, and single-platform profiles appear in the region where both coefficients are large. Full compliance is not user-optimal at any point of the slice.

between the platforms, the strategic mechanism that drives the fee increase at (A, A) does not apply, and equilibrium fees remain at their benchmark values, so users on the compliant side obtain the additional reach at unchanged prices. The user-welfare planner therefore favours same-side compliance whenever the gain in reach is too small to compensate users for the fee increase full interoperability would entail.

The size of this divergence has a transparent closed form: comparing the symmetric configurations, the user welfare that the planner can add over the market outcome (A, A) equals

$$\max\left\{0,\ \tfrac{1}{2}\mu\left(\alpha_b + \alpha_s\right) - \min\left\{\alpha_b S_u, \alpha_s B_u\right\}\right\},$$

up to the narrow parameter bands in which an asymmetric profile does slightly better. The first term is the network-wide fee increase that interoperability triggers, and the second is the reach gain that same-side compliance must forgo on the side it leaves closed. The divergence between the market and the planner therefore grows linearly with the interoperability intensity μ and shrinks with the value-weighted unattached pools. When μ is small, the market and the planner coincide: full compliance then delivers the pools' reach with almost no fee increase. As μ rises, the gap widens, reaching its maximum at complete interoperability, $\mu = 1$, the point at which the cross-side discount in the access fees is fully extinguished. At the benchmark calibration the expression evaluates to $0.175 - 0.15 = 0.025$, which is precisely the user-welfare advantage of (S, S) over (A, A) in Appendix C, and it rises to 0.2 at $\mu = 1$.

Figure 4 displays the same object over the plane of unattached pools at the benchmark interoperability $\mu = 0.5$. The surface is highest where the pools are thin, which is where the withdrawn fee discount dominates the added reach, and it falls to zero where both pools are large, the region in which full compliance is also what the planner would choose.

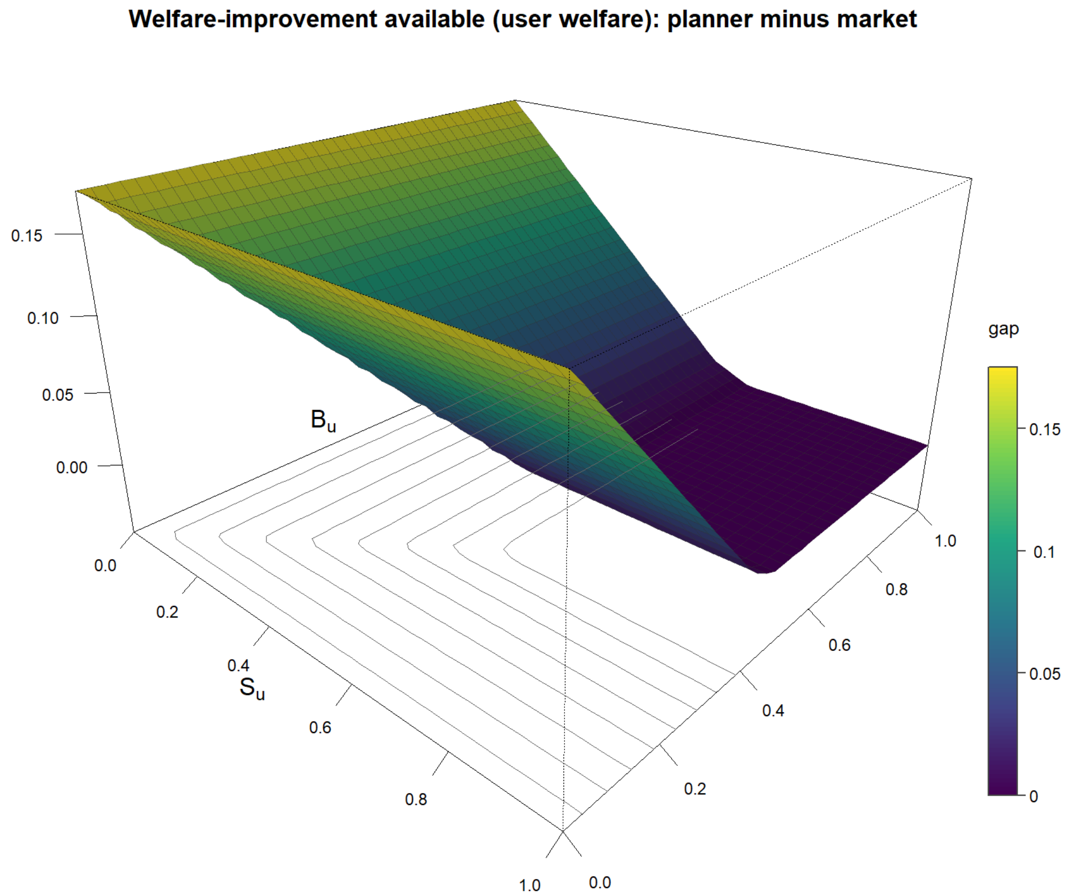

Figure 4: User welfare a planner could add over the market outcome, as a function of the unattached pools S_u and B_u, at the benchmark interoperability intensity $\mu = 0.5$. The surface is highest where the pools are thin and falls to zero where both pools are large, the region in which the market outcome (A, A) is also user-optimal.

The size of the divergence can also be read from the closed forms at the limit. Aggregating the surplus changes on the two sides, the comparison reduces to a single condition.

Finding 6. *At $\mu = 1$, the move from the closed benchmark (N, N) to full compliance (A, A) changes the buyer and seller surpluses by*

$$UB_{AA} - UB_{NN} = \alpha_b\Big(\tfrac{1}{2} + S_u\Big) - \alpha_s,$$
$$US_{AA} - US_{NN} = \alpha_s\Big(\tfrac{1}{2} + B_u\Big) - \alpha_b,$$

and the aggregate user welfare changes by

$$(UB + US)_{AA} - (UB + US)_{NN} = \alpha_b\Big(S_u - \tfrac{1}{2}\Big) + \alpha_s\Big(B_u - \tfrac{1}{2}\Big).$$

In particular, total user welfare in (A, A) exceeds total user welfare in (N, N) if and only if the right-hand side of the last expression is positive.

Finding 6 makes the trade-off explicit. On the buyer side, the reach gain $\alpha_b(\frac{1}{2} + S_u)$ has two pieces: $\alpha_b/2$ from the rival's seller mass that the open cross-platform channel makes accessible, and $\alpha_b S_u$ from the unattached seller pool. The fee loss α_s is the cross-side subsidy that platforms used to extend to buyers at (N, N) but recover under full interoperability, and the seller side is its mirror image. Aggregating, the reach gains scale with the cross-side coefficients and the unattached pools, the fee losses scale with the cross-side coefficients alone, and users gain in aggregate exactly when the pools, weighted by the cross-side coefficients, exceed the half-mass threshold on average across the two sides. For buyers alone to gain at (A, A) over (N, N), S_u must exceed $\alpha_s/\alpha_b - \frac{1}{2}$, which at our calibration requires $S_u > 0.83$. For sellers alone to gain, $B_u > \alpha_b/\alpha_s - \frac{1}{2} = 0.25$. Outside the symmetric benchmark, the sign of the aggregate effect tracks the asymmetry between the two pools weighted by the relevant cross-side coefficients.

The distributional consequence motivates the choice of welfare standard. Under full compliance, each platform's profit rises by $\frac{1}{2}\mu(\alpha_b + \alpha_s)$ over $\pi_{NN} = \frac{1}{2}(\tau_b + \tau_s - \alpha_b - \alpha_s)$. At complete interoperability, $\mu = 1$, the increase is exactly the cross-side discount the platforms recover, and profits reach $\pi_{AA} = \frac{1}{2}(\tau_b + \tau_s)$. The accompanying change in user welfare takes the ambiguous sign of the condition in Finding 6. With the pools at their benchmark values $B_u = S_u = \frac{1}{2}$ and $\mu = 1$, the condition holds with equality: buyers lose $\alpha_s - \alpha_b = 0.1$ at our calibration, sellers gain the same amount, the net change in user welfare is exactly nil, and the entire surplus that full compliance creates accrues to the platforms. More generally, platforms always gain from full compliance, while users gain only when the unattached pools are sufficiently large. The user welfare a planner could recover is accordingly largest where the pools are thin, which is also the environment in which open transaction networks have most often been deployed. The numerical values that support these statements at the benchmark calibration are collected in Appendix C.

The implication, developed in the next section, is that compliance and access pricing are complementary instruments. An intervention that opens the architecture but leaves access pricing untouched reaches the efficient outcome on a total-welfare standard yet falls short of the consumer-welfare optimum, and the shortfall is concentrated in markets where the unattached pool on one side is thin.

6 Policy implications

Our analysis suggests that an OTN is best understood as a market-based instrument of redistribution whose effects reach well beyond the enlargement of existing networks. Opening the network changes who pays for participation and who benefits from it at the same time as it widens the set of people who can trade with one another, and the value at stake moves between three constituencies: the platforms that operate the market, the buyers and sellers already attached to those platforms, and the participants whom the protocol newly brings within reach. This section walks through the mechanics of that redistribution in plain terms and closes with the limits of what the model can support.

How closed platforms price: The starting point is how closed platforms set their fees. In a two-sided market each side is worth money to the other, since an extra buyer makes the platform more attractive to sellers and an extra seller makes it more attractive to buyers. Economists call this a *cross-side externality*, the benefit that one side's participation confers on the other side, and a platform competing for users prices it in by discounting its fees below the level that product differentiation alone would support, a level we refer to as the Hotelling benchmark of fees.[8] Buyers are charged less because each buyer attracts sellers, sellers are charged less because each seller attracts buyers, and competition for exclusive access to users thereby forces platforms to hand part of the network surplus back to the very users who create it.

Why opening the network raises fees: Full interoperability removes the reason for that discount. Once both platforms comply on both sides and the protocol connects them completely, a buyer's custom stops being exclusive to the platform she joins, because the sellers she can trade with sit on either platform or outside both. Neither platform can any longer monopolise the benefit her participation creates, so neither needs to pay for it, and the equilibrium fee climbs back to the Hotelling benchmark, with each platform's profit rising by the amount of the discount it used to concede.[9] What this extra profit is made of matters for policy: expanded trade does create new value, but in the model that value accrues to users as wider reach, while the platforms' gain is exactly the old competitive subsidy, withdrawn from users who were already transacting profitably before the protocol existed. The channel that makes interoperability privately attractive to platforms is therefore the same channel that extracts surplus from incumbent users, and this is the central sense in which an OTN redistributes.

[8] In the model, the closed-benchmark fees are $b_{NN} = \tau_b - \alpha_s$ and $s_{NN} = \tau_s - \alpha_b$: each side pays the mismatch-cost parameter, which measures how differentiated the two platforms are, minus a discount equal to the value that side creates for the other.

[9] Each platform's profit rises by $\frac{1}{2}\mu(\alpha_b + \alpha_s)$ at interoperability degree μ. At complete interoperability, $\mu = 1$, the fees reach $b_{AA} = \tau_b$ and $s_{AA} = \tau_s$ and the recovery of the discount is exact.

Who gains, and when: Users receive two things from full compliance, wider reach and higher fees, and they come out ahead only if the reach is worth more than the discount they lose. Finding 6 states the exact condition, which weighs the pools of previously unattached buyers and sellers by the strength of the corresponding network effects.[10] The weighting carries the substance of the condition, because the comparison turns on how much trading value the excluded population brings within reach. The pools in the model are effective masses, meaning the number of unattached participants weighted by the quality and intensity of the trade they can sustain, so a million occasional window-shoppers may add less value than a hundred thousand active, well-integrated sellers. Opening the network benefits users where the excluded population is large in value terms, and an assessment based on registration counts alone will overstate the case wherever newly connected participants transact rarely or poorly.

The inclusion reading: The unattached pools have a natural interpretation in the settings where OTNs have actually been deployed: they are the rural and informal periphery of the market, from the small rural seller who never joined a national marketplace because listing fees and logistics requirements did not fit her scale, to the kirana store that serves its neighbourhood in cash and has no presence on any proprietary marketplace, the informal trader whose practices fit no platform's onboarding logic, and the buyer in a peripheral town whom no platform found worth acquiring. This population would like to trade and is rationed out of proprietary platforms by fees and other constraints, and when platforms comply, incumbent users come to share their trading environment with it. The model makes the financing of that inclusion explicit: incumbent, predominantly urban, users lose the fee subsidy they previously enjoyed, while the newcomers gain access without any bargaining power over the fees the platforms set, and at the benchmark calibration the platforms capture about two thirds of the total gain that full compliance generates, with the users' third shrinking to nothing as interoperability becomes complete. Inclusion is real, and it is financed by incumbent users through the withdrawn subsidy, with the platforms capturing much of the remainder.

Opening one side can serve users better: A planner, meaning a hypothetical policymaker who chooses the market configuration to maximise a chosen measure of welfare, would often stop short of full interoperability if the measure counts only the surplus of buyers and sellers, the convention known as the *consumer-welfare standard*. Compliance by both platforms on a single side delivers the reach of the unattached pool on that side while leaving fees at their benchmark level, because no channel opens between the platforms and the competitive subsidy survives, and at the benchmark calibration users do better under seller-side compliance alone than under full compliance, with both improving on the closed benchmark.[11] The market never selects this configuration on its own. Full

[10] At $\mu = 1$, aggregate user welfare rises above the closed benchmark exactly when $\alpha_b(S_u - \frac{1}{2}) + \alpha_s(B_u - \frac{1}{2}) > 0$.

[11] In units of the mismatch cost, user welfare at the benchmark calibration is 2.750 under seller-side compliance by both platforms, 2.725 under full compliance, and 2.550 at the closed benchmark.

compliance is each platform's dominant strategy, and the fee increase that harms users is part of what makes it privately attractive. The divergence between what the market does and what a user-welfare planner would choose grows with the degree of interoperability and shrinks with the value of the unattached pools, and Figures 2 and 3 map where each configuration serves users best.

Architecture and pricing are complements: The practical conclusion is that open architecture and discipline on access pricing work together, with neither able to substitute for the other. An intervention that opens the network while leaving pricing untouched achieves the total-welfare optimum in the costless model yet lets the platforms convert the competitive subsidy into margin, so users, and especially the small merchants the policy is meant to serve, can end up no better off than before, while an intervention that disciplines prices without opening the network forgoes the reach the protocol creates. Namma Yatri, the Bengaluru mobility network built on the same protocol family as ONDC, illustrates the pairing through an open architecture combined with a zero-commission model for drivers, a direct form of access-price restraint. In the language of the model, it opens the network while preventing fees from rising to the Hotelling benchmark, which is the combination under which the newly included and the incumbents can gain at the same time.

Limits of these conclusions: Two limitations bound these conclusions, and both matter for how the results should be used. Compliance is costless in the present model, whereas joining an open network in reality requires technical integration and a measure of surrendered control over data and customer relationships, and once compliance is costly, full compliance loses its automatic dominance while the total-welfare case for full interoperability must weigh the cost as a deadweight loss. Market structure is likewise held fixed, since the number of platforms does not respond to compliance choices, so the analysis is silent on entry and exit. For both reasons the model cannot explain why some large incumbents have stayed out of live OTNs, a first-order fact about deployments such as ONDC. Within the model opening is profitable, so the explanation must lie in what the model omits, whether compliance costs, asymmetries between platforms, or margins such as per-transaction commissions that an open protocol would erode. These questions are the subject of the cost and asymmetry extensions in companion work, and the conclusions above should be read as conditional on the costless, fixed-structure environment.

7 Bringing the model to the data

This section aims to identify what aspects of the model can be made amenable to empirical testing. Conditional on the type, quantity and availability of data that can be sourced, there are a number of variables, assumptions and predictions from the model that can be potentially identified and tested econometrically.

7.1 Potential data sources

The degree to which the model can be brought to the data and to which its assumptions and predictions can be tested depends crucially on the type and quality of the data that will be made available. A reasonable lower bound on the availability of data would be participation rates, which allow one to test in reduced form the change in participation incentives on both seller and buyer sides, especially if paired with administrative-level data. The testable predictions discussed in this section can be studied upon availability of ONDC's order-level transaction data, since orders record buyer-app and seller-app pairs alongside category, city, timestamp (potentially) and value, allowing one to identify compliance configurations, interoperability degree and participation counts. Another key data source would be fee panels, recovering b and s from seller-side schedules and buyer-side charges, as well as a side-specific compliance registry with dates, which apps went live as buyer apps and as seller apps, in which categories and cities, and when. The latter two sources would benefit greatly from primary data collection. The generation of such a merchant panel would achieve different goals at once: the degree of single- or multi-homing activity, pre-ONDC attachment statuses, transaction volumes, and potentially net revenues for welfare calculations. On the buyer side, surveys would greatly enhance the typical thin structure of commercial panels.

7.2 Mapping model variables to observables

- b_i, s_i: Total effective take per user per period. The access fee must be read as expected total payment equals membership charge, plus per-transaction fee times expected usage. These quantities would include delivery fees, ad spend requirements, payment charges.
- B_i, S_i: Buyer / seller counts per buyer-app / seller-app. Number of attached, transacting users.
- B_u, S_u: Effective mass of unattached participants. Possible to quantify for sellers, difficult for users. The identification of plausible proxies is a key aspect of future analysis.
- β_i, σ_i: whether app i is live on the protocol as a buyer app / seller app, by category. Binary, potentially observable from network participant registries.
- α_b, α_s: cross-side participation elasticities. Potentially recoverable: see Testable Prediction 1.
- μ: Realised cross-network order share. Potentially recoverable: see Testable Prediction 2.

Setting: For the reasons discussed above, ONDC would offer the sharpest test bed. Namma Yatri is a natural candidate as well, together with driver and rider surveys. The natural unit of observation would be category (vehicle type) $\times$ city $\times$ month, since protocol rollout was staggered across categories.

7.3 Testable assumptions

Assumption 1: Singlehoming. Buyers and sellers are active on at most one platform per period. This assumption could be tested on the seller side with data on presence across seller apps. For buyers the data requirements are more stringent as one would need some sort of account-level information showing app usage across apps. The ideal test is a direct measurement of the multi-homing rate on each side. However, it is likely that sellers multi-home extensively in Indian e-commerce, which means one would have to quantify the degree of "distance" from singlehoming.

Assumption 2: Exogenous unattached pools. The masses B_u and S_u are fixed, unresponsive to platform conduct. Having onboarding flows of previously unattached participants alongside the fee panel would allow one to test whether unattached participation responds to fee changes and subsidy campaigns.

Assumption 3: No same-side network effects. This assumption implies that more sellers means variety for buyers but competition for sellers. Upon observing incumbent seller prices, order counts, and margins around network opening, a staggered event study on incumbent sellers can identify changes as newly onboarded sellers arrive in their category and city.

7.4 Testable predictions

Prediction 1: Same-side compliance leaves fees unchanged, and full compliance raises them. The model implies that the change in fees from the closed or same-side configurations to full compliance equals $\Delta b = \mu\alpha_s$ for buyers and $\Delta s = \mu\alpha_b$ for sellers, and exactly zero under same-side compliance. If a panel of effective take rates by app $\times$ category $\times$ time is available, as well as a category rollout calendar, then the model presents a clean design, using a difference-in-differences on category $\times$ city units with the three configurations as treatment, and the same-side arm is a theory-driven placebo. The three arms (same-side only, opposite sides, both sides) would predict a null in the first arm and positive magnitudes in the third. To avoid interference, treatment should be defined at the market level (category - city) rather than at the app level.

Prediction 2: Over-identifying ratio restriction. The model implies that the ratio of fee changes $\Delta b/\Delta s = \alpha_s/\alpha_b$. But α_s, α_b are separately recoverable as cross-side elasticities from participation responses to counts on the other side. Using the fee panel from Prediction 1, together with app-level participation counts on both sides at sufficient frequency to estimate cross-side elasticities, one obtains two independent estimates of the same ratio, which must agree. This is an over-identification test. The ratio is estimated both from the fee-change ratio, as well as from a participation system using the demand functions in equation (3). The comparison can then be done with GMM using both moment conditions and the J-overid test.

Prediction 3: Linearity of fees vs. μ. The interoperability μ between the two proprietary platforms has a specific linear relationship with fees. Fee increases $\mu\alpha_s, \mu\alpha_b$ and per-platform profit gains $0.5\mu(\alpha_s + \alpha_b)$ are increasing and linear in μ. At $\mu = 1$, cross-side discounts vanish entirely and fees reach the Hotelling level $b = \tau_b, s = \tau_s$. This exercise requires order-level information with buyer-app / seller-app pairs, by category and city. Then μ can be identified as the realised cross-network order share, and fee changes can be regressed on it, testing linearity (non-parametric, or a t-test on the coefficient of a squared μ) and imposing the 0-intercept restriction. There needs to be enough cross-category variation (product attributes). There might be concerns of μ including individual choices (awareness, search behaviour), so appropriate instruments should be sought. Possible candidates could be logistics-based or network coverage.

Prediction 4: Share-neutrality under symmetric compliance. When both platforms make identical compliance choices, participation is entirely unaffected by B_u and S_u, and shares stay at 0.5 regardless of the unattached pools. Compliance on the same side by both platforms neutralises itself. Upon building a reliable proxy for the unattached pool size (for sellers, for example), a regression of share changes on pool size, interacted with symmetric versus asymmetric configuration, tests the model prediction of a zero coefficient in the symmetric arm. This test has a natural interpretation, but risks conflating a true null with an underpowered test.

Prediction 5: Adoption hazard decreases in opposing platform compliance. Finding 3 implies that the incentive to comply is largest when the other platform is closed and smallest when the other platform is fully open. Upon availability of onboarding dates for buyer and seller apps, by category, a measure of time-varying compliance status for the competing platform, and some app characteristics, one can estimate duration models and hazard rates for time-to-adoption. Two tests can be done, the first showing two opposite signs for the duration model, and the second showing two statistically different impacts using ONDC introduction as treatment, using the hazard DiD method proposed by Deaner and Ku (2024).

Prediction 6: Using a prediction rejection as a structural lower bound. One can use the observed rejection of a model prediction: the fact that large incumbents stayed out of ONDC partly contradicts Finding 5, which is likely due to specific functional forms/assumptions. One can treat the rejection as a measurement: the model's predicted profit gain from compliance, $1/2\mu(\alpha_b + \alpha_s)$ plus the external-reach terms, are a structural lower bound on the compliance cost for each non-adopter. Estimation of such a bound and comparisons across firm types can potentially allow one to discriminate among the candidate explanations that might lie behind the rejection, such as integration cost, surrendered data and customer control, and commission margins an open protocol would erode.

8 Conclusion

This paper has studied how Open Transaction Networks reshape competition between two-sided platforms. We extended the canonical two-sided singlehoming framework of Armstrong (2006) by allowing each platform to adopt an open transaction protocol on the buyer side, the seller side, both sides, or neither side. Compliance has an external effect, granting a platform's users access to pools of unattached trading partners, and an internal effect, making rival platforms partially interoperable when they are compliant on opposite sides.

The costless model overturns the standard underinvestment logic of the compatibility literature: openness emerges endogenously, with full compliance a dominant strategy for each platform and the fully open configuration the unique pure-strategy Nash equilibrium of the compliance game. The private attractiveness of openness, however, has a redistributive source. Full interoperability extinguishes the cross-side discounts that platforms concede when competing for exclusive access to users, so equilibrium fees rise toward the Hotelling level and platforms recapture the subsidy they previously extended. The welfare assessment therefore depends on the standard applied: total welfare is maximised under full compliance, while user welfare rises relative to the closed benchmark only when the unattached pools, weighted by the cross-side coefficients, are sufficiently large. Otherwise a user-welfare planner prefers same-side compliance, which widens users' reach without the accompanying rise in fees.

Several extensions merit attention. Compliance is costless in the present model. Introducing compliance costs turns the adoption decision into a genuine trade-off, gives content to the total-welfare comparison, and is the subject of companion work. The model also treats the unattached pools as exogenous, holds the duopoly market structure fixed, and restricts platforms to access fees. Endogenising participation in the protocol layer, allowing entry, and admitting per-transaction fees are natural next steps. Finally, the framework generates predictions on fee structures, participation, and the distribution of surplus that can, in principle, be confronted with data from deployed networks such as ONDC and Namma Yatri.

References

Alonso, C., T. Bhojwani, E. Hanedar, D. Prihardini, G. Uña, and K. Zhabska (2023). Stacking up the benefits: Lessons from india's digital journey. IMF Working Paper WP/23/78, International Monetary Fund, Washington, DC.

Armstrong, M. (2006). Competition in two-sided markets. *RAND Journal of Economics 37*(3), 668–691.

Baldwin, C. Y. and C. J. Woodard (2009). The architecture of platforms: a unified view.

In A. Gawer (Ed.), *Platforms, Markets and Innovation*, pp. 19–44. Cheltenham, UK: Edward Elgar.

Belleflamme, P. and M. Peitz (2021). *The Economics of Platforms: Concepts and Strategy.* Cambridge: Cambridge University Press.

Belleflamme, P., M. Peitz, and E. Toulemonde (2022). The tension between market shares and profit under platform competition. *International Journal of Industrial Organization 81*, 102807.

Belleflamme, P. and E. Toulemonde (2009). Negative intra-group externalities in two-sided markets. *International Economic Review 50*(1), 245–272.

Boudreau, K. (2010). Open platform strategies and innovation: granting access versus devolving control. *Management Science 56*(10), 1849–1872.

Caillaud, B. and B. Jullien (2003). Chicken & egg: competition among intermediation service providers. *RAND Journal of Economics 34*(2), 309–328.

Crémer, J., Y.-A. de Montjoye, and H. Schweitzer (2019). Competition policy for the digital era. Technical report, European Commission, Brussels.

Deaner, B. and H. Ku (2024). Causal duration analysis with diff-in-diff. Technical report, Technical report, Working Paper.

D'Silva, D., Z. Filková, F. Packer, and S. Tiwari (2019). The design of digital financial infrastructure: Lessons from india. BIS Papers 106, Bank for International Settlements, Basel.

Ekmekci, M., A. White, and L. Wu (2025). Platform competition and interoperability: The net fee model. *Management Science 71*(10), 8842–8864.

Evans, D. S. (2003). Some empirical aspects of multi-sided platform industries. *Review of Network Economics 2*(3), 191–209.

Farrell, J. and G. Saloner (1985). Standardization, compatibility, and innovation. *RAND Journal of Economics 16*(1), 70–83.

Farrell, J. and G. Saloner (1992). Converters, compatibility, and the control of interfaces. *Journal of Industrial Economics 40*(1), 9–35.

Furman, J., D. Coyle, A. Fletcher, P. Marsden, and D. McAuley (2019). Unlocking digital competition: Report of the digital competition expert panel. Technical report, HM Treasury, London.

Hagiu, A. (2006). Pricing and commitment by two-sided platforms. *RAND Journal of Economics 37*(3), 720–737.

Hotelling, H. (1929). Stability in competition. *Economic Journal 39*(153), 41–57.

Katz, M. L. and C. Shapiro (1985). Network externalities, competition, and compatibility. *American Economic Review 75*(3), 424–440.

Maruyama, M. and Y. Zennyo (2015). Application compatibility and affiliation in two-sided markets. *Economics Letters 130*, 39–42.

Rasch, A. and T. Wenzel (2014). Content provision and compatibility in a platform market. *Economics Letters 124*(3), 478–481.

Reisinger, M. (2014). Two-part tariff competition between two-sided platforms. *European Economic Review 68*, 168–180.

Rochet, J.-C. and J. Tirole (2003). Platform competition in two-sided markets. *Journal of the European Economic Association 1*(4), 990–1029.

Rochet, J.-C. and J. Tirole (2006). Two-sided markets: a progress report. *RAND Journal of Economics 37*(3), 645–667.

West, J. (2003). How open is open enough? melding proprietary and open source platform strategies. *Research Policy 32*(7), 1259–1285.

World Bank (2023). G20 policy recommendations for advancing financial inclusion and productivity gains through digital public infrastructure. Technical report, Report prepared for the G20 Global Partnership for Financial Inclusion, World Bank, Washington, DC.

A Compliance choices

A.1 Proof of Finding 3

Set $\mu = 0$ and define:

$$\begin{aligned} Z_s &\equiv 6\tau_b\tau_s + \tau_s\left(\alpha_b - \alpha_s\right) - \left(2\alpha_b + \alpha_s\right)\left(\alpha_b + \alpha_s\right), \\ Z_b &\equiv 6\tau_b\tau_s - \tau_b\left(\alpha_b - \alpha_s\right) - \left(2\alpha_s + \alpha_b\right)\left(\alpha_b + \alpha_s\right), \end{aligned}$$

Note that $Z_s - Z_b = \left(\alpha_b - \alpha_s\right)\left(\tau_b - \alpha_s + \tau_s - \alpha_b\right) > 0 \iff \alpha_b > \alpha_s$.

- Suppose that platform 2 is not compliant (strategy N). We compute platform 1's incentive to be compliant (on one or two sides) rather than not. We find:

$$\begin{cases} \Delta_{B,N} \equiv \pi_{BN} - \pi_{NN} = \frac{\alpha_b S_u}{2D_0}\left(Z_s + \alpha_b\tau_s S_u\right), \\ \Delta_{S,N} \equiv \pi_{SN} - \pi_{NN} = \frac{\alpha_s B_u}{2D_0}\left(Z_b + \alpha_s\tau_b B_u\right), \\ \Delta_{A,N} = \pi_{AN} - \pi_{NN} = \Delta_{B,N} + \Delta_{S,N} + \frac{\alpha_b\alpha_s(\alpha_b+\alpha_s)B_u S_u}{2D_0}. \end{cases} \tag{11}$$

- Suppose that platform 2 is compliant on the buyer side (strategy B). We compute:

$$\begin{cases} \Delta_{B,B} \equiv \pi_{BB} - \pi_{NB} = \frac{\alpha_b S_u}{2D_0}\left(Z_s - \alpha_b \tau_s S_u\right), \\ \Delta_{S,B} \equiv \pi_{SB} - \pi_{NB} = \frac{\alpha_s B_u}{2D_0}\left(Z_b + \tau_b \alpha_s B_u - \alpha_b\left(\alpha_b + \alpha_s\right) S_u\right), \\ \Delta_{A,B} = \pi_{AB} - \pi_{NB} = \Delta_{B,B} + \Delta_{S,B} + \frac{\alpha_b \alpha_s (\alpha_b + \alpha_s) B_u S_u}{2D_0}. \end{cases} \tag{12}$$

- Suppose that platform 2 is compliant on the seller side (strategy S). We compute:

$$\begin{cases} \Delta_{B,S} \equiv \pi_{BS} - \pi_{NS} = \frac{\alpha_b S_u}{2D_0}\left(Z_s + \tau_s \alpha_b S_u - \alpha_s\left(\alpha_b + \alpha_s\right) B_u\right), \\ \Delta_{S,S} \equiv \pi_{SS} - \pi_{NS} = \frac{\alpha_s B_u}{2D_0}\left(Z_b - \alpha_s \tau_b B_u\right), \\ \Delta_{A,S} = \pi_{AS} - \pi_{NS} = \Delta_{B,S} + \Delta_{S,S} + \frac{\alpha_b \alpha_s (\alpha_b + \alpha_s) B_u S_u}{2D_0}. \end{cases} \tag{13}$$

- Suppose that platform 2 is compliant on both sides (strategy A). We compute:

$$\begin{cases} \Delta_{B,A} \equiv \pi_{BA} - \pi_{NA} = \frac{\alpha_b S_u}{2D_0}\left(Z_s - \tau_s \alpha_b S_u - \alpha_s\left(\alpha_b + \alpha_s\right) B_u\right), \\ \Delta_{S,A} \equiv \pi_{SA} - \pi_{NA} = \frac{\alpha_s B_u}{2D_0}\left(Z_b - \alpha_s \tau_b B_u - \alpha_b\left(\alpha_b + \alpha_s\right) S_u\right), \\ \Delta_{A,A} = \pi_{AA} - \pi_{NA} = \Delta_{B,A} + \Delta_{S,A} + \frac{\alpha_b \alpha_s (\alpha_b + \alpha_s) B_u S_u}{2D_0}. \end{cases} \tag{14}$$

- Comparing the first lines of expressions (11) to (14), we get:

$$\begin{aligned} \Delta_{B,N} - \Delta_{B,B} &= \Delta_{B,S} - \Delta_{B,A} = \frac{\alpha_b^2 \tau_s}{D_0} S_u^2 > 0, \\ \Delta_{B,N} - \Delta_{B,S} &= \Delta_{B,B} - \Delta_{B,A} = \frac{\alpha_b \alpha_s \left(\alpha_b + \alpha_s\right)}{2D_0} B_u S_u > 0, \\ \Delta_{B,N} - \Delta_{B,A} &= \alpha_b \frac{\alpha_s^2 B_u + \alpha_b \alpha_s B_u + 2\alpha_b \tau_s S_u}{2D_0} S_u > 0. \end{aligned}$$

  This implies that $\Delta_{B,N} > \Delta_{B,B}, \Delta_{B,S} > \Delta_{B,A}$. In words the incentive to be compliant on the buyer side rather than not compliant is the largest when the other platform is not compliant ($\Delta_{B,N} = \pi_{BN} - \pi_{NN}$) and the lowest when the other platform is compliant on both sides ($\Delta_{B,A} = \pi_{BA} - \pi_{NA}$).

- Comparing, in the same way, the second, third, and fourth lines of the four expressions, we can show that the same ranking of incentives stands for the other strategies of platform 2, which proves Finding 3.

A.2 Proof of Finding 4

Set $\mu = 0$. Computing a platform's incentive to be compliant on two sides instead of one, we find:

$$\begin{aligned} \pi_{AS} - \pi_{SS} &= \frac{\alpha_b S_u}{2D_0}\left(Z_s + \tau_s \alpha_b S_u\right) > \pi_{AA} - \pi_{SA} = \frac{\alpha_b S_u}{2D_0}\left(Z_s - \tau_s \alpha_b S_u\right) \\ \pi_{AB} - \pi_{BB} &= \frac{\alpha_s B_u}{2D_0}\left(Z_b + \tau_b \alpha_s B_u\right) > \pi_{AA} - \pi_{BA} = \frac{\alpha_s B_u}{2D_0}\left(Z_b - \tau_b \alpha_s B_u\right) \end{aligned}$$

Comparing the previous expressions with (11), (12), and (13), we observe that $\pi_{AA} - \pi_{BA} = \pi_{SS} - \pi_{NS}$, $\pi_{AA} - \pi_{SA} = \pi_{BB} - \pi_{NB}$, $\pi_{AS} - \pi_{SS} = \pi_{BN} - \pi_{NN}$, and $\pi_{AB} - \pi_{BB} = \pi_{SN} - \pi_{NN}$, which proves Finding 4.

A.3 Definition of (Y_b, Y_s)

Let $B_u = 0, S_u = 0$. We compute:

$$\begin{aligned}\pi_{BS} &= \pi_{AS} = \pi_{BA} = \pi_{NN} + \tfrac{1}{2}\mu Y_b,\\ \pi_{SB} &= \pi_{AB} = \pi_{SA} = \pi_{NN} + \frac{1}{2}\mu Y_s,\end{aligned}$$

with

$$\begin{aligned}Y_b &= \frac{1}{D}\Big[-2\alpha_b(\alpha_b+\alpha_s)^2\mu^2 + (\alpha_b(9\alpha_s+8\alpha_b)(\alpha_b+\alpha_s) + \alpha_s(2\alpha_s\tau_b + \tau_s\alpha_b - \alpha_b\tau_b))\mu\\ &\quad + (6\tau_b\tau_s(\alpha_s+5\alpha_b) - 2\alpha_b(5\alpha_s+4\alpha_b)(\alpha_b+\alpha_s) + 2(\alpha_b-\alpha_s)(\alpha_b\tau_s+\tau_b\alpha_s))\Big],\\ Y_s &= \frac{1}{D}\Big[-2\alpha_s(\alpha_b+\alpha_s)^2\mu^2 + (\alpha_s(9\alpha_b+8\alpha_s)(\alpha_b+\alpha_s) + \alpha_b(2\alpha_b\tau_s + \tau_b\alpha_s - \alpha_s\tau_s))\mu\\ &\quad + (6\tau_b\tau_s(\alpha_b+5\alpha_s) - 2\alpha_s(5\alpha_b+4\alpha_s)(\alpha_b+\alpha_s) - 2(\alpha_b-\alpha_s)(\alpha_b\tau_s+\tau_b\alpha_s))\Big].\end{aligned}$$

We compute:

$$Y_b + Y_s = \frac{1}{D}\Big[2\mu(\alpha_b+\alpha_s)\Big((4-\mu)(\alpha_b+\alpha_s)^2 + \alpha_b\alpha_s\Big) + 2\mu\left(\alpha_s^2\tau_b + \alpha_b^2\tau_s\right) + 4(\alpha_b+\alpha_s)D_0\Big] > 0.$$

B Numerical procedure

The numerical results of Section 5 are computed from the closed-form expressions of Section 3. For a given vector of primitives, the equilibrium of each of the sixteen compliance profiles is obtained by substituting the profile's compliance indicators $(\beta_1, \sigma_1, \beta_2, \sigma_2)$ into the equilibrium fee, participation, and profit expressions, and the profile is admitted only if it satisfies the second-order condition (6), the interiority conditions (7), and delivers interior participations $0 < B_i, S_i < 1$ on both sides at the resulting fees. The Stage 1 equilibrium at the parameter point is the pure-strategy Nash equilibrium of the 4×4 matrix of profits so obtained, identified by verifying, for each profile, that neither platform can raise its profit through a unilateral change of compliance among the admitted profiles, and the planner's configuration is the admitted profile that maximises the chosen welfare objective at the same point. Because the set of compliance profiles is finite and the Stage 2 equilibrium is available in closed form, the computation requires no numerical optimisation beyond the evaluation of the per-profile closed forms. The comparative-statics figures are produced by repeating this calculation on a uniform grid of 41×41 points over the two parameters that are varied, with the remaining primitives held at the benchmark. At each node we record the Nash configuration, the planner's configuration and the difference in user welfare between them. The continuous surfaces are obtained by linear interpolation of the gridded values.

C Equilibrium objects at the benchmark calibration

Table 1 reports the equilibrium fees, participations, profits, and user surpluses across the ten distinct compliance subgames at the benchmark calibration $\tau_b = \tau_s = 1$, $\alpha_b = 0.3$, $\alpha_s = 0.4$, $r_b = r_s = 2$, $B_u = S_u = 0.5$, and $\mu = 0.5$. The values are computed from the closed-form expressions of Section 3 and they anchor the qualitative claims made in Section 5. The symmetric classes are reported above the asymmetric profiles. The columns are platform 1's buyer-side fee, platform 2's buyer-side fee, platform 1's seller-side fee, platform 2's seller-side fee, platform 1's buyer share, platform 1's seller share, the two per-platform profits, the total buyer surplus, the total seller surplus and total welfare.

Table 1: Equilibrium objects across the ten distinct compliance subgames at the benchmark calibration. Fees and participations are reported per platform. Profits are per platform. User surpluses and welfare are reported over the full mass on each side.

(K,L)	b_1	b_2	s_1	s_2	B_1	S_1	π_1	π_2	UB	US	W
(N,N)	0.600	0.600	0.700	0.700	0.500	0.500	0.650	0.650	1.300	1.250	3.850
(B,B)	0.600	0.600	0.700	0.700	0.500	0.500	0.650	0.650	1.450	1.250	4.000
(S,S)	0.600	0.600	0.700	0.700	0.500	0.500	0.650	0.650	1.300	1.450	4.050
(A,A)	0.800	0.800	0.850	0.850	0.500	0.500	0.825	0.825	1.325	1.400	4.375
(B,N)	0.649	0.551	0.702	0.698	0.528	0.510	0.701	0.602	1.376	1.250	3.929
(S,N)	0.597	0.603	0.768	0.632	0.514	0.538	0.720	0.585	1.300	1.351	3.957
(A,N)	0.647	0.553	0.769	0.631	0.542	0.547	0.772	0.538	1.377	1.352	4.040
(S,B)	0.623	0.777	0.874	0.676	0.474	0.544	0.771	0.717	1.316	1.330	4.133
(A,B)	0.672	0.728	0.875	0.675	0.501	0.550	0.819	0.666	1.391	1.327	4.204
(A,S)	0.775	0.625	0.743	0.807	0.535	0.492	0.781	0.700	1.314	1.429	4.224

The table makes several of the body's qualitative claims concrete at the benchmark. The Stage 1 dominance result is visible row by row: holding the rival's strategy fixed in the columns labelled by L, platform 1's profit at (A, L) exceeds its profit at the corresponding (K, L) for every $K \in \{N, B, S\}$. For instance, at $L = N$, $\pi_1(A, N) = 0.772$ exceeds $\pi_1(B, N) = 0.701$, $\pi_1(S, N) = 0.720$, and $\pi_1(N, N) = 0.650$. The same ordering holds across the other columns by the symmetry $\pi_2(K, L) = \pi_1(L, K)$, so platform 2 also has A as its best response to any rival action.

The fee structure that drives the welfare comparison is visible in the first four columns, where the closed-benchmark fees $b_{NN} = 0.6$ and $s_{NN} = 0.7$ carry the cross-side discount. The same-side configurations (B, B) and (S, S) leave fees unchanged from the benchmark, as expected, while (A, A) raises fees toward the Hotelling level. At the benchmark $\mu = 0.5$ the rise is partial: $b_{AA} = 0.8 = \tau_b - (1 - \mu)\alpha_s$ and $s_{AA} = 0.85$. At full interoperability $\mu = 1$ the fees would rise the full way to $\tau_b = 1$ and $\tau_s = 1$, eliminating the cross-side discount entirely. The asymmetric subgames show the compliant platform charging more on its open sides and capturing more participation.

The user-welfare comparison foreshadowed by Finding 6 is likewise consistent with the table. At the benchmark $\mu = 0.5$, the user-welfare aggregate $UB+US$ at (A, A) is 2.725, against 2.550 at (N, N), while the same-side configurations achieve 2.700 at (B, B) and 2.750 at (S, S). The numerical advantage of (S, S) over (A, A) for users at this calibration matches the partition figure in Section 5, where the user-welfare planner chooses (S, S) when the seller-side externality dominates and the unattached pools are at their benchmark values. At full interoperability $\mu = 1$, the comparison sharpens: $UB+US$ at (A, A) falls to exactly 2.550, matching (N, N) at this calibration, while the platforms' profits rise. The closed-form expressions in Finding 6 pin down the limit values analytically. The values in the table are the comparative-statics evaluation of the same expressions at $\mu = 0.5$.

Finally, total welfare is maximised at (A, A), with the values in the last column peaking at $W_{AA} = 4.375$. The total-welfare planner agrees with the market, but the distribution of the gain is uneven: of the 0.525 welfare gain over (N, N), platforms capture 0.350 and users capture the remaining 0.175. Under full interoperability, $\mu = 1$, the share captured by users falls to zero at this calibration.