

The Economics of Platforms in a Walrasian Framework

Anil K. Jain* and Robert M. Townsend†

August 2020

Abstract We present a tractable model of platform competition in a general equilibrium setting. We endogenize the size, number, and type of each platform, while allowing for different user types in utility and impact on platform costs. The model is applicable to the recent growth in digital currency platforms. The economy is Pareto efficient because platforms internalize the network effects of adding more or different types of users by offering type-specific contracts that state both the number and composition of platform users. Using the Walrasian equilibrium concept, the sum of type-specific fees paid cover platform costs. Given the Pareto efficiency of our environment, we argue against the presumption that platforms with externalities need be regulated.

JEL codes: D50, D62

Keywords: Two-sided markets, First and second welfare theorems, Externalities

We would like to thank the anonymous reviewer and the editor for their careful reading of our manuscript and their suggestions. We would also like to thank Glen Weyl, Marc Rysman, and seminar participants at the the Bank for International Settlements, Boston Federal Reserve, the International Industrial Organization Conference, and Princeton University for helpful comments. Anil Jain is grateful for financial support from the Consortium on Financial Systems and Poverty (CFSP). Robert Townsend is grateful for financial support from the Centre for Economic Policy Research (CEPR) and the Department for International Development (DFID) (MRG002_1255), the Eunice Kennedy Shriver National Institute of Child Health and Human Development (NICHD) (R01 HD027638), and the Consortium on Financial Systems and Poverty at the University of Chicago, funded by Bill Melinda Gates Foundation (51935). Townsend consults with the Federal Reserve Banks of New York and Boston. The findings and conclusions in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, the views of any other person associated with the Federal Reserve System or the views of the funders. All mistakes are our own.

*Federal Reserve Board of Governors, Washington D.C., USA E-mail: anil.k.jain@frb.gov · †Massachusetts Institute of Technology, Department of Economics, Cambridge, M.A., USA E-mail: rtownsen@mit.edu

1 Introduction

We are interested in economic platforms that inherently depend on attracting multiple different types of users. For instance, the quality or usefulness of a payment platform (such as digital currencies, credit cards, and mobile payments) depend on which merchants accept the payment and which consumers use the payment. Each party cares about the other. A mobile phone network is attractive only if it allows a user to message her contacts on the same technological platform. A dark pool for the trading of financial instruments needs to attract both buyers and sellers in somewhat proportionate numbers if it is to allow trade and coexist with other public exchanges. Even traditional financial intermediaries can be thought of in this way, as they intermediate between savers and borrowers, transforming the risk, liquidity, and maturity structure of funds.

We ask, in these types of markets with multiple competing platforms how one defines a Walrasian equilibrium. Does it typically exist, or are there inherent problems? If an equilibrium exists, is it allocatively efficient or is there a case for the regulation of prices? Finally, what is the relationship between competitive equilibria and the distribution of welfare; specifically, does one party on the platform have an inherent advantage?

Platform competition has rightly attracted significant academic and regulatory attention, particularly *the interchange fee*. The interchange fee is a charge for the acquiring bank (the bank that processes a credit card payment on behalf of the merchant) levied by the issuing bank (the bank that issues a consumer's credit card) to balance the credit card's costs between the merchant and the consumer's bank. Consumers' utility and merchants' profits from using or accepting a credit card depend on the number of users of both types, as well as their respective costs via the interchange fee. In this environment, because a user's utility or profit depends on the composition of the card's users, does this dependence cause a network externality? Can the interchange fee correct this network externality? Would a market-determined interchange fee require regulation? Finally, how does the size of the interchange fee distribute costs between consumers and merchants?

These questions have been asked by the academic and policy literature but only partially addressed. "The large volume of theoretical literature on interchange fees has arisen for the simplest of reasons: understanding their termination and effect is intellectually challenging", according to Schmalensee and Evans [2005]. Baxter [1983]'s seminal work, seems to be the first paper to model the multi-sided nature of payment systems: "In the case of transactional services, although consumer P's marginal valuation of the additional use of a particular payment mechanism may differ markedly from consumer M's marginal valuation, these valuations cannot be independent of one another." Baxter models these interdependencies across platform users in a competitive framework. In more recent work, Rochet and Tirole [2003], Schmalensee and Evans [2005], Armstrong [2006], Hagiu [2006], Rochet and Tirole [2006], Rysman [2009], Weyl [2010], Weyl and White [2016] argue that platforms are unable to fully internalize the merchant's marginal utility gain from an extra consumer, which leads to an unpriced "externality", and consequently, a market inefficiency.

Though the credit card payment system links directly to policy issues at stake, it also conjures up the image of imperfect competition, as the credit card industry is relatively concentrated with the leading

companies being Visa, MasterCard, American Express, and Discover. However, crypto currencies are entering, and beginning to compete. There are number of crypto currencies that feature payments, including not only Bitcoin, but also Dogecoin, Litecoin, Monero, Ripple, Stellar, and Zcash.

This inevitably raises the question of whether there is scope for the existence of multiple coins, or in the language of this paper, can there be multiple payment platforms co-existing in equilibrium. This competitive battle is being waged against the key constraint—the problem of scaling up. Bitcoin’s miners are validators using a proof-of-work protocol which consumes significant electricity, yet has limited capacity and slow transaction speed—an estimated 7 transactions per second for bitcoin (Croman et al. [2016]). But new entrants with alternative protocols are faster and cheaper, for example, Stellar with its Federated Byzantine Agreement; or Algorand’s with its proof-of-stake protocol. Visa is typically held up as the “gold standard” on these dimensions with a stated capacity to handle more than 24,000 transactions per second. No doubt facing competitive pressures, Facebook, along with a consortium of financial intermediaries, have announced a new digital currency for payments. In short, the payments industry has become much more competitive. Notably, these different currencies have in-built product differentiation across them and these various currencies typically appeal to different customer bases—both in characteristics of the platform’s users and the relative size of the platform.

Likewise, the financial intermediation industry has become more competitive, with financial technology firms entering with peer-to-peer (P2P) and business-to-business (B2B) platforms competing against more traditional commercial banks (Frost et al. [2019]). More broadly, competition in financial intermediation has been noted in other markets. For instance, there were more than 40 different platforms for trading listed securities available to traders in 2008 (O’Hara and Ye [2011]). Our paper follows recent research into financial structure and efficiency (for example, Lin et al. [2020]). Our paper is related to a theory literature on efficiency with a given segmented market structure, except that we turn the question around and ask what segmented structure of markets will support efficiency (for example, Koutsougeras [2020]). Combining these two is a literature on endogenous market segmentation and inequality (for example, Yang [2020]).

Given this background, with credit cards, coins and P2P interfaces as examples of platforms, and in contrast to the recent literature on two-sided markets, we return to the competitive framework first introduced by Baxter. We extend the original literature in two key ways: (i) We use tools from general equilibrium theory to model platform competition, and (ii) we allow platforms to offer bundles that detail the composition of a platform’s users—that is, we are clear about the platform characteristics, the commodity point, as it were. Through the use of this modified contract, we show that the prices for the platform membership, varying by type of user, overcome the inherent externality, in a similar manner as that suggested by Arrow [1969]. That is, the competitive equilibrium is Pareto optimal, and the usual first and second welfare theorems hold.

Our paper has four main results: First, building on Prescott and Townsend [2006], who analyzed firms as clubs in general equilibrium, we provide a framework that shows that platform contracts and competition among platforms can internalize the previously described externality. To be consistent with the presumption of a competitive market, we make the key assumption that platforms experience some form of decreasing returns. In reality, payment platforms face the problem of scale—the difficulty of processing many transactions in a limited amount of time. The validation protocols of distributed ledgers

require the ledgers to be synchronized in real-time, meaning every node is connected to every other, so the complexity of synchronization increase combinatorially with the number of nodes. That is related to the size of the blocks for underlying transactions to be validated and hence to speed. See Mallet [2009] and Townsend [Forthcoming]. In a different context, Altinkiliç and Hansen [2000] and O’Hara and Ye [2011] document non-increasing returns to scale in equity underwriting and equity exchanges, respectively.

Second, and more specifically, we prove that both the first and second welfare theorems hold in our model environment; the competitive equilibrium is Pareto optimal; and, any optimal allocation of resources can be achieved by lump-sum taxes and transfers on underlying wealth. To prove these results, we model that each basic user type faces a user price for each of a (infinite) number of potential platforms, which vary in the number of own-type participants and other-type participants. In equilibrium, at given prices, the solution to these decentralized problems delivers the mix and number of participants in active platforms that each user anticipated when they chose their platforms. That is, in equilibrium each active platform is populated with user types and numbers exactly as anticipated. Multiple types of platforms can coexist simultaneously, though far fewer than the potential number one can envision.

The solution is efficient because the type-specific Walrasian prices for joining platforms, which each user takes as given, change across the many potential platforms in a way that internalizes the marginal effect of altering the composition of the platform. Put differently, each agent of each type (having tiny, negligible influence), is buying a bundle that includes the composition and number of total participants, and buying the right to interface with her own and other types in known numbers. Although the problem is decentralized and each type independently determines the platform they want to join, type-specific Walrasian prices to join a platform of a given size direct traffic, so that for each and every type, the composition of active platforms will have the membership that was purchased. In short, the commodity space is expanded to include the intrinsic externality feature of the platform and prices on that commodity space decentralize the problem.

Third, we use this framework to do some comparative statics: We characterize how the equilibrium prices for each type of user to join the platform and the composition of a platform’s users change as we alter parameters of the underlying economic environment. We make a distinction between a fundamental type of user (consumer and merchants) versus within-a-type users that differ only in wealth or preferences. In this way we can examine how the equilibrium changes as we alter different consumers’ wealth. The latter allows us to see how higher wealth for a certain type leads to more advantageous matches for that type but subsequently spills over to others’ and hence to their own utility. Moreover, our comparative statics allow us to explore how the equilibrium allocation and utility may change in response to future developments in the payment industry and consequently examine questions of interest to policymakers. For example, if the cost of building platforms falls, our model demonstrates that this could cause both larger platforms and relatively greater utility for poorer users—a potentially important implication for policymakers examining financial and payment access in less economically developed economies.

We exploit in these latter comparative static exercises the fact that changing Pareto weights is equivalent to changing wealth—that is, we use a programming problem to maximize the Pareto weighted sums of utilities, and then change the weights, and subsequently tracing out all the Pareto optimal equilibria. A given optimum requires lump-sum taxes and transfers, or equivalently, a change in the initial underlying distribution of wealth.

Fourth, we demonstrate the generality of our framework for modeling platforms. We extend the model to allow for heterogeneous agent preferences. In addition, we compare a competitive and a monopoly framework to demonstrate the marked difference in their platform allocations. In this sense, competition matters even though both modes of organization internalize the externality.

Our Walrasian framework offers a compelling model for approximating the outcome of competition among financial intermediaries, holding the type of intermediation fixed. In the limit, as the number of platforms becomes arbitrarily large, we can ignore strategic aspects; of course, this approach sets us apart from the industrial organization literature, which focuses on smaller numbers and imperfect competition. To ensure the outcome that each platform is essentially a price-taker, we do need to assume in our setting that platforms do not have ever increasing returns to scale, though again we find that realistic in the context of some of the key examples.

Our framework builds heavily on club theory and in particular, the firms as club literature. Koopmans and Beckmann [1957] discuss the problem of assigning indivisible plants to a finite number of locations and its link to more general linear assignment or programming problems. A system of rents sustains an optimal assignment in the sense that the profit from each plant-location pair can be split into an imputed rent to the plant and an imputed rent to the location. At these prices landowners and factory owners would not wish to change the mix of tenants or location. As Koopmans and Beckman point out, the key to this beyond linear programming is Gale et al. [1951]’s theorem that delivers Lagrange multipliers on constraints. Every location has a match and the firms and location are not over- or under- subscribed. A linear program ignores the intrinsic indivisibilities—the integer nature of the actual problem—yet nevertheless achieves the solution.

Utilizing two related methods developed in the firms as club literature, we likewise overcome the non-convexity in our production set (assigning individuals to platforms). First, with a large number of agents one can approximate the environment with a production set that has constant returns, that is, when the non-convexity is small relative to the size of the economy. Essentially, the production set becomes a convex cone, as in McKenzie [1959, 1981]’s formulation of general equilibrium. Second, more specifically, we use lotteries (at the aggregate level) as developed in Prescott and Townsend [2006], to assign fractions of agent types to contracts and platforms even though the individual assignments are discrete. The firms as clubs methodology is well suited for our setting because it allows us to solve for which platforms emerge in equilibrium, the size of each platform, and who is part of each platform.

The closest literature to our work on platforms is the two-sided markets’ literature. This literature considers platforms that sell to at least two different user groups, and whose utility is dependent on who else uses the platform. In contrast to our paper, in general, the two-sided markets literature uses an industrial organization, partial equilibrium framework. The main finding in Rochet and Tirole [2003, 2006], Armstrong [2006], Weyl [2010], Weyl and White [2016] is that two-sided markets can lead to market failure. In particular, in the two-sided market literature, a key concern is how the distribution of users’ fees will cover the platform’s fixed and marginal costs.

Rysman [2009]’s comprehensive overview of the empirical and theoretical work on two-sided market states that “the main result (in the two-sided market literature) is that pricing to one side of the market depends not only on the demand and costs that those consumers bring but also on how their participation

affects participation on the other side.” This statement highlights three of the main advantages of our general equilibrium framework with a Walrasian allocation mechanism: First, we show that net prices are appropriate—the indirect effect on the ‘other side’ is priced in, as the price is on the composition of the platform which all believe they have a right to buy—and outcomes are efficient. Second, we show how the equilibrium changes—the prices for joining a platform, the size of platforms, and the resulting agent utilities—as we alter the underlying wealth distribution or the cost of building a platform. For instance, in one of our examples, we show that as we increase the fixed cost of building a platform, the relative cost for joining a platform rises the most for the poorest individuals, thereby increasing inequality in the system. Third, the Walrasian framework has a vastly different allocation relative to the monopoly framework.

Weyl and White [2016] consider the general equilibrium implications of two-sided markets with imperfect competition. They provide a new solution concept—Insulated Equilibrium—to explain how platforms may induce agents to coordinate over which platforms to join. Our paper focuses on modeling perfect competition with the full observability of an agent’s type. In contrast to our paper, Weyl and White argue that there remains a potential for market failure due to an unpriced consumption externality. The key difference in our papers’ predictions arises from our differing modeling choices. Our paper’s economy is perfectly competitive, whereas Weyl and White assume an oligopolistic platform economy where each platform has market power and cannot extract the full consumer surplus. This oligopolistic competition potentially leads the platforms to charge socially inefficient prices. In our economy, the platforms are perfectly competitive and earn no rents, removing this potential source of social inefficiency.¹

Indeed, we need to emphasize the limitations of what we are doing and, specifically, what we are not doing. We do not consider agents to have any pricing power. We do not consider the problem of establishing new products or platforms in the sense of innovation and entry into an existing equilibrium outcome and the problem of changing client expectations. Relatedly, we do not discuss the historical development of platforms or consider current regulatory restrictions, including well-intended but potentially misguided regulations that may limit our ideal market design. Nor do we model oligopolistic competition, though we do allow our platforms to be configured with different compositions of customers so there is clear product differentiation (just no market power). Finally, we do not allow ever increasing economies of scale in platform size.

2 Model

There are two types of individuals, merchants (A) and consumers (B).² There is a continuum of measure one of each type. There is variation within these types – namely, there are sub-types of merchants and consumers who differ in their endowment levels, wealth. We index each agent by (T, s) where T is the type (merchant or consumer) and s is the sub-type. There are I subtypes of merchants, A (indexed by i), and J subtypes of consumers, B (indexed by j). By introducing variation in an agent’s wealth, we

¹ Within our framework we can also model a monopolist platform sector. In this model, the monopolist will maximize profits by severely restricting supply and producing a negligible mass of platforms.

² For clarity, we restrict our model to two types, but our model is sufficiently general to accommodate multiple types.

can analyze how changes in the economic environment affect both the composition of a platform and an agent's utility.³

There is a fraction $\alpha_{T,s}$ of each type T and subtype s , and there is a measure of each one of each type T , $\sum_s \alpha_{T,s} = 1 \forall T \in \{A, B\}$.⁴ Clearly, the fraction of each subtype $\alpha_{T,s}$ are arbitrary real numbers on the unit interval, not integers. Each agent has an endowment of capital, denoted by $\kappa_{T,s} > 0$. Capital will be the numeraire.

We model utility at a reduced-form level. That is, we assume agents procure utility from being matched with other agents. Although this assumption is not realistic per se, we presume the process of being matched with other agents facilitates trade over the platform, and trade gives final allocations, resulting in utility. We do not model that underlying environment explicitly—and in some ways that makes our model more general. Further, we could generalize and introduce a term for any private benefit the platform provides over and above its matching service. In short, the utilities we use are to be thought of as indirect.

We only allow non-negative integers of merchants and consumers to join a platform. The utility of a merchant, A (of any subtype s), matched with N_A merchants (including the merchant herself) and N_B consumers is:

$$U_{A,s}(N_A, N_B) = U_A(N_A, N_B) = \begin{cases} 0 & \text{if } N_B = 0 \\ \left[\left(\frac{N_B}{N_A} \right)^{\gamma_A} + N_B^{\epsilon_A} \right] & \text{else} \end{cases}$$

Note that the baseline utility of not being on any platform and being matched with none of the other type is zero.⁵ This situation is the “opt-out” or autarky option, and it is always available.

Symmetrically, for a consumer, B (of any subtype s), it is:

$$U_{B,s}(N_A, N_B) = U_B(N_A, N_B) = \begin{cases} 0 & \text{if } N_A = 0 \\ \left[\left(\frac{N_A}{N_B} \right)^{\gamma_B} + N_A^{\epsilon_B} \right] & \text{else} \end{cases}$$

Where $\{\gamma_A, \gamma_B, \epsilon_A, \epsilon_B\} \in (0, 1)^4$ are the key parameters.

This utility function exhibits two important features which Ellison and Fudenberg [2003] highlight:

- (1) **Market Impact Effects:** Each type prefers more of the other type and less of it's own

$$U_A(N_A, N_B + 1) - U_A(N_A, N_B) = \frac{[(N_B + 1)^{\gamma_A} - N_B^{\gamma_A}]}{N_A^{\gamma_A}} + [(N_B + 1)^{\epsilon_A} - N_B^{\epsilon_A}] > 0$$

$$U_A(N_A + 1, N_B) - U_A(N_A, N_B) = \left[\left(\frac{1}{N_A + 1} \right)^{\gamma_A} - \left(\frac{1}{N_A} \right)^{\gamma_A} \right] N_B^{\gamma_A} < 0$$

³ Section (6.1) extends the baseline model to allow subtypes to have different preferences.

⁴ We use this assumption for computational ease but it is straightforward to allow different measures of each type. Additionally, some insights can be drawn from varying agents' wealth endowments.

⁵ This is a natural assumption for the opt-out utility because the lower bound for N_A is one (because as soon as a merchant joins a platform, there must be at least one merchant on the platform), and, if N_A is positive, the limit of $U_A(N_A, N_B)$ as N_B goes to zero is zero ($\lim_{N_A \geq 1, N_B \rightarrow 0} \left(\frac{N_B}{N_A} \right)^{\gamma_A} + N_B^{\epsilon_A} = 0$).

Individuals will compete between agents of their own type, though they prefer more of the other type. For example, in the general merchant and consumer case, we are presuming that merchants dislike more merchants, as this situation would lead to greater competition and possibly reduce the good's price. Therefore, we are modeling a reduced form specification for competition between agents of the same type.

- (2) **Scale effects:** An individual prefers larger platforms for a given ratio
 - assume $\tau > 1$, therefore:

$$U_A(\tau N_A, \tau N_B) - U_A(N_A, N_B) = (\tau^{\epsilon_A} - 1)N_B^{\epsilon_A} > 0$$

For a given ratio of participants, individuals prefer to be on larger platforms, as such platforms provide more possibilities for trade and may promote economies of agglomeration.

Symmetrically, both effects also apply for the type B utility function.

The assumption that epsilon is greater than zero might seem to bias us in the direction of having large platforms, undercutting our results, or make the price taking assumption unrealistic, but we shall see this does not happen.

In the model, agents buy personal contracts that stipulate the number of merchants and consumers on the platform.

We denote the contract by $d_T(N_A, N_B)$, where N_A and N_B are the number of merchants and consumers, respectively, on the given platform, and T denotes the type of individual the contract is for—a dummy as it were—indicating whether it is for merchants (A) or consumers (B). Types are observed, and Type T cannot buy a contract indexed by T' . Further, one can think of any individual agent of a given type T and subtype s as allowed to join only one platform. Thus, we can create a function $x_{T,s}[d_T(N_A, N_B)] \geq 0$ such that $\sum_s x_{T,s}[d_T(N_A, N_B)] = 1$, which is an indicator (or, more generally, a probability distribution, on which in the following) for the assignment of an agent (T, s) to contract $d_T(N_A, N_B)$.

The set of all possible contracts for type A , the space in which $d_A(N_A, N_B)$ lies, is denoted as D_A and, similarly, the set of contracts for type B is denoted D_B .

The consumption set of type A, s agents can be written as the following:

$$X_{A,s} = \left\{ x_{A,s}[d_A(N_A, N_B)] \geq 0 \forall d_A \in D_A, \sum_{d_A \in D_A} x_{A,s}[d_A(N_A, N_B)] = 1, x_{A,s}[d_B(N_A, N_B)] = 0 \forall d_B \in D_B \right\}$$

The above condition states that type A, s agents can purchase any non-negative amount of contract $d_A \in D_A$, but none of the type B contracts, and in total, must purchase a total of one unit of contracts, that is, the individual must join a platform with probability one. Since, the agent's consumption set is convex, and since the agent's utility function is linear, the utility function is concave.

Symmetrically the consumption set of type B, s agents can be written as:

$$X_{B,s} = \left\{ x_{B,s}[d_B(N_A, N_B)] \geq 0 \forall d_B \in D_B, \sum_{d_B \in D_B} x_{B,s}[d_B(N_A, N_B)] = 1, x_{B,s}[d_A(N_A, N_B)] = 0 \forall d_A \in D_A \right\}$$

As individuals can join only a single platform, this constraint introduces an indivisibility into an agent's consumption space. To overcome this problem we allow individuals to purchase mixtures, or probabilities of being assigned to a platform of a certain size including the opt-out option.⁶ For example, consider an agent who buys two different contracts: the first contract assigns the agent to a platform consisting of four merchants and three consumers with probability one-third, and the second contract assigns the agent to a platform consisting of three merchants and one consumer with probability two-thirds. The deterministic case, where an agent buys a contract that matches them with a platform of size (N_A, N_B) with certainty, can be seen as a special case. We do not insist that there be mixing in a competitive equilibrium but it can happen as a special case. For instance, there can be mixing between a given platform and an opt-out contract when agents are poor because the poor agent has insufficient wealth to buy fully into a platform.

As a technical assumption, we assume there is a maximal platform of size $(\overline{N}_A, \overline{N}_B)$, and any platform up to this size can be created. Assuming there is a maximal platform size bounds the possible set of platforms and hence makes the commodity space finite. This maximal platform size is for simplicity of our proofs because we can choose \overline{N}_A and \overline{N}_B arbitrarily large such that this condition does not bind.

The commodity space is thus:

$$L = \mathbb{R}^{2(\overline{N}_A \times \overline{N}_B + 1) + 1}$$

There are contracts for every possible platform size, in turn indexed by the two types. Therefore, there are $(\overline{N}_A \times \overline{N}_B + 1)$ contracts for each type because we defined the maximal platform size to be $(\overline{N}_A, \overline{N}_B)$ and there is always the opt-out contract. Because there are two types, we multiply this number by two to calculate the total number of contracts available. Finally, there is a market for capital, as we describe in the following paragraph.⁷

All contracts $d_T(N_A, N_B)$ are priced in units of the capital good and the price for contract $d_T(N_A, N_B)$ is denoted as $p[d_T(N_A, N_B)]$ for types A and B (where $T \in \{A, B\}$).⁸

⁶ A similar modeling approach is used in Prescott and Townsend [1984], Prescott and Townsend [2005], Pawasutipaisit [2010].

⁷ For instance, assume \overline{N}_A and \overline{N}_B are equal to 2. Then for type A agents there would be five possible contracts to join a platform. There are platforms of size and composition (N_A, N_B) : (1,0) [opt-out contract], (1,1), (2,1), (1,2) and (2,2). Similarly for type B agents, the agents could join platforms of composition (N_A, N_B) : (0,1) [opt-out contract], (1,1), (2,1), (1,2) and (2,2). Finally, agents have their capital endowment, κ . Therefore, in total there are eleven contracts.

⁸ Note we could denote the price function as $p_T[d_T(N_A, N_B)]$ but since agents can only buy contracts of their specific type we can reduce notation by denoting the pricing function as $p[d_T(N_A, N_B)]$.

2.1 Agent's Problem

In summary, agent T, s takes prices $p[d_T(N_A, N_B)] \forall d_T \in D_T$ as given and solves the maximization problem:

$$\max_{x_{T,s} \in X_{T,s}} \sum_{d_T(N_A, N_B)} x_{T,s}[d_T(N_A, N_B)] U_T[d_T(N_A, N_B)] \quad (1)$$

$$\text{s.t.} \quad \sum_{d_T(N_A, N_B)} x_{T,s}[d_T(N_A, N_B)] p[d_T(N_A, N_B)] \leq \kappa_{T,s} \quad (2)$$

$$\sum_{d_T(N_A, N_B)} x_{T,s}[d_T(N_A, N_B)] = 1 \quad (3)$$

where each type of individual has an endowment of $\kappa_{T,s}$ (that is strictly positive) of capital and the price of capital is normalized to one—that is, capital is the numeraire.

Equation (1) is the agent's expected utility from the assignment problem. Equation (2) is the agent's budget constraint. Equation (3) is the agent's probability or matching constraint, which requires the agent to join a platform or opt-out.⁹

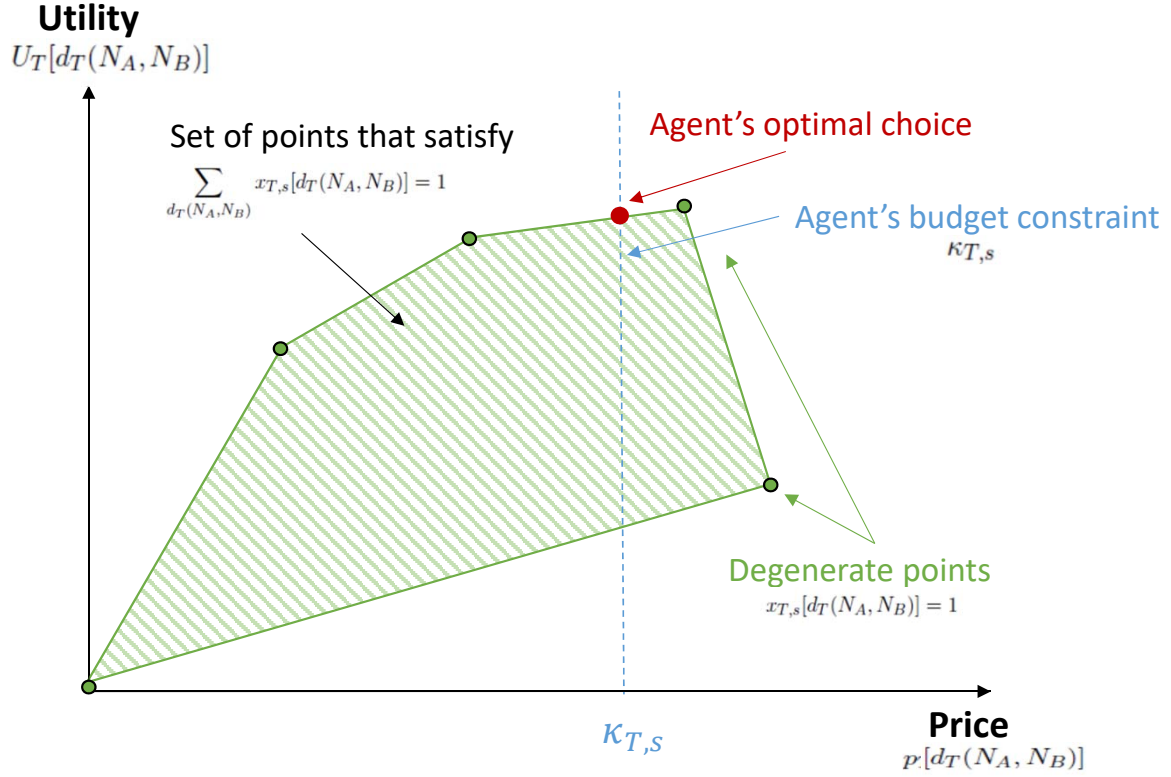
To ease exposition, figure 1 graphically represents the agent's maximization problem and optimal choice. The green dots represent the utility and price (recall that we have normalized the price of the capital good to one) for each degenerate platform choice (that is, when $x_{T,s}[d_T(N_A, N_B)] = 1$), the dashed blue line represents the agent's hypothetical budget constraint, $\kappa_{T,s}$, and the dashed green area represents the set of points that satisfy the agent's probability constraint (equation (3)). In this example, the agent's optimal choice and expected utility is represented by the red dot. As shown in the figure, as the agent has the option to buy contracts that assign the agent to a platform with some probability, the agent is able to purchase a bundle of contracts that provides higher expected utility than if the agent were able to buy only those contract that assigns the agent to a platform with certainty.

2.2 Platforms

We assume there are intermediaries or marketmakers who create platforms and sell contracts for each type to join platforms. As will be evident, there are constant returns to scale for the intermediaries, so for simplicity we can envision that just one marketmaker is needed in equilibrium. We denote $y_A[d_A(N_A, N_B)]$, as the *number of contracts* produced for type A of platform size and composition (N_A, N_B) and $y_B[d_B(N_A, N_B)]$ as the *number of contracts* produced for type B of platform size and composition (N_A, N_B) . These are counting measures and there is nothing random. Further, we denote the *number of platforms* of size and composition (N_A, N_B) as $y(N_A, N_B)$. Thus $N_A \times y(N_A, N_B)$ is the number of type A 's in total on the type of platform counted as $y(N_A, N_B)$. Similarly, $N_B \times y(N_A, N_B)$

⁹ Because an agent can join a platform that is only populated by that agent (autarky or a singleton platform), this constraint essentially requires agents to join at most one platform in equilibrium.

Fig. 1: The agent's maximization problem



is the number of type B 's in total on the type of platform counted as $y(N_A, N_B)$. In turn, each of these must equal the numbers $y_A[d_A(N_A, N_B)]$ and $y_B[d_B(N_A, N_B)]$ respectively, as defined earlier.

The intermediary must satisfy the following consistency constraint:

$$\frac{y_A[d_A(N_A, N_B)]}{N_A} = \frac{y_B[d_B(N_A, N_B)]}{N_B} = y(N_A, N_B) \quad \forall d_A \in D_A, \forall d_B \in D_B \quad (4)$$

In other words, this constraint states that the number of platforms created for type A of size and composition (N_A, N_B) , relative to the number of platforms created for type B of the same size and composition (N_A, N_B) must equal the relative number of type B to the number of type A on the platform.¹⁰ Also, all of the numbers dealing with the number of platforms are on a continuum and so do not have to take on integer values. This is because the mathematics takes into account the continuum measure of each type of mass.¹¹

¹⁰ For instance, assume that there are 0.1 platforms of size and composition $(N_A = 2, N_B = 1)$ (that is, $y(2, 1)$ equals 0.1). Then, to ensure there are the appropriate number of type As to type Bs on the platform, we require that the number of contracts for type of A for platform of composition $(2, 1)$ to be equal to 0.2 (that is, $y_A[d_A(2, 1)]$ equals 0.2), and that the number of contracts for type B for platform of composition $(2, 1)$ to be equal to 0.1 (that is, $y_B[d_B(2, 1)]$ equals 0.1).

¹¹ For example, if we multiply each type by 100, we will have larger numbers for each type but the same proportions. Now consider if 0.1 platforms are created that match three merchants and two consumers, a platform of composition $(3, 2)$.

A platform of size (N_A, N_B) requires the following amount of capital:

$$C(N_A, N_B) = \begin{cases} 0 & \text{if } N_A = 0 \text{ or } N_B = 0 \\ c_A N_A + c_B N_B + c N_A N_B + K & \text{else} \end{cases}$$

The capital requirement of a singleton or opt-out platform is normalized to zero, as it costs nothing to produce and is always available. In creating a platform, there is positive marginal costs for an extra agent on each side of the platform (captured by c_A and c_B) and for the multiple of agents on both sides (captured by the interaction term c). Additionally, there can be some positive fixed cost, K , in creating a platform. For a more flexible specification we allow c_A and c_B to be different. We assume that $c_A, c_B, c \in (0, \infty)$ and $K \in [0, \infty)$. We require c_A, c_B and c to be strictly larger than zero; this assumption ensures we can bound the size of the equilibrium platforms. In particular, with c strictly greater than zero, the cost of doubling the size of any given platform rises more than proportionally (excluding the fixed cost)—that is, there are decreasing returns to scale in this technology after some size.¹²

We denote the amount of capital input purchased by the intermediary as y_κ – this amount has to be sufficient to build the proposed platforms, as counted in $y(N_A, N_B)$. Recall again that there can be a variety of platforms. Thus we can write the intermediary's capital constraint as the following:

$$\sum_{N_A, N_B} y(N_A, N_B) [C(N_A, N_B)] \leq y_\kappa \quad (5)$$

Hence, the intermediary's production set is:

$$Y = \left\{ (y, y_A, y_B, y_\kappa) \in \mathbb{R}^{2(\overline{N_A \times N_B} + 1) + 1} \mid (4) \text{ and } (5) \text{ are satisfied} \right\}$$

It is a convex cone as in McKenzie [1959]. Note the choice objects are the $y(N_A, N_B)$'s, while the cost function enters only as a weighting coefficient. Hence there are constant returns to scale in the cost function as in constraint (5).

We explore the role of market power in platform supply and agent welfare by modeling two different environments: First, we model a price-taking intermediary, and second, we model a price-setting intermediary who has market power and who can set both the quantity and price of each platform contract.

This matching would require $0.1 \times 3 = 0.3$ merchant contracts and $0.1 \times 2 = 0.2$ consumer contracts. We could multiply this figure by the common factor of 100 to have 10 platforms each with the composition (3, 2) in total, hence with 30 type A s and 20 type B s in total. If the fraction of platforms created were 0.135 and we used 100 as the base, we would end up with 13.5 platforms, but multiplying by 1000, and we are back to integers, with 135 platforms, and so on. The point is that the counting measures y are a more general way to do the math and do not require integers. Also, any single individual, or any single platform, for that matter, has zero mass.

¹² Recall that the production of platform of a certain size and composition has constant returns to scale (for example, creating ten platforms of size (N_A, N_B) will require ten times the amount of capital as creating one platform of size (N_A, N_B) and the stated cost function shows that creating a larger platform (that is, a platform with more individuals on that platform) has decreasing returns to scale after some size (due to the strictly positive coefficient c).

2.3 Competition: price-taking intermediary.

The intermediary takes the Walrasian prices $p[d_T(N_A, N_B)] \forall d_T \in D_T, T \in \{A, B\}$ as given parametrically and maximizes profits by constructing platforms and selling type-specific matchings (as before, the price of capital is normalized to one):

$$\pi = \max_{y, y_A, y_B, y_\kappa \in Y} \sum_{N_A, N_B} \{p[d_A(N_A, N_B)] \times y_A[d_A(N_A, N_B)] + p[d_B(N_A, N_B)] \times y_B[d_B(N_A, N_B)]\} - y_\kappa \quad (6)$$

such that:

$$\frac{y_A[d_A(N_A, N_B)]}{N_A} = \frac{y_B[d_B(N_A, N_B)]}{N_B} = y(N_A, N_B) \quad \forall d_A \in D_A, \forall d_B \in D_B \quad (7)$$

$$\sum_{N_A, N_B} y(N_A, N_B)[C(N_A, N_B)] \leq y_\kappa \quad (8)$$

Equation (6) states that the intermediary maximizes the number of platforms of a given size (N_A, N_B) to produce given the prices for each position in the platform. The intermediary's profits are equal to the number of contracts the intermediary constructs multiplied by their respective price, minus the cost of the capital input.

The intermediary's Lagrange problem is:

$$\begin{aligned} L(y_A, y_B, y, y_\kappa, \mu) = & \sum_{N_A, N_B} \{p[d_A(N_A, N_B)] \times y_A[d_A(N_A, N_B)] + p[d_B(N_A, N_B)] \times y_B[d_B(N_A, N_B)]\} \\ & - y_\kappa + \sum_{N_A, N_B} \left[\mu_{N_A, N_B}^A \left(\frac{y_A[d_A(N_A, N_B)]}{N_A} - y(N_A, N_B) \right) \right] \\ & + \sum_{N_A, N_B} \left[\mu_{N_A, N_B}^B \left(\frac{y_B[d_B(N_A, N_B)]}{N_B} - y(N_A, N_B) \right) \right] \\ & + \mu^k \left(\sum_{N_A, N_B} y(N_A, N_B)[C(N_A, N_B)] - y_\kappa \right) \end{aligned}$$

where μ_{N_A, N_B}^A and μ_{N_A, N_B}^B are the Lagrange multipliers for the intermediary's consistency constraints for a platform of size (N_A, N_B) and for types A and B respectively (equation (7)). μ^k is the Lagrange multiplier for the intermediary's capital constraint (equation (8)).

Solving the intermediary's Lagrange problem gives the following first order condition for creating a platform of size $y(N_A, N_B)$:

$$C(N_A, N_B) \geq p[d_A(N_A, N_B)] * N_A + p[d_B(N_A, N_B)] * N_B \quad (9)$$

where equation (9) holds with equality if there is a positive number of active platforms of that size (N_A, N_B) in equilibrium. If equation (9) is a strict inequality then no such platform exists in equilibrium. Notice this natural condition requires that the payments received by the platform must cover all of the platform's costs. The payments come from type-specific Walrasian prices and in that sense, the interchange fee in the credit card example is emerging endogenously. Also note when the inequality in equation (9) is strict, such platforms of that type do not exist. One could raise the price marginally; however, this increase would only discourage demand and still would not cover the supply side's costs. The market for inactive platforms is thus clearing at a zero quantity with the minimum price the intermediary is willing to accept, which is greater than the maximum sum of prices the household types are willing to pay.

2.3.1 Competition: Market Clearing.

For market clearing we require the following conditions to hold

$$\sum_s \alpha_{T,s} x_{T,s} [d_T(N_A, N_B)] = y_T [d_T(N_A, N_B)] \quad \forall N_A, N_B, T \in \{A, B\} \quad (10)$$

$$\sum_{T,s} \alpha_{T,s} \kappa_{T,s} = y_\kappa \quad (11)$$

Equation (10) ensures that the (decentralized) amount of demand for each contract for each type equals the (decentralized) supply of that contract. Equation (11) states that the total endowment of capital (the supply) must equal the amount of capital used by the intermediary.

2.3.2 Competitive Equilibrium.

Let us define x as the vector of contracts bought $x_{T,s} [d_T(N_A, N_B)]$ for all subtypes (T, s) , then a competitive equilibrium in this economy is $(p, x, \{y, y_A, y_B, y_\kappa\})$ such that for given prices $p [d_T(N_A, N_B)]$:

1. The allocation $\{x_{T,s} [d_T(N_A, N_B)]\}$ solves the agent's maximization problem [that is, $x_{T,s} [d_T(N_A, N_B)]$ solves equation (1) subject to equations (2 and 3)].
2. The allocation $\{y, y_A, y_B, y_\kappa\}$ solves the platform's maximization problem [that is, $\{y, y_A, y_B, y_\kappa\}$ solves equation (6) subject to $\{y, y_A, y_B, y_\kappa\} \in Y$].
3. The market clearing conditions hold [equations (10) and (11) hold].

In equilibrium, the pricing mechanism will determine the size and number of each platform and subsequently the relative proportions of merchants and consumers on each platform.

2.4 Monopoly: price-setting intermediary.

In contrast to Section (2.3), we model the intermediary as a price-setting monopolist, who sets prices $p [d_T(N_A, N_B)] \quad \forall d_T \in D_T, T \in \{A, B\}$ and quantities $y_T (d_T(N_A, N_B)) \quad \forall d_T \in D_T, T \in \{A, B\}$ to

maximize profits subject to aggregate demand equals aggregate supply, where aggregate demand is derived from the consumer's problem.

$$\pi = \max_{p, y_A, y_B, y_\kappa \in L \times Y} \sum_{N_A, N_B} \{p[d_A(N_A, N_B)] \times y_A[d_A(N_A, N_B)] + p[d_B(N_A, N_B)] \times y_B[d_B(N_A, N_B)]\} - y_\kappa \quad (12)$$

$$\text{s.t. } \sum_s \alpha_{T,s} x_{T,s} [d_T(N_A, N_B)] \geq y_T [d_T(N_A, N_B)] \quad \forall N_A, N_B, T \in \{A, B\} \quad (13)$$

Equation (12) states that the intermediary problem maximizes how many platforms of a given size to produce and the price to charge each side of the market ($p[d_A(N_A, N_B)]$ and $p[d_B(N_A, N_B)]$) for each position in the platform; subject to, quantity supplied being less than or equal to total demand for each contract and the allocation being within the intermediary's production set, Y .¹³

2.4.1 Monopoly Equilibrium.

Then a monopoly equilibrium in this economy is $(p, x, \{y, y_A, y_B, y_\kappa\})$ such that:

1. The allocation $\{x_{T,s}[d_T(N_A, N_B)]\}$ solves the agent's maximization problem [that is, $x_{T,s}[d_T(N_A, N_B)]$ solves equation (1) subject to equations (2 and 3)].
2. The allocation $\{y, y_A, y_B, y_\kappa\}$ and prices p solves the platform's maximization problem [that is, $(p\{y, y_A, y_B, y_\kappa\})$ solves equation (12) subject to $(p\{y, y_A, y_B, y_\kappa\}) \in L \times Y$ and equation (13)].

3 Social Planner's Problem

First, we set up the social planner's problem and determine the set of all Pareto optimal contracts. We show (i) a competitive equilibrium is Pareto optimal, (ii) *any* Pareto optimal allocation can be achieved with lump-sum transfers and taxes among agents and (iii) there exists a competitive equilibrium. These results have two important implications: (i) the decentralized problem is Pareto optimal, and (ii) when solving for the competitive equilibrium, we can use the simpler social planner's problem to compute the allocation. Consequently, we can use the Lagrange multipliers to impute the competitive equilibrium prices and wealth associated with that allocation.

The social planner's welfare maximizing problem with Pareto weights $\lambda_{A,s}$ and $\lambda_{B,s}$ for types (A, s) and (B, s) , respectively, is¹⁴

¹³ Recall that the production set requires the intermediary's consistency constraint (equation 4) and capital constraint to be satisfied (equation 5).

¹⁴ Note that the social planner's problem can also be written as a maximization problem in only x , that is, the program does not need to be written in both x and y . Specifically, for all platforms, we could rewrite equation (15) in the following way, on the left hand side, as the sum of contracts for a platform of a specific size (N_A, N_B) purchased by agents of type A. This has to be in proportion to the sum of contracts purchased by agents of type B, where the proportion is determined by the respective population (N_A/N_B) on that platform, written on the right hand side (as currently written, in equation (15) these two sides have to be equal to each other because they are equal to y). For equation (16), we can use the left hand side of equation (15) to substitute in for y (for either type A or type B). For ease of exposition and consistency with the proofs in the appendix, we have retained y in our description of the problem.

$$\begin{aligned}
& \max_{x \geq 0, y \geq 0} \sum_s \lambda_{A,s} \left\{ \sum_{d_A(N_A, N_B)} \alpha_{A,s} x_{A,s} [d_A(N_A, N_B)] U_A(N_A, N_B) \right\} \\
& \quad + \sum_s \lambda_{B,s} \left\{ \sum_{d_B(N_A, N_B)} \alpha_{B,s} x_{B,s} [d_B(N_A, N_B)] U_B(N_A, N_B) \right\} \\
\text{s.t.} \quad & \sum_{d_T(N_A, N_B)} x_{T,s} [d_T(N_A, N_B)] = 1 \quad \forall T, s \tag{14} \\
& \sum_s \alpha_{T,s} x_{T,s} [d_T(N_A, N_B)] = y(N_A, N_B) \times N_T \quad \forall d_T \in D_T, \forall T \in \{A, B\} \tag{15} \\
& \sum_{N_A, N_B} y(N_A, N_B) [C(N_A, N_B)] \leq \sum_{T,s} \alpha_{T,s} \kappa_{T,s} \tag{16}
\end{aligned}$$

Equation (14) ensures that each individual is assigned to a platform, equation (15) ensures that the total purchase of contracts equals the number of contracts produced, and equation (16) ensures the total number of contracts produced is resource feasible.

3.1 Dual

The Pareto problem can also be written in terms of its dual equivalent:

$$\begin{aligned}
& \min_p \sum_{T,s} (p_{T,s} + p_\kappa \alpha_{T,s} \kappa_{T,s}) \\
& \text{s.t. } p_{T,s} + \alpha_{T,s} p [d_T(N_A, N_B)] \geq \lambda_{T,s} \alpha_{T,s} U_T(N_A, N_B) \quad \forall i, \forall T, \forall (N_A, N_B) \tag{17} \\
& p_\kappa C(N_A, N_B) - \{p [d_A(N_A, N_B)] \times N_A + p [d_B(N_A, N_B)] \times N_B\} \geq 0 \quad \forall (N_A, N_B)
\end{aligned}$$

In this formulation $p_{T,s}$, $p [d_T(N_A, N_B)]$, and p_κ are the Lagrangian multipliers associated with the participation constraint for the agent of type T, s (equation 14), the Lagrangian multiplier associated with the consistency constraint for type T for all platforms (equation 15), and the Lagrangian multiplier associated with the economy's resource constraint (equation 16), respectively.

The dual minimizes the aggregate cost of the economy (in terms of prices of each type and total capital) such that each type of agent receives a given level of Pareto weighted utility. Whereas, the primal problem maximizes the Pareto weighted expected utility of each type subject to the matching and resource constraints.

The Pareto problem is well defined in both the primal and dual form therefore, by the ‘‘strong duality property’’¹⁵ there must exist an optimal solution (p^*, x^*, y^*) such that:

¹⁵ See Bradley et al. [1977] pages 142-143 for more details.

$$\sum_{T,s} \lambda_{T,s} \left\{ \sum_{d_T(N_A, N_B)} \alpha_{T,s} x_{T,s}^* [d_T(N_A, N_B)] U_T(N_A, N_B) \right\} = \sum_{T,s} \alpha_{T,s} (p_{T,s}^* + p_{\kappa T,s}^*)$$

In the proofs of going between the Pareto allocation and competitive equilibrium we will assume that individuals are non-satiated, but this assumption is not crucial because we can always expand the commodity space such that this assumption holds.

The following Theorems 1 to Theorem 3 prove that for all Pareto weights, there is a competitive equilibrium that replicates the social planner's problem.

Theorem 1. *If all agents are non-satiated, a competitive equilibrium (p^*, x^*, y^*) is a Pareto optimal allocation (x^*, y^*) . [First Welfare Theorem]*

The proof is standard, and is in the Appendix.

Theorem 2. *Any Pareto optimal allocation (x^*, y^*) can be achieved through a competitive equilibrium with transfers between agents subject to there being a cheaper point for all agents and agents are non-satiated.*

The proof is in the Appendix. The proof relies on first showing that any solution to Pareto program can be supported as a compensated equilibrium and subsequently showing that we can map the Lagrange multipliers and other variables from any Pareto optimal allocation into a compensated equilibrium. Finally, the proof shows that a compensated equilibrium is a competitive equilibrium, subject to the existence of a cheaper point. The proof follows from Prescott and Townsend [2005].

Theorem 3. *For any given distribution of endowments, if the Pareto weights at a fixed point of the mapping are non-zero, then a competitive equilibrium exists.*

The proof is in the Appendix. The proof relies on finding a mapping that satisfies the conditions of Kakutani's fixed point theorem and subsequently showing that this fixed point is a competitive equilibrium, subject to the existence of a cheaper point. The proof follows from Prescott and Townsend [2005].

4 Results

4.1 How does market power affect the allocation of resources and rent?

In section (3) we showed that the competitive equilibrium is a Pareto optimal allocation. Here we will analyze the monopoly equilibrium. The main difference between the competitive equilibrium and the monopoly equilibrium is the number and type of platforms created. In particular, the price-setting intermediary in the monopoly equilibrium will restrict the supply of platforms to maximize its rent.

Theorem 4. *The price-setting intermediary in the monopoly equilibrium will capture all the rent in the economy and will produce fewer slots than the price-taking intermediary in the competitive equilibrium.*

The proof shows that in the monopoly equilibrium, the intermediary will use its price setting power to charge higher prices (than the competitive equilibrium), thereby reducing the number of platforms created in equilibrium. Moreover, in our environment the price-setting intermediary can set prices in such a way that it captures the whole rent. In contrast, in the competitive equilibrium, the prices to join a platform adjust such that in equilibrium, the economy's total resources will be fully utilized to build platforms. The full proof is provided in Section (8.4) in the Appendix.

Overall, market structure changes both the allocation of rents and the allocation of resources within the economy. In particular, competition ensures that the intermediary makes no profits, and that surplus is accrued by the agents. Further, competition ensures that all the resources in the economy are used to produce platforms.

4.2 Prices for joining a platform in a competitive equilibrium.

To better understand how prices are determined in the competitive equilibrium, we analyze the agent's maximization problem. The agents' maximization problem can be written as the following Lagrange maximization problem. We can use this problem to show which contracts the agent of type T buys.

$$L = \sum_{d_T} x_{T,s}[d_T(N_A, N_B)] U_T[d_T(N_A, N_B)] - \mu_{T,s}^W \left(\sum_{d_T} x_{T,s}[d_T(N_A, N_B)] p[d_T(N_A, N_B)] - \kappa_{T,s} \right) - \mu_{T,s}^P \left(\sum_{d_T} x_{T,s}[d_T(N_A, N_B)] - 1 \right)$$

The first order condition for Type T and contract $x_{T,s}[d_T(N_A, N_B)]$ is the following:

$$U_T(N_A, N_B) - \mu_{T,s}^P - \mu_{T,s}^W * p[d_T(N_A, N_B)] \leq 0 \quad (18)$$

Where $\mu_{T,s}^P$ is the Lagrange multiplier associated with the individual being assigned to some platform, and $\mu_{T,s}^W$ is the Lagrange multiplier associated with the agent's budget constraint. Furthermore, for any platform the agent buys with positive probability ($x_{T,s}[d_T(N_A, N_B)] > 0$), the equation will hold with equality. If the left-hand side of equation (18) is strictly less than zero, that agent will not purchase that contract.

Let us analyse what equation (18) implies; consider an agent of type T who purchases strictly positive measures of two different contracts, $d_T(N_A, N_B)$ and $d_T(N'_A, N'_B)$. Let us define the variables $\Delta U \equiv U_T(N_A, N_B) - U_T(N'_A, N'_B)$ and $\Delta p \equiv p[d_T(N_A, N_B)] - p[d_T(N'_A, N'_B)]$, then we can state:

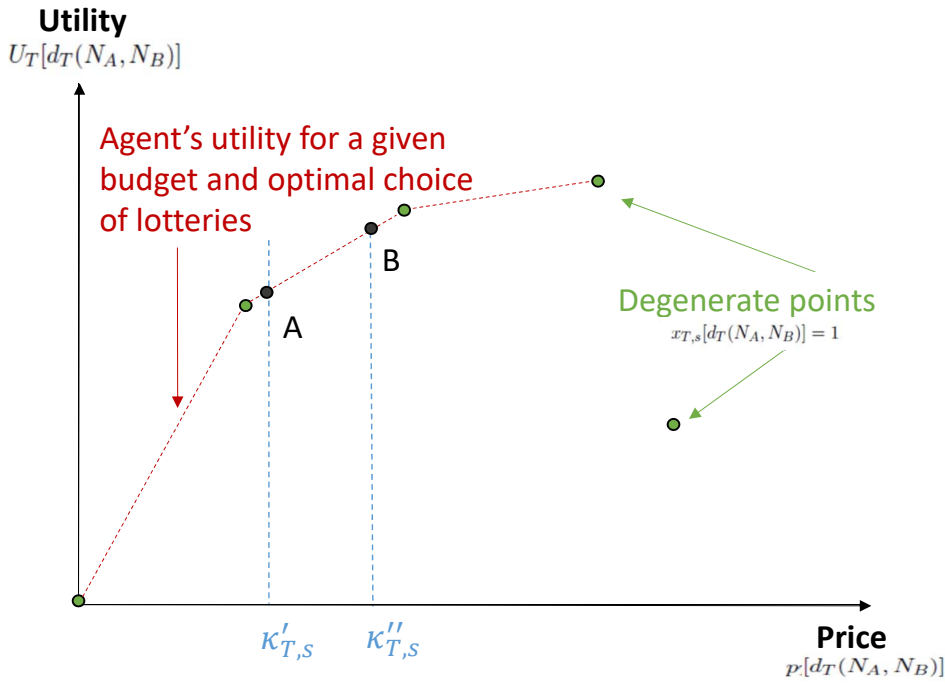
$$\Delta U = \mu_{T,s}^W * \Delta p$$

Therefore, if an agent buys strictly positive measures of two contracts, the difference in utility between the two contracts will be a constant multiplied by the difference in price.

When an agent’s probability constraint is binding ($\mu_{T,s}^P > 0$) and the agent buys two different contracts, the more expensive contract will confer less “utils per dollar” than the less expensive contract.¹⁶ Intuitively, when an agent’s probability constraint binds, this agent would prefer to join multiple platforms of the same type, but is constrained by the requirement to join only one platform.

Figure 2 demonstrates this result. If the agent buys bundle “B”, the agent’s utility per dollar is lower than if the agent buys bundle “A”—raising the question, why does the agent not just buy more of bundle A? The agent would like to buy more of bundle “A” but is constrained by joining only one platform in equilibrium (the probability constraint in the agent’s maximization problem, equation 3).

Fig. 2: Agent’s optimal bundle choice



To compute the prices paid in equilibrium, we explain our procedure in greater detail in Section 8.5 (in the Appendix), specifically we solve the social planner’s problem and subsequently solve for the prices paid by each agent in equilibrium.

5 Competitive Equilibrium Examples

In a general equilibrium framework we can analyze how both the composition of platforms and the resulting utilities change as we alter parameters. To develop intuition, we provide five experiments in our environment.

¹⁶ Conversely, when the probability constraint does not bind and the agent buys multiple different contracts, all purchased contracts will confer the same utils per dollar.

First, as a useful benchmark we examine an equilibrium where we have symmetric parameters for both sides of the market—that is, the same costs, preferences, and wealth.

Second, we analyze an example that varies the wealth *within* and *across* types. Leveraging our general equilibrium approach, we demonstrate that even with symmetric preferences and costs, a subtype with *lower* wealth may actually be *better* off than an alternative subtype.

Third, we explore more generally how the equilibrium—and, subsequently, agents’ utilities—change, as we redistribute wealth within our economy.

Fourth, we examine how the equilibrium utilities change as we alter the Pareto weights. Counter-intuitively, we show that even if an agent’s relative Pareto weight falls, their equilibrium utility can actually rise, depending on the general equilibrium matching effects.

Fifth, given that the cost of producing platforms changes over time (for instance because of technological improvement), we demonstrate how the equilibrium utilities change as we alter the fixed cost of producing platforms. We show that increasing fixed costs leads to heterogeneous effects and, potentially, to increases in inequality.

To improve intuition, in the following examples, let us apply our model to digital currency platforms. Each digital currency platform connects merchants to consumers. We have in mind platforms such as Bitcoin, Ethereum, Kodakcoin, and Zcash. For ease of exposition, we limit attention to only two subtypes of merchants—‘Small’ merchants ($A, 1$) and ‘Big’ merchants ($A, 2$)—and to only two subtypes of consumers—‘Lay’ consumers ($B, 1$) and ‘Tech’ consumers ($B, 2$)—who also may vary in wealth. Each consumer would prefer to be on a platform with more merchants (more advantageous terms) and fewer consumers (less advantageous terms). Similarly, merchants want many consumers to be on their platform but would like fewer rival merchants on their platform, holding the number of consumers fixed.

5.1 Example 1: Symmetric wealth, preferences, population proportions, and cost parameters.

To simplify the exercise, our initial example is symmetric—there are equal fractions of each subtype ($\alpha_{A1} = \alpha_{A2} = \alpha_{B1} = \alpha_{B2}$), each subtype has the same Pareto weight ($\lambda_{A1} = \lambda_{A2} = \lambda_{B1} = \lambda_{B2}$), the cost function is the same for both types ($c_A = c_B$) and the utility functions’ parameters are the same ($\gamma_A = \gamma_B$ and $\epsilon_A = \epsilon_B$).¹⁷ In this initial example, although there are nominally two subtypes of merchants and consumers in the notation, they are in fact identical, and therefore there is no variation by subtype.

In this equilibrium, all users, merchants and consumers, pay a price of two units of capital to join a digital payment platform that matches them with two users of the other type, and one more user of their own type so the total number of people on each platform is four. All platforms are of equal size and producing each platform requires 8 units of capital. In total, there is a measure of 0.5 of these platforms, which makes sense, as that number multiplied by the number of each type on the platform, 2, delivers the measure of each type, unity.

¹⁷ The parameter values are: $\alpha_{A1} = \alpha_{A2} = \alpha_{B1} = \alpha_{B2} = \frac{1}{2}$; $c_A = c_B = c = 1$, $K = 0$; $\gamma_A = \gamma_B = \epsilon_A = \epsilon_B = \frac{1}{2}$; $\lambda_{A1} = \lambda_{A2} = \lambda_{B1} = \lambda_{B2} = \frac{1}{4}$.

Table 1: Equilibrium platforms and user utility for Example 1

Equilibrium platforms					
Platform Size (N_A, N_B)	Number of Platforms Created $y(N_A, N_B)$	Cost of Production $C(N_A, N_B)$			
(2,2)	0.5	8			

Equilibrium user utility and platform choice					
Type (T, s)	Wealth ($\kappa_{T,s}$)	Platform Joined (N_A, N_B)	Price of Joining $p(d_T[N_A, N_B])$	Pr(joining) $x_{T,s}(d_T[N_A, N_B])$	Utility on Platform $U_T(N_A, N_B)$
A,1	2	(2,2)	2	1	2.41
A,2	2	(2,2)	2	1	2.41
B,1	2	(2,2)	2	1	2.41
B,2	2	(2,2)	2	1	2.41

5.2 Example 2: Different wealth but otherwise same preferences, population proportions and cost parameters as earlier.

Our second example varies wealth both *within* and *across* types but otherwise keeps all parameters and demographics the same.¹⁸

In this equilibrium, two different sizes of platforms are created. One set of platforms are larger than the other, of size 5, and are populated with relatively more merchants than consumers. Its existence is due to the richer tech consumers—the wealthiest group in the entire population. Tech consumers obtain the highest utility, as they join platforms that are bigger and have a more favorable ratio of merchants to consumers. The poorer lay consumers join smaller platforms, of size 3; these platforms are populated with a less favorable ratio of merchants to consumers. This less favorable ratio causes lower utility for lay consumers (and the lower utility for consumers is reflected by lower prices for consumers to join that platform). Thus, the poorer lay consumers absorb the population imbalance created by platforms catering to the richer tech consumers.

Turning to the type of contracts purchased. The tech consumers, lay consumers, and the small merchants all buy contracts, where they are assigned to a particular platform with probability one. Whereas, the big merchants buy a mixture of probabilities in two different platforms; 50 percent are allocated to the platforms of size (3, 2), and 50 percent are allocated to the platforms of size (1, 2). The respective prices for these two different contracts are 1.37 and 1.91. Note that solely buying the platform (1, 2) for merchants costs 1.91, which is beyond the big merchants' budget of 1.64.

¹⁸ There are equal fractions of each type ($\alpha_{A1} = \alpha_{A2} = \alpha_{B1} = \alpha_{B2}$), the cost function is the same for both types ($c_A = c_B$), and the utility functions' parameters are the same ($\gamma_A = \gamma_B$ and $\epsilon_A = \epsilon_B$); however, the agents vary in wealth. The parameter values are: $\alpha_{A1} = \alpha_{A2} = \alpha_{B1} = \alpha_{B2} = \frac{1}{2}$; $c_A = c_B = c = 1$, $K = 0$; $\gamma_A = \gamma_B = \epsilon_A = \epsilon_B = \frac{1}{2}$

Table 2: Equilibrium platforms and user utility for Example 2

Equilibrium platforms						
Platform Size (N_A, N_B)	Number of Platforms Created $y(N_A, N_B)$	Cost of Production $C(N_A, N_B)$				
(3,2)	0.25	11				
(1,2)	0.25	5				

Equilibrium user utility and platform choice						
Type	Wealth	Platform Joined	Price of Joining	Pr(joining)	Platform Utility	Expected Utility
T, s	$\kappa_{T,s}$	(N_A, N_B)	$p(d_T[N_A, N_B])$	$x_{T,s}(d_T[N_A, N_B])$	$U_T(N_A, N_B)$	
<i>Merchant (A)</i>						
Small (A,1)	1.37	(3,2)	1.37	1.0	2.23	2.23
Big (A,2)	1.64	(3,2)	1.37	0.5	2.23	2.52
		(1,2)	1.91	0.5	2.80	
<i>Consumer (B)</i>						
Lay (B,1)	1.54	(1,2)	1.54	1.0	1.70	1.70
Tech (B,2)	3.45	(3,2)	3.54	1.0	2.96	2.96

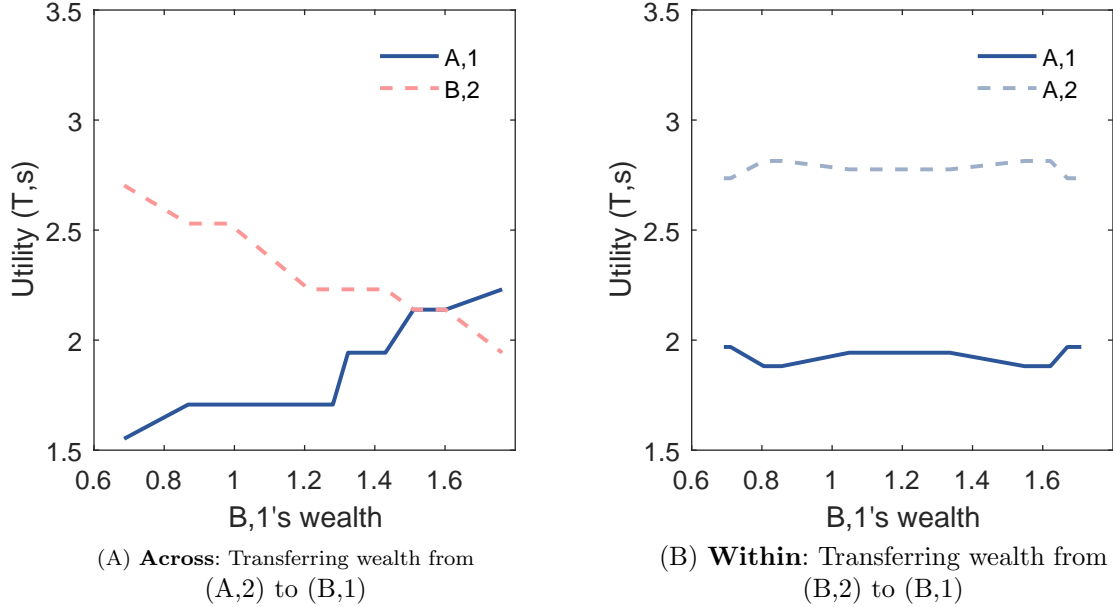
In line with intuition, those that benefit the most from the platforms pay the majority of the total costs (here the consumers). For example, for platforms of size (3, 2), consumers pay 63 percent of the cost of making that platform, yet, consumers are only 40 percent of that platform's population.¹⁹

Finally, even though preferences and costs are symmetric, and that lay consumers (B,1) are *wealthier* than small merchants (A,1), note that lay consumers are relatively *worse off* than small merchants. This result follows from the general equilibrium set-up and that tech consumers are relatively rich and want to join platforms populated with a high number and fraction of merchants. This demand for merchants ensures that small merchants (A,1) are compensated for joining these platforms by contributing relatively less. An alternate way to examine this distribution of costs is that the merchants are in scarcer supply (because consumers are so much wealthier; average consumer wealth is 2.5, and average merchant wealth is only 1.5), yet they need to participate equally on platforms, on average. Therefore, the merchants' price schedule is lower than the consumers' price schedule.

5.3 How does the competitive equilibrium change as we redistribute endowments?

If we redistribute wealth in our economy, this redistribution will change the relative demand for merchants and consumers and subsequently change the relative prices to join a given platform. To examine the

¹⁹ 2 consumers pay 3.45 units of capital each and the 3 merchants pay 1.37 units each. Therefore, consumers pay 6.9 units to join the platform out of the 11 units of capital required to produce the platform.

Fig. 3: Redistributing wealth *across*- and *within*-agent type

general equilibrium effects of redistributing wealth, we construct two placebo interventions that reallocate wealth within our economy while holding the total resources constant.

Figure (3) uses the same cost and preferences as in the previous examples but varies agents' wealth. The left panel shows the effects on the utility of $(A, 1)$ and $(B, 2)$ from redistributing wealth *across* types, in particular from $(A, 2)$ to $(B, 1)$ —that is, $\kappa_{A,2} + \kappa_{B,1} \approx 2.4$.²⁰ The right panel shows the effects on the utility of $(A, 1)$ and $(A, 2)$ from redistributing wealth *within* a type, in particular between $(B, 1)$ and $(B, 2)$.²¹

Starting with the intervention that reallocates wealth across types. Recall our payment platform example that connects merchants to consumers. As we increase lay consumers' wealth ($\kappa_{B,1}$) (at the expense of small merchants, $A,2$), their willingness to pay to join platforms with a higher number and ratio of merchants rises. Subsequently, the price schedule for consumers to join platforms for a given number of merchants will also rise. Therefore, because tech consumers' wealth ($\kappa_{B,2}$) is a constant and they now face higher prices, their utility must fall. We help the subtype receiving more wealth (*in this case* $B,1$) at the expense of the other subtype in the same type ($B,2$) (and symmetrically, hurt the subtype losing wealth $(A,2)$ and help the other subtype in the same type $(A,1)$).

Further (in the right-panel of figure 3), we consider how the equilibrium changes as we adjust the endowments within a type (consumers) and hold the endowments of the other types (merchants) fixed.

²⁰ We solve the model using the Pareto problem and then impute the wealth and prices which replicate the same allocation. We simulate 2880 equilibria for different Pareto weights, and then collect only the equilibria in which $0.51 < \kappa_{A,1} < 0.59$ and $1.01 < \kappa_{B,2} < 1.19$. We then 'join up' all the points to plot a smooth curve.

²¹ We solve the model using the Pareto problem and then impute the wealth and prices which replicate the same allocation. We simulate 2880 equilibria for different Pareto weights, and then collected only the equilibria in which $0.51 < \kappa_{A,1} < 0.59$ and $1.01 < \kappa_{A,2} < 1.1$. We are approximately holding the endowment of $(A, 1)$ and $(A, 2)$ constant.

There is no effect on merchants' utilities because any reduction in purchasing power by one of the consumer subtypes is compensated by an equal change in the other slightly richer consumer subtype.²²

5.4 How does the competitive equilibrium change as we alter the Pareto weights?

We can also consider how the equilibrium changes as we adjust the Pareto weights on only one subtype ($B, 2$).²³

Fig. 4: How does the utility for each subtype change as we alter the Pareto weight (λ) for Tech Consumers ($B,2$)?




Figure (4) demonstrates (for given parameters) how the resulting utilities change as the Pareto weight for type ($B, 2$) increases. First, it is clear and intuitive the utility of ($B, 2$) monotonically weakly increases with their respective Pareto weight. This is a general result.

As figure (4) shows, type ($B, 1$) is clearly disadvantaged. This is a general result and follows from the utility of subtype ($B, 2$) rising.

Recall our merchant and consumer example from before. If we increase the Pareto weight on tech consumers ($\lambda_{B,2}$), the allocation will match them in both larger platforms and with more merchants.

²² There is a tiny change in the utility of ($A, 2$) because of the discrete nature of the possible platform combinations and the changes in the platforms type (B) can purchase.

²³ The parameter values are:
 $\alpha_{A1} = \alpha_{A2} = \alpha_{B1} = \alpha_{B2} = \frac{1}{2}$; $c_A = c_B = c = 1, K = 0$; $\gamma_A = \gamma_B = \epsilon_A = \epsilon_B = \frac{1}{2}$
 $\lambda_{A1} = \frac{1.01-x}{3}, \lambda_{A2} = \frac{0.99-x}{3}, \lambda_{B1} = \frac{1-x}{3}, \lambda_{B2} = x$. We introduce a tiny wedge between ($A, 1$) and ($A, 2$) to highlight the effects on a favored subtype.

This Pareto weight change has two effects in equilibrium: First, there are fewer resources left for the lay consumers, and second, there are fewer merchants left unmatched.

The story is more complicated for the merchants. An increase in the tech consumers' Pareto weight can lead to lower or *higher* utility for merchants. One of the merchants subtypes will always be made worse off ($A, 2$) by the rise in $(\lambda_{B,2})$ because some platforms are composed of relatively more merchants, favoring consumers on those platforms.

As seen in figure (4), it is possible for one of the merchant subtypes ($A, 1$) to be made better off—even as their relative Pareto weight *falls* (as the Pareto weight for $B, 2$ increases from 0.3 to 0.35). This result occurs because the most favored consumer subtype (tech) is matched to proportionally more merchants as $\lambda_{B,2}$ increases, and so the proportion of remaining merchants to consumers declines. Subsequently, those merchants who are not matched with tech consumers may be matched at favorable ratios of consumers to merchants, increasing their utility.

5.5 How does the competitive equilibrium change as we alter the costs of building platforms?


A further important consideration is how the equilibrium changes as we adjust costs; for instance, lower electricity costs or technological innovations may decrease the costs of creating a platform. Figure (5) shows how the equilibrium changes as the fixed cost of building a platform changes.²⁴ As one would expect for a given distribution of wealth, as the fixed cost rises, utility falls. However, interestingly, the distribution of utility also changes. Figure (5) shows that the richest subtype ($A, 2$) is barely affected by the rise in platform costs. In contrast, the poorest subtype's utility, ($A, 1$)'s utility, falls about 50 percent as we increase the fixed cost of building a platform from 0.2 units of capital to 2 units of capital. In this example, the poorest agents are most adversely affected by increasing the platforms' fixed costs.

For larger fixed costs of producing a platform, the distribution becomes more dispersed, and inequality between different subtypes becomes more pronounced. To gain intuition for this result, recall our interpretation that agents are endowed with two assets: labor and capital. As we increase the costs of producing platforms of a given size, the relative value of capital to labor becomes larger. Therefore, agents who are endowed with more capital are less hurt by the rise in costs, leading to greater inequality.

In figure (6) we examine how the characteristics of the equilibrium change as we increase the fixed cost from 0 to 2 units of capital. Specifically, we plot the number of platforms (top left); the average size of the platforms (top right); the total number of slots (bottom left); and the number of agents in autarky (bottom right). As the fixed cost rises, the number of active platforms falls by nearly 50 percent (top-left panel), the average size of platforms increases by nearly 25 percent (top-right panel), and, on net, participation falls (bottom left and bottom right panels). That is, the number of agents in autarky, not participating in platforms in any way, is increasing. The intuition for this result follows from two key effects. First, the cost of building platforms of any size rises, hitting the resource constraint and causing

²⁴ The economy's parameters are $\alpha_{A1} = \alpha_{A2} = \alpha_{B1} = \alpha_{B2} = \frac{1}{2}$; $c_A = c_B = c = 1$; $\gamma_A = \gamma_B = \epsilon_A = \epsilon_B = \frac{1}{2}$; $\kappa_{A1} = 0.5, \kappa_{A2} = 1.5, \kappa_{B1} = 0.8, \kappa_{B2} = 1.1$. For computational simplicity, we allow the equilibrium wealth levels to be close to the desired wealth levels ($\kappa_{A1} = 0.5, \kappa_{A2} = 1.5, \kappa_{B1} = 0.8, \kappa_{B2} = 1.1$). We only plot the equilibrium utilities for those equilibria such that the maximum difference between the desired wealth endowment and the plotted capital endowment is less than 0.1 units of capital for each subtype.

Fig. 5: How does the utility for each subtype change as we alter the fixed cost of building a platform?



the set of platforms to become smaller. The mass of platforms drops from 0.52 to 0.31 as the fixed cost of building a platform rises from 0 to 2 units of capital (upper left panel). Second, note that the relative cost per slot increases more for small platforms than for big platforms, causing larger platforms to be produced in equilibrium. The number of users per platform rises to 5 (upper-right panel). Again, the net effect is that the number of available contracts drops to 1.55 (bottom-left panel) and the number of agents in autarky rises to 0.45 (bottom-left panel).

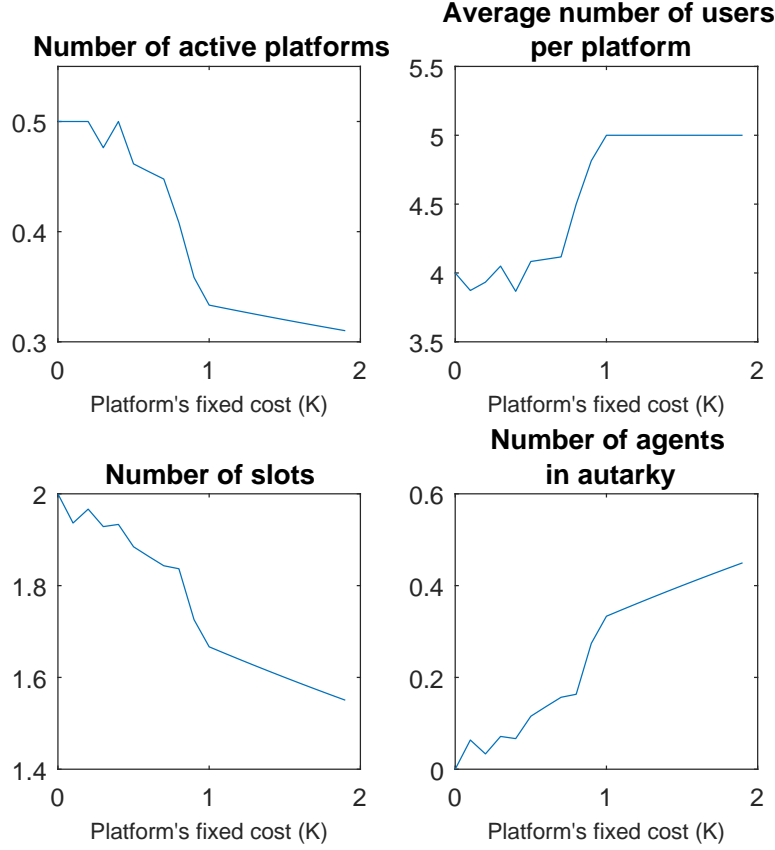
To examine the distributional effect on participation as the fixed cost rises, figure (7) shows participation by subtype. As the fixed cost rises—and, subsequently, the cost of joining platforms rises—the poorest subtypes (A,1 and B,1) become less likely to participate, whereas the richer subtypes (A,2 and B,2) continue to always join a platform.

Do these distributional effects suggest a rationale for regulating prices? No—the equilibrium outcome is Pareto optimal, so the optimal government intervention would be to introduce lump-sum taxation on the rich and transfers to the poor. This transfer of wealth would increase the utility of the poorest, achieving a more equitable division of utility while maintaining a Pareto optimal allocation. Alternative interventions would be distorting.

6 Extensions to the model

Our model is relatively general and can be extended in multiple ways. In this section, we present one extension. Specifically, We allow subtypes to have differences in preferences.

Fig. 6: How do the platform characteristics vary as we alter the fixed cost of building a platform?



6.1 How does user heterogeneity in preferences (within type) affect the competitive equilibrium?

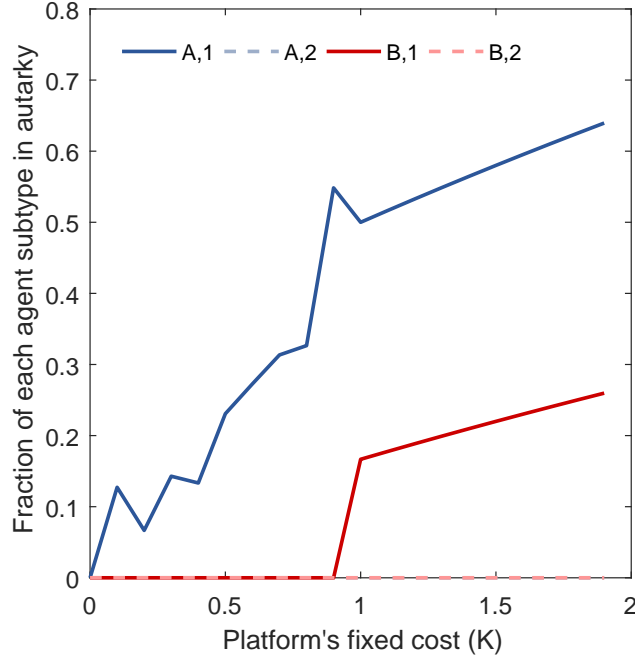
We have concentrated on all types having the same preferences, but potentially varying in their wealth endowments. In this subsection, we consider how varying preferences *within* type affect the competitive equilibrium.

In our reformulated economy we introduce three new parameters $(\beta_1^{T,s}, \beta_2^{T,s}, \beta_3^{T,s})$, which potentially vary across type (T) and subtype (s). Further, we had allowed the preference parameters γ , and ϵ to vary across types in previous sections, in these future experiments we are varying across both type and subtype. The merchant (A, i) 's utility function is now:

$$U_{A,i}(N_A, N_B) = \begin{cases} 0 & \text{if } N_A \text{ or } N_B = 0 \\ \left[\beta_1^{A,i} \left(\frac{N_B}{N_A} \right)^{\gamma_A} + \beta_2^{A,i} N_B^{\epsilon_A} + \beta_3^{A,i} \right] & \text{else} \end{cases}$$

In particular note that: $\beta_1^{A,i}$ alters the merchant (A, i) 's utility with respect to the ratio of consumers and merchants on the platform. $\beta_2^{A,i}$ alters the merchant (A, i) 's utility with respect to the size of the platform (holding the ratio of consumers and merchants constant). Finally, $\beta_3^{A,i}$ is the merchant (A, i) 's

Fig. 7: How does participation by subtype change as the fixed cost of building a platform rises?



intrinsic value from joining a platform. Therefore, the introduction of the parameters $(\beta_1^{T,s}, \beta_2^{T,s}, \beta_3^{T,s})$ facilitates the comparison of how users who vary in their preferences alter the resulting equilibrium.²⁵

Symmetrically, consumer (B, j) 's utility function is:

$$U_{B,j}(N_A, N_B) = \begin{cases} 0 & \text{if } N_A \text{ or } N_B = 0 \\ \left[\beta_1^{B,j} \left(\frac{N_A}{N_B} \right)^{\gamma_A} + \beta_2^{B,j} N_A^{\epsilon_B} + \beta_3^{B,j} \right] & \text{else} \end{cases}$$

Recalling our prior example describing a payment platform, it is natural to consider that lay and tech consumers will vary in preferences as well as wealth. For instance, a lay consumer may both be poor and prefer to be on *any* platform (high $\beta_3^{B,j}$), whereas the tech consumer may prefer to have a choice of merchants (higher $\beta_1^{B,j}$).

In contrast, to Armstrong [2006], Weyl and White [2016], who show that heterogeneity in user preferences leads to market failure, our economy's competitive equilibrium remains Pareto efficient. The main difference in our papers' results is caused by our differing modeling choices. In Armstrong [2006] and Weyl and White [2016]'s models, each oligopolistic platform potentially serves users with varying preferences and can only partially extract consumer surplus, leading to potentially socially inefficient prices, whereas our model has free entry for platforms (as opposed to exogenously fixing the number of platforms), which (i) allows the possibility of complete platform differentiation, and (ii) prevents pricing distortions due to market power. Subsequently, users may separate according to their preferences; for instance, if there


²⁵ Note if $\beta_1^{T,s} = 1, \beta_2^{T,s} = 1$ and $\beta_3^{T,s} = 0$ for all types and subtypes, we have the same utility function as previous sections.

are subtypes who strongly prefer larger platforms, they can join other agents who strongly prefer larger platforms. Note, that once we have converted our economy to a standard looking Walrasian one, it is less surprising that heterogeneity in preferences is not a source of problems.

The introduction of consumer heterogeneity in preferences does lead to interesting comparative statics. To understand in greater detail how the differences in preferences affect the competitive equilibrium, we apply the new utility function to the experiment from Section (5.5).

In figure (8) we plot how the equilibrium utilities (for the new utility function) for each subtype vary, and, in addition, we alter the fixed cost of building a platform for some given parameters.²⁶ To introduce differences in user preferences, we alter the tech consumers', subtype (B,2), preferences, while keeping all other subtypes preferences unchanged from the previous section. We change the tech consumers' preferences in two ways: (i) tech consumers (B,2) strongly prefer to be on a platform with a large number of merchants ($\beta_1^{B,2} = 3$) relative to lay consumers ($\beta_1^{B,1} = 1$) and (ii) tech consumers are relatively indifferent about the size of the platform ($\beta_2^{B,2} = 0.01$), whereas lay consumers prefer larger platforms ($\beta_2^{B,1} = 1$).

Fig. 8: How does the utility for each subtype change as we alter the fixed cost of building a platform with the new utility function?



²⁶ The economy's parameters are $\alpha_{A1} = \alpha_{A2} = \alpha_{B1} = \alpha_{B2} = \frac{1}{2}$; $c_A = c_B = c = 1$; $\gamma_A = \gamma_B = \epsilon_A = \epsilon_B = \frac{1}{2}$; $\kappa_{A1} = 0.7, \kappa_{A2} = 1.3, \kappa_{B1} = 1, \kappa_{B2} = 1$.

Further we make tech consumers (B,2) strongly prefer platforms that have a favorable ratio of consumers to merchants ($\beta_1^{B,2} = 3$), be mostly indifferent about the size of the platform ($\beta_2^{B,2} = 0.01$) and have little to no benefit from being on a platform ($\beta_3^{B,2} = 0.01$). For all other types, we maintain the previous utility function ($\beta_1^{A,1} = \beta_1^{A,2} = \beta_1^{B,1} = 1$), ($\beta_2^{A,1} = \beta_2^{A,2} = \beta_2^{B,1} = 1$) and ($\beta_3^{A,1} = \beta_3^{A,2} = \beta_3^{B,1} = 0$).

For computational simplicity, we allow the equilibrium wealth levels to be close to the desired wealth levels.

Comparing figure (5) and figure (8), we see that the tech consumers (B,2) are the most adversely affected by increasing the platform's fixed cost with the new utility function, whereas, the other subtypes are significantly less affected.

Intuitively, as the platform's fixed cost increases, the relative price of smaller platforms becomes higher. Therefore, the competitive equilibrium is composed of larger platforms but with a small number of active platforms. Consequently, even though the cost of building platforms is larger (and, subsequently, the production possibility frontier of the economy is shrinking), the equilibrium utility of lay consumers and the merchants are relatively unchanged. The big losers in this experiment are the tech consumers—who relatively prefer smaller platforms with a higher fraction of merchants. In contrast, the other subtypes relatively prefer larger platforms and subsequently are less affected by the rise in the fixed cost.

7 Conclusion

There are many economic platforms that must cater to multiple, differentiated users who, in turn, care about who else the platform serves—for instance, credit cards, clearinghouses, and dark pools, to name but a few (Rochet and Tirole [2003], Ellison and Fudenberg [2003], Rochet and Tirole [2006], Caillaud and Jullien [2003], Armstrong [2006], Rysman [2009], Weyl [2010] and Weyl and White [2016]). Over-the-counter markets can also be conceptualized in this way—who is trading with whom, what is the network architecture, and what is the overall degree of direct and indirect connectedness (Allen and Gale [2000], Leitner [2005], Allen et al. [2012], Acemoglu et al. [2015], Cohen-Cole et al. [2014], Elliott et al. [2014]). Modeling each of these arrangements is inherently difficult and there is much more to be done. Here we try to capture each of the applications in a stylized way by building a common conceptual framework for analysis.

Our paper has four main contributions.

Our first contribution is methodological. As in the prior work on firms as clubs by Prescott and Townsend [2006], which builds on Koopmans and Beckmann [1957], Sattinger [1993], Hornstein and Prescott [1993], Prescott and Townsend [1984] Hansen [1985] and Rogerson [1988], we model an economy with competing platforms in a general equilibrium framework, with platforms as clubs. Our framework is relatively general; we can analyze an economy with many (that is, more than two) types of users, who may have heterogeneous preferences; an economy with heterogeneous costs for servicing different users; or an economy with inherent differences within a type's wealth.

Second, our economy incorporates the fact that an individual's utility may be contingent on the actions of others—in short an externality. But we show how to internalize interdependencies so that they do not lead to an inefficient equilibrium overall. In particular, the potential externality is “priced” – in a manner suggested by Arrow [1969]. The competitive equilibrium is efficient.

Third, we demonstrate how changes in one agent's wealth (or Pareto weight) have interesting general equilibrium effects both within- and across-types. The matching in the economy is endogenous, and the math of assignment has to work out in the general equilibrium. For instance, consider a payment platform for consumers and merchants where there are two subtypes of consumers, lay and tech. An increase in the

lay consumer's wealth will lead to *decreases* in the *tech* consumer's welfare and ambiguous effects on the merchant's welfare. This result follows from our assumption that agents do not like to be on a platform with more of their own type, and therefore as we increase the lay consumer's wealth, the lay consumers will prefer platforms with more merchants (and fewer consumers). This increase in lay consumers' wealth is also bad for some merchants with low wealth, as they are now on platforms with fewer consumers and relatively more merchants. Further, the rise in the lay consumer's wealth will lead lay consumers to pay a greater fraction of the costs of being on a platform.

Fourth, we show how technological progress may reduce economic inequality and relatively increase the participation of poorer individuals—an important implication for policymakers looking to increase financial access. Specifically, in our model, a reduction in the fixed cost of building a platform reduces the relative value of capital (that is wealth) and subsequently allows both bigger and more platforms to be created which in turn creates more demand from the various subtypes to join the platform. The biggest utility gain is for the lowest wealth subtypes, who can now join some platforms rather than reside in autarky/non participation.

We should make clear at the same time the limitations of our framework. First, our model is purely static, and we exclude any coordination failures (Caillaud and Jullien [2003], Ellison and Fudenberg [2003], Ellison et al. [2004], Ambrus and Argenziano [2009], Lee [2013], Weyl and White [2016]) and any possibility of innovation in platform design as an intrinsic part of the model.

Second, no platforms or agents have any pricing power in our model, which as Weyl [2010] and Weyl and White [2016] show may interact with the agent's preferences over other agents' actions to exacerbate or minimize market failures. A key question in the two-sided market literature is the allocation of fees. In our Walrasian set up there is no rationale for the regulation of prices on a platform – if a social planner wishes to implement a more equitable allocation, a social planner should redistribute wealth and not regulate prices.

Third, the only source of platform differentiation arises from the size and composition of a platform's users. Relatedly, we also require the characteristics of agents to be clearly identified and rules enforced (that is, no adverse selection or false advertising). Some might find it implausible that the neighborhood composition can be so tightly controlled.

Fourth, we do not allow ever increasing economies of scale in platform size. The existence of economies of scale remains an empirical matter, depending on the particular platform and the market one has in mind. But for some there is no presumption of ever increasing returns. Duffie and Zhu [2011] argue there are economies of scales for central counterparty clearinghouse (CCP) platforms but O'Hara and Ye [2011] for equity market platforms and Altinkiliç and Hansen [2000] for capital issuance find contrary evidence.

Our work's most significant difference to the existing two-sided market literature and the macro financial literature is our methodology. We concentrate on modeling platforms in a Walrasian equilibrium with an extended commodity space with complete contracts and exclusivity. In contrast, the two-sided market literature concentrates on modeling platforms in a partial equilibrium environment and the macro financial literature typically imposes incomplete contracts or a particular institutional arrangement or game. Moreover, the two sided market literature focuses on how market power and imperfect competition affect platform economics, while our framework considers perfect competition between platforms. The

macro financial literature argues explicitly or implicitly for regulation, to ensure stability, and sometimes, externalities is the key rationale. Whereas, we argue for the appropriate design of markets ex ante and letting rights to trade be priced in equilibrium to remove externalities (see also Kilenthong and Townsend [2014]). Therefore, our alternative modelling methodology—explicitly looking at perfect competition with complete contracts—ensures that we can analyze different questions (such as whether the outcome is Pareto optimal), and examine different comparative statics (such as how does inequality change as we increase the fixed cost of building a platform).

We do not view our paper as the final word. In some sense we are trying to arbitrage across distinct literatures, bringing some general equilibrium insights to applied problems in industrial organization and market design/ regulation. Ultimately, modeling and understanding platform economies with more nuanced, but important, details is crucial. We hope this paper ignites a discussion on how to model and analyze multiple, competing platforms.

References

- D. Acemoglu, A. Ozdaglar, and A. Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015.
- F. Allen and D. Gale. Financial contagion. *Journal of Political Economy*, 108(1):1–33, 2000.
- F. Allen, A. Babus, and E. Carletti. Asset commonality, debt maturity and systemic risk. *Journal of Financial Economics*, 104(3):519–534, 2012.
- O. Altinkiliç and R. Hansen. Are there economies of scale in underwriting fees? evidence of rising external financing costs. *Review of Financial Studies*, 13(1):191–218, 2000.
- A. Ambrus and R. Argenziano. Asymmetric networks in two-sided markets. *American Economic Journal: Microeconomics*, 1(1):17–52, 2009.
- M. Armstrong. Competition in two-sided markets. *The RAND Journal of Economics*, 37(3):668–691, 2006.
- K. J. Arrow. The organization of economic activity: issues pertinent to the choice of market versus nonmarket allocation. *The Analysis and Evaluation of Public Expenditure: the PPB system*, 1:59–73, 1969.
- W. F. Baxter. Bank interchange of transactional paper: Legal and economic perspectives. *The Journal of Law & Economics*, 26(3):541–588, 1983.
- S. Bradley, A. Hax, and T. Magnanti. Applied mathematical programming. 1977.
- B. Caillaud and B. Jullien. Chicken & egg: Competition among intermediation service providers. *RAND Journal of Economics*, pages 309–328, 2003.
- E. Cohen-Cole, A. Kirilenko, and E. Patacchini. Trading networks and liquidity provision. *Journal of Financial Economics*, 2014.
- K. Croman, C. Decker, I. Eyal, A. E. Gencer, A. Juels, A. Kosba, A. Miller, P. Saxena, E. Shi, E. G. Sirer, et al. On scaling decentralized blockchains. In *International Conference on Financial Cryptography and Data Security*, pages 106–125. Springer, 2016.
- D. Duffie and H. Zhu. Does a central clearing counterparty reduce counterparty risk? *Review of Asset Pricing Studies*, 1(1):74–95, 2011.
- M. Elliott, B. Golub, and M. O. Jackson. Financial networks and contagion. *American Economic Review*, 104(10):3115–53, 2014.
- G. Ellison and D. Fudenberg. Knife-edge or plateau: When do market models tip? *The Quarterly Journal of Economics*, 118(4):1249–1278, 2003.
- G. Ellison, D. Fudenberg, and M. Möbius. Competing auctions. *Journal of the European Economic Association*, 2(1):30–66, 2004.
- J. Frost, L. Gambacorta, Y. Huang, H. S. Shin, and P. Zbinden. Bigtech and the changing structure of financial intermediation. *Bank for International Settlements Working Paper*, 2019.
- D. Gale, H. W. Kuhn, and A. W. Tucker. Linear programming and the theory of games. *Activity Analysis of Production and Allocation*, 13:317–335, 1951.
- A. Hagiu. Pricing and commitment by two-sided platforms. *The RAND Journal of Economics*, 37(3):720–737, 2006.
- G. D. Hansen. Indivisible labor and the business cycle. *Journal of Monetary Economics*, 16(3):309–327, 1985.

- A. Hornstein and E. C. Prescott. The firm and the plant in general equilibrium theory. *General equilibrium, growth, and trade*, 2:393–410, 1993.
- W. T. Kilenthong and R. M. Townsend. A market based solution to price externalities: A generalized framework. *National Bureau of Economic Research Working Paper No. 20275*, 2014.
- T. C. Koopmans and M. Beckmann. Assignment problems and the location of economic activities. *Econometrica: Journal of the Econometric Society*, pages 53–76, 1957.
- L. C. Koutsougeras. Coalitions with limited coordination. *Economic Theory*, pages 1–18, 2020.
- R. S. Lee. Vertical integration and exclusivity in platform and two-sided markets. *The American Economic Review*, 103(7):2960–3000, 2013.
- Y. Leitner. Financial networks: Contagion, commitment, and private sector bailouts. *The Journal of Finance*, 60(6):2925–2953, 2005.
- J. Y. Lin, J. Miao, and P. Wang. Convergence, financial development, and policy analysis. *Economic Theory*, 69(3):523–568, 2020.
- J. Mallet. Limits on the communication of knowledge in human organisations. *Studies in Emergent Order 2: 1-18*, 2009.
- L. W. McKenzie. On the existence of general equilibrium for a competitive market. *Econometrica: Journal of the Econometric Society*, pages 54–71, 1959.
- L. W. McKenzie. The classical theorem on existence of competitive equilibrium. *Econometrica: Journal of the Econometric Society*, pages 819–841, 1981.
- M. O’Hara and M. Ye. Is market fragmentation harming market quality? *Journal of Financial Economics*, 100(3):459–474, 2011.
- A. Pawasutipaisit. Family formation in walrasian markets. *University of Chicago PhD Thesis*, 2010.
- E. C. Prescott and R. M. Townsend. Pareto optima and competitive equilibria with adverse selection and moral hazard. *Econometrica: Journal of the Econometric Society*, pages 21–45, 1984.
- E. S. Prescott and R. M. Townsend. Firms as clubs in walrasian markets with private information: Technical appendix. *Federal Reserve Bank Richmond*, 05-11, 2005.
- E. S. Prescott and R. M. Townsend. Firms as clubs in walrasian markets with private information. *Journal of Political Economy*, 114(4):644–671, 2006.
- J.-C. Rochet and J. Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003.
- J.-C. Rochet and J. Tirole. Two-sided markets: a progress report. *The RAND Journal of Economics*, 37(3):645–667, 2006.
- R. Rogerson. Indivisible labor, lotteries and equilibrium. *Journal of Monetary Economics*, 21(1):3–16, 1988.
- M. Rysman. The economics of two-sided markets. *The Journal of Economic Perspectives*, pages 125–143, 2009.
- M. Sattinger. Assignment models of the distribution of earnings. *Journal of Economic Literature*, pages 831–880, 1993.
- R. Schmalensee and D. S. Evans. The economics of interchange fees and their regulation: An overview. 2005.
- R. M. Townsend. Distributed ledgers: Innovation and regulation in financial infrastructure and payment systems. *MIT Press*, Forthcoming.

- E. G. Weyl. A price theory of multi-sided platforms. *The American Economic Review*, pages 1642–1672, 2010.
- E. G. Weyl and A. White. Insulated platform competition. *mimeo*, 2016.
- H. Yang. Targeted search, endogenous market segmentation, and wage inequality. *Economic Theory*, 69 (2):367–414, 2020.

8 Appendix

To simplify the exposition of the proofs for Theorems (1), (2) and (3), we transform the model presented in Section (2) to an environment that is more general with more abstract and simpler notation. Without loss of generality, we denote subagent type (T, s) as agent i and let there be I types.²⁷ As before there are α_i of each agent type i . Let there be n number of possible commodities (or contracts or goods) available to agents. We denote an agent of type i 's purchase of each commodity j as x_{ij} , summarized in vector x_i .²⁸ Moreover, one element of the vector, x_i includes an agent of type i 's consumption of the resource endowment (in the model in the text, this element is each agent type i 's capital endowment, κ_i), but as in the earlier text, agents' receive no direct utility from the consumption of this capital good.²⁹ Finally, we let ξ_i denote the vector endowment of a type i . In the model described in the text, this vector for type i would be zero for all goods except the capital good, where it would be equal to κ_i .

We can write the agent's consumption set using a finite number of linear inequalities (which may hold as an equality or as an inequality), that is,

$$X_i = \{x_i \in R_+^n | g_i x_i - b_i \leq 0\} \quad (19)$$

where g_i is a matrix. The linear constraints $(g_i x_i - b_i \leq 0)$ are the constraints on the agent's choice of contracts. In the model in the text, the agent's consumption set consisted of three constraints: (i) each agent must purchase exactly one measure of contracts (that is, the agent must join a platform with certainty but recall this set of platforms includes the "opt-out" platform, where the agent is on a platform by themselves), (ii) each agent could buy only non-negative measures of contracts, and (iii) each agent could only buy contracts ear-marked for that type. It is trivial to map these constraints into the more general notation above.

Turning to the intermediaries. We can denote the intermediary's production by the vector y , where each element of y is a measure of the commodity produced for agent type i , or the capital input good, y_κ . As such, we represent the intermediary's production set by a finite number of linear inequalities, specifically as:

$$Y = \{y \in R_+^n | f y \leq 0\} \quad (20)$$

where f is a matrix and the constraints may hold with equality or as an inequality.

Using dot-product notation, the Pareto program can be written as:

²⁷ This new formulation reduces the number of subscripts and eases the exposition of the proof.

²⁸ x_{ij} is the equivalent of subtype (T, s) 's purchase of contract $x_{T,s}[d_T(N_A, N_B)]$ in the earlier model.

²⁹ In our earlier model, n would be equal to $2(\bar{N}_A \times \bar{N}_B + 1) + 1$, for further explanation, see footnote ??.

$$\max_{\{x_i\} \geq 0, y} \sum_i \lambda_i \alpha_i u_i x_i \quad (21)$$

$$\text{s.t. } \sum_i \alpha_i (x_i - \xi_i) - y = 0 \quad (22)$$

$$fy \leq 0 \quad (23)$$

$$g_i x_i - b_i \leq 0, \forall i \quad (24)$$

where λ_i is the Pareto weight for agent i and u_i is the agent of type i 's utility from commodity bundle x_i .

To complete our proofs of the welfare theorems and existence of an equilibrium, we make three additional assumptions.

For all i , X_i is bounded. $\sum_i \alpha_i (x_i - \xi_i) \cap Y \neq 0$. For all i , for any x_i component of a feasible allocation, there exists $x'_i \in X_i$ such that $u_i x'_i \geq u_i x_i$

The first assumption is satisfied when the consumption set has a probability measure constraint, ensuring that the consumption set is convex and closed. The second assumption merely requires a feasible allocation exists. The third assumption is a no-satiation assumption, and as described earlier, we can increase the possible set of platforms (specifically, increasing the maximal sized platform) to ensure that this assumption is satisfied.

8.1 Proof of Theorem 1

If all agents are non-satiated, a competitive equilibrium (p^, x^*, y^*) is a Pareto optimal allocation (x^*, y^*) . [First Welfare Theorem]*

Proof The proof for this theorem is standard, a proof by contradiction.

Assume there is some competitive equilibrium (p^*, x^*, y^*) which is not Pareto-efficient, it follows that x^* must be Pareto-dominated by some other feasible consumption plan x' and production plan y' . Therefore, for at least one agent of type i , the bundle x'_i , must be strictly preferred to the bundle x_i^* . Since agents are maximising at x^* for given prices p subject to the constraints on that agent's consumption set, then for this agent type, px'_i must be strictly greater than px_i^* . Summing over all agents, we must then have the following strict inequality:

$$\sum_i \alpha_i p x'_i > \sum_i \alpha_i p x_i^* \quad (25)$$

That is, the total spent by the consumers for allocation x' must be strictly greater than the total spent for allocation x^* .

Similarly, from the intermediary's problem, we know that the intermediary maximises profits taking prices as given. Hence it follows that the intermediary's profits cannot be greater for allocation y' than allocation y^* for prices p , implying:

$$p \cdot y^* \geq p \cdot y' \quad (26)$$

Using these inequalities (25) and (26) and the fact that markets clear, we have the following relationship:

$$\sum_i \alpha_i p x'_i > \sum_i \alpha_i p x_i^* = p \cdot y^* \geq p y' \quad (27)$$

Implying that:

$$\sum_i \alpha_i x'_i > y' \quad (28)$$

But inequality (28) violates the feasibility condition, therefore, allocation (x^*, y^*) must be Pareto optimal.

Our paper's modelling environment is similar to Prescott and Townsend [2005], who analyzed firms as clubs in a general equilibrium setting (whereas our paper analyses platform competition in a general equilibrium setting), and specifically used lotteries to convexify the commodity space. As such, the proofs for Theorem (2) and Theorem (3) follow from their paper. We provide similar proofs to their paper, but try to provide more exposition and intuition for the underlying results.

In proving the second welfare theorem, it is useful to define a compensated equilibrium and a competitive equilibrium in our economy.

Definition 1 A compensated equilibrium (x^*, y^*, p^*) in this economy is defined as:

1. $\forall i, x_i^*$ minimizes $p^* x_i$ subject to $x_i \in X_i$ and $u_i x_i \geq u_i x_i^*$ (all agents minimize cost subject to attaining a required level of utility)
2. y^* maximises $p^* y$ subject to $y \in Y$ (all intermediaries maximise profit subject to feasibility constraint)
3. $\sum_i \alpha_i (x_i^* - \xi_i) = y^*$ (market clearing)

Definition 2 A competitive equilibrium (x^*, y^*, p^*) in this economy is defined as:

1. $\forall i, x_i^*$ maximizes $u_i x_i$ subject to $x_i \in X_i$ and $p^* x_i \leq p^* \xi_i$ (all agents maximise utility subject to their feasibility and budget constraint)
2. y^* maximises $p^* y$ subject to $y \in Y$ (all intermediaries maximise profit subject to feasibility constraint)
3. $\sum_i \alpha_i (x_i^* - \xi_i) = y^*$ (market clearing)

8.2 Proof of Theorem 2

Any Pareto optimal allocation (x^, y^*) can be achieved through a competitive equilibrium with transfers between agents subject to there being a cheaper point for all agents and agents are non-satiated.*

Proof To prove this theorem, we first show that any solution to the Pareto program can be supported as a compensated equilibrium, and second we show that a compensated equilibrium is a competitive equilibrium, subject to the existence of a cheaper point.

To show that any solution to Pareto program can be supported as a compensated equilibrium, we start by showing the first order conditions from the Lagrangian problem are necessary and sufficient for optimized solutions. Then to prove that any Pareto optimal allocation can be achieved with transfers, we show that we can map the Lagrange multipliers and other variables from any Pareto optimal allocation into a compensated equilibrium.

Earlier we described the Pareto problem (equations 21 to 24), now let us start by characterizing the solution to the Pareto program using the first-order conditions. Let u_{ij} be agent type i 's utility from good j . Let p_j be the Lagrangian variable on the market clearing constraint for the j th commodity, let μ be the vector of Lagrangian variables on the production set constraints, and finally let γ_i be the Lagrangian variables on agent i 's consumption set constraints. Let f_j and $g_{i,j}$ correspond to the j th column of matrix f and g_i respectively

The Lagrangian for this problem is:

$$\mathcal{L}(x, y, p, \mu, \gamma) = \sum_i \lambda_i \alpha_i u_i x_i + \sum_j p_j \left(y_j - \sum_i \alpha_i (x_{ij} - \xi_{ij}) \right) + \mu f y + \sum_i \gamma_i (b_i - g_i x_i) \quad (29)$$

Lemma 1 *The allocation (x^*, y^*) , with $x^* \geq 0$, is a solution to the Pareto program if and only if x^* , y^* , and the Lagrangian variables $(p^*, \mu^*, \gamma^* \geq 0)$ satisfy:*

$$\forall i, j, \quad \lambda_i \alpha_i u_{ij} - p_j^* \alpha_i - \gamma_i^* g_{i,j} = 0 \text{ if } x_{ij}^* > 0 \quad (30)$$

$$\forall i, j, \quad \lambda_i \alpha_i u_{ij} - p_j^* \alpha_i - \gamma_i^* g_{i,j} \leq 0 \text{ if } x_{ij}^* = 0 \quad (31)$$

$$\forall j, \quad p_j^* - \mu^* f_j = 0 \quad (32)$$

$$p^* \left(\sum_i \alpha_i (x_i^* - \xi_i) - y^* \right) = 0 \quad (33)$$

$$\sum_i \alpha_i (x_i^* - \xi_i) - y^* = 0 \quad (34)$$

$$\mu^* f y^* = 0 \quad (35)$$

$$f y^* \leq 0 \quad (36)$$

$$\forall i, \quad \gamma_i^* (g_i x_i^* - b_i) = 0 \quad (37)$$

$$\forall i, \quad g_i x_i^* \leq b_i \quad (38)$$

Proof Since the Pareto program is a linear program, the Kuhn-Tucker conditions are necessary and sufficient

To reduce the set of possible results from equations (30) to (38), let us prove a couple of preliminary results:

Lemma 2 *For (p^*, y^*) to be part of a solution to the Pareto program, $p^* y^* = 0$*

Proof Multiplying equation (32) by y_j and summing over all j gives:

$$p^* y^* = \mu^* f y^*. \quad (39)$$

Substituting in equation (36) for the right hand side of equation (39) gives $p^* y^* = 0$. Intuitively, this result follows from constant returns to scale in the intermediary's production function.

Lemma 3 *With non-satiation, there does not exist a solution to the Pareto program such that $p^* = 0$.*

Proof We prove this lemma by contradiction. By non-satiation there must exist for each i , $x'_i \in X_i$ such that $u_i x'_i > u_i x_i^*$. If $p^* = 0$, then by equations (30) and (37), for some $x_{ij}^* > 0$:

$$\lambda_i \alpha_i u_{ij} x_{ij}^* = x_{ij}^* \gamma_i^* g_{i,j} = \gamma_i^* b_i \quad (40)$$

Since $u_i x'_i > u_i x_i^*$, for some $x_{ij}^* > 0$:

$$\lambda_i \alpha_i u_{ij} x'_{ij} > \lambda_i \alpha_i u_{ij} x_{ij}^* = \gamma_i^* b_i \quad (41)$$

where $\lambda_i \alpha_i u_{ij} x_{ij}^* = \gamma_i^* b_i$ follows from equation (40).

However, by the feasibility inequality (38), we have $\gamma_i^* b_i \geq \gamma_i^* g_i x'_i$, therefore, for some $x_{ij}^* > 0$:

$$\lambda_i \alpha_i u_{ij} x'_{ij} > \lambda_i \alpha_i u_{ij} x_{ij}^* = \gamma_i^* b_i \geq \gamma_i^* g_i x'_i \quad (42)$$

But if $p^* = 0$, the inequality $\lambda_i \alpha_i u_{ij} x'_{ij} > \gamma_i^* g_i x'_i$ contradicts equation (30). Therefore, $p^* \neq 0$.

To demonstrate that we can map the Lagrange multipliers and other variables from any Pareto optimal allocation into a compensated equilibrium, it is useful to rewrite conditions 1 and 2 from the definition of the compensated equilibrium (definition 1) in terms of the necessary and sufficiency conditions. For the consumer's minimization problem let $\beta_i \geq 0$ be the Lagrangian multiplier on the constraint $u_i x_i \geq u_i x_i^*$ and let μ be the vector of Lagrangian multipliers on the production constraints.

Lemma 4 *Conditions 1 and 2 in the definition of the compensated equilibrium (definition 1) can be rewritten in the following form:*

1. $\forall i, x_i^* \geq 0$ and dual variables $(\beta_i \geq 0, v_i)$ satisfy condition 1 of a compensated equilibrium if and only if they satisfy:

$$\begin{aligned} \forall j, \quad & \beta_i u_{ij} - p_j - v_i g_{i,j} = 0, (\leq 0 \text{ if } x_{ij} = 0) \\ & \beta_i (u_i x_i - u_i x_i^*) = 0 \\ & u_i x_i - u_i x_i^* \geq 0 \\ & v_i (g_i x_i - b_i) = 0 \\ & g_i x_i - b_i \leq 0 \end{aligned}$$

2. y and Lagrangian multiplier μ satisfy condition 2 of a compensated equilibrium if and only if:

$$\begin{aligned} \forall j, \quad & p_j - \mu f_j = 0 \\ & \mu f y = 0 \\ & f y \leq 0 \end{aligned}$$

Proof This lemma follows from the agent's and the intermediary's problem being linear programs, and as such, the Kuhn-Tucker conditions are sufficient and necessary.

Lemma 5 *Any solution to the Pareto program, with $\lambda \geq 0$ can be supported as a compensated equilibrium.*

Proof Lemma 1 shows that the allocation (x^*, y^*) and Lagrangian variables $(p^*, \mu^*, \gamma^* \geq 0)$ must satisfy certain necessary and sufficient conditions to be solutions to the Pareto problem. Now we show that we can

map the Lagrange multipliers and other variables from any Pareto optimal allocation into a compensated equilibrium.

Let $\beta_i = \lambda_i$, let $p = p^*$, let $\mu = \mu^*$ and let $v_i = \gamma_i^*/\alpha_i$. Using these relabelled variables shows that condition (1) of the compensated equilibrium holds (compare the necessary and sufficient conditions in condition (1) with equations (30), (31) and (38) and use that $x_i = x_i^*$). Using these relabelled variables shows that condition (2) of the compensated equilibrium holds (compare the necessary and sufficient conditions in condition (2) with equations (32) and (36)). Finally, condition (3) of the compensated equilibrium holds (compare the necessary and sufficient conditions in condition (3) with equation (34)).

Lemma 6 Any solution to the Pareto program, with $\lambda > 0$ can be supported as a competitive equilibrium.

Proof Similar to lemma (4) we can rewrite the first condition of the competitive equilibrium (where $\theta_i \geq 0$ is the Lagrangian variable on agent i 's budget constraint and w_i the Lagrangian variable on agent i 's consumption set constraints) as:

$$\begin{aligned} \forall j, \quad u_{ij} - \theta_i p_j - w_i g_{i,j} &= 0, (\leq 0 \text{ if } x_{ij} = 0) \\ \theta_i p(x_i - \xi_i) &= 0 \\ p(x_i - \xi_i) &\leq 0 \\ w_i(g_i x_i - b_i) &= 0 \\ g_i x_i - b_i &\leq 0 \end{aligned}$$

As in the proof of lemma (5), let $p = p^*$ and $\mu = \mu^*$. Additionally, for all i , let $\theta_i = 1/\lambda_i$ and $w_i = \gamma_i^*/\lambda_i \alpha_i$. Using a similar strategy as the previous lemma, it is trivial to show that we can map the Pareto problem into the conditions from the competitive equilibrium.

The last step in the proof is to show that, if there exists a cheaper point, a compensated equilibrium is a competitive equilibrium.

Lemma 7 Take a solution to the Pareto program for $\lambda \geq 0$. If at the corresponding compensated equilibrium, there exists for all i , a cheaper point satisfying $x_i \in X_i$, then $\lambda > 0$.

Proof We prove this lemma using a proof by contradiction. Let x^* be the solution to the Pareto problem, then from equations (30) and (31) we have for all i :

$$x_i^*(\lambda_i \alpha_i u_i - p^* \alpha_i - \gamma_i^* g_i) = 0 \tag{43}$$

Take any $x'_i \in X_i$, then again from equations (30) and (31) we have:

$$x'_i(\lambda_i \alpha_i u_i - p^* \alpha_i - \gamma_i^* g_i) \leq 0 \tag{44}$$

Therefore implying:

$$x_i^*(\lambda_i \alpha_i u_i - p^* \alpha_i - \gamma_i^* g_i) \geq x_i'(\lambda_i \alpha_i u_i - p^* \alpha_i - \gamma_i^* g_i) \quad \forall x_i' \in X_i$$

Assume there exists a cheaper point $x_i' \in X_i$ and consider the case $\lambda_i = 0$, then rearranging the last inequality gives:

$$p^* x_i' \geq p^* x_i^* + \frac{\gamma_i^* g_i}{\alpha_i} (x_i^* - x_i') \quad \forall x_i' \in X_i \quad (45)$$

By equations (37) and (38), we have that:

$$\gamma_i^* (g_i x_i^* - b_i) = 0 \quad (46)$$

$$\gamma_i^* (g_i x_i' - b_i) \leq 0 \quad \forall x_i' \in X_i \quad (47)$$

Taking equations (46) and (47) together suggest that:

$$\gamma_i^* g_i x_i^* \geq \gamma_i^* g_i x_i' \quad \forall x_i' \in X_i \quad (48)$$

Plugging inequality (48) into inequality (45) implies that $p^* x_i' \geq p^* x_i^*$, for all $x_i' \in X_i$, hence contradicting the assumption that a cheaper point exists and $\lambda = 0$.

Taking all the lemmas together, we have proved that any Pareto optimal allocation (x^*, y^*) can be achieved through a competitive equilibrium with transfers between agents subject to there being a cheaper point for all agents and agents are non-satiated.

8.3 Proof of Theorem 3

For any given distribution of endowments, if the Pareto weights at a fixed point of the mapping are non-zero, then a competitive equilibrium exists.

This proof follows from Prescott and Townsend [2005].

Proof To prove this that a competitive equilibrium exists, first, we find a mapping that satisfies the conditions of Kakutani's fixed point theorem and second, we show that this fixed point is a competitive equilibrium. For convenience, Kakutani's theorem is stated below.

Theorem 1 (Kakutani's fixed point theorem) *Suppose that $A \in R^N$ is non-empty, compact, convex set, and that $f : A \rightarrow A$ is an upper hemicontinuous correspondence from A into itself with the property that the set $f(x) \subset A$ is nonempty and convex for every $x \in A$. Then $f(\cdot)$ has a fixed point; that is, there is an $x \in A$ such that $x \in f(x)$.*

First, we normalize possible prices to a bounded, convex, and compact set. Specifically, as agents do not have free disposal (they have to join one and only one platform with certainty), we have not ruled out the possibility that some prices are negative. To allow for negative prices, we restrict prices to the closed unit ball, a compact and convex set:

$$P = \{p \in \mathbb{R}^n \mid \sqrt{p \cdot p} \leq 1\} \quad (49)$$

Second, we want to show there exists a correspondence from (λ, x, y, p) to (λ', x, y, p') , that satisfies the conditions of Kakutani's fixed point theorem. Our correspondence consists of two parts:

$$\begin{aligned} (\lambda) &\rightarrow (x', y', p') \\ (\lambda, x, y, p) &\rightarrow (\lambda') \end{aligned}$$

We restrict each variable to lie in a compact, convex set. Recall there are I types of agents, so we can restrict (i) $\lambda \in S^{I-1}$ (the unit simplex), (ii) $x \in X$, (iii) $y \in \Gamma$, (iv) $p \in P$, where X is the cross product of I and X_i , and $\Gamma = \{y \in Y \mid y \text{ is feasible}\}$. The set Γ is convex, and compactness is given by the compactness of X_i and market clearing.

For the first part of the mapping, from the proof of theorem 8.2, for any $\lambda \in S^{I-1}$, there is a solution to the Pareto program $(x^*, y^*, p^*, \mu^*, \gamma^*)$. We normalize p to lie in the unit circle with³⁰:

$$\tilde{p} = \left\{ \frac{p^*}{\sqrt{p^* \cdot p^*}} \right\}$$

Next, we include in the mapping, the convex hull of prices calculated from the normalization, specifically:

$$p' \in \text{co}\tilde{P}$$

where \tilde{P} is the set of normalized prices.³¹ This mapping is non-empty, compact-valued, convex-valued, and upper hemicontinuous.

The second part of the mapping calculates the new Pareto weights, λ' , as a function of the transfers needed to support an allocation x from prices p . Consider the following function for any $\lambda \in S^{I-1}$:

$$\hat{\lambda}_i = \max \left\{ 0, \lambda_i + \frac{p(\xi_i - x_i)}{A} \right\}, \quad \text{and} \quad \hat{\lambda}'_i = \frac{\hat{\lambda}'_i}{\sum_i \hat{\lambda}'_i} \quad (50)$$

³⁰ Recall that by lemma (3), we ensured that $p^* \neq 0$, thereby allowing this normalization.

³¹ This step ensures that we preserve the convexity of the mapping while keeping prices in P , and since the set of \tilde{p} is convex, this step does not add any relative prices to the mapping.

where A is a large positive number such that:³²

$$A > \sum_i |p(\xi_i - x_i)|$$

This function is continuous, so it is a non-empty, compact-valued, convex valued, and upper-hemicontinuous correspondence.

Therefore, we have defined a mapping from $S^{I-1} \times X \times \Gamma \times P \rightarrow S^{I-1} \times X \times \Gamma \times P$.³³ Each of these sets is convex and compact (and cross-products of convex and compact sets are also convex and compact). Both parts of the mapping are non-empty, convex-valued, compact-valued, and upper hemicontinuous (and cross-products of correspondences preserve these properties), therefore by Kakutani's fixed point theorem, a fixed point (λ, x, y, p) exists.

Notice that the first part of the mapping is a solution to the Pareto program, then by Lemma (4.5), the solution can be supported as a compensated equilibrium.

Now to complete the proof, we need to show that the fixed point satisfies the necessary and sufficient conditions from the Consumer's problem. Assume that $\lambda > 0$, from the second part of the mapping, at a fixed point, $p(\xi_i - x_i)$, must be the same sign for all i . By combining equation (33) from the necessary and sufficient conditions of the Pareto problem, (specifically, $p^*(\sum_i \alpha_i(x_i^* - \xi_i) - y^*) = 0$), and Lemma (2) (specifically, that producers make zero profits, $py = 0$), shows that that $p(\xi_i - x_i) = 0, \forall i$, thereby satisfying one of the conditions. For the other two conditions of the Consumer's problem, they require setting $\theta_i = 1/\lambda_i$ and $w_i = \gamma_i/(\lambda_i \alpha_i)$.

8.4 Proof of Theorem 4

The price-setting intermediary in the monopolistic equilibrium will capture all the rent in the economy and will produce less slots than the price-taking intermediary in the competitive equilibrium.

Proof To begin we show that in the monopolistic equilibrium, the intermediary will produce a negligible amount of platforms. Then we show that the competitive equilibrium will produce platforms that use the entire endowment in the economy.

For simplicity, let us assume there are only two types of agents A and B with no subtypes. Assume a monopolistic intermediary produces X (where X is less than one) platforms of size³⁴ $(1, 1)$ and sells each contract to type T at a price of κ_T/X , where κ_T is the agent T 's wealth. The agents can either participate (that is, buy contracts) or not buy. If the agent does not buy any contracts, their resultant utility is zero.

Let us assume each agent buys X contracts of the platform of size $(1, 1)$. Then type T 's utility will be $XU_T(1, 1)$, that is, the utility of being on a platform of size $(1, 1)$ multiplied by the probability of being on that platform, X .

³² Note because the sets P and X are bounded, there exists such a positive number A .

³³ y is not explicitly used in the portion of this mapping since by Lemma 4.2, any solution to the Pareto program satisfied $py = 0$, and as such, scaling p has no impact.

³⁴ We restrict attention to the platform of size $(1,1)$ for expositional ease, although the intermediary could construct platforms of any given size. Additionally, even though in equilibrium the platform will produce only a negligible amount of this platform, the platform of size $(1, 1)$ would be the cheapest platform to produce.

Could the agent buy any other contract? No, because the monopolist only produces one type of platform. Could the agent buy less of the contract? Yes, but utility is increasing in the purchase of this contract, X , therefore not optimal. Could the agent buy more of the contract? No, because the agent is constrained by their wealth endowment, κ_T .

The intermediary's profit is equal to: $\kappa_A + \kappa_B - X(c_A + c_B + c)$. Therefore, the intermediary's profit is decreasing in X . Therefore, the intermediary will produce the smallest positive number of platforms, X , as possible to maximize profits. Therefore, in the monopolistic equilibrium only a negligible number of platforms will be produced.

In the competitive equilibrium, from theorem (1) – the First Welfare Theorem – we know that the competitive equilibrium is a Pareto optimal allocation, second, given the intermediary's constant returns to scale technology, we know the intermediary makes zero profits. Combining these two results, we know in the competitive equilibrium there will be a positive number of platforms and that the total cost of producing these platforms will be κ_A plus κ_B , the total amount of resources in the economy.

8.5 Computation

Attempting to compute the Pareto problem is difficult due to the large commodity space and the high number of constraints. Therefore, we transform the above Pareto problem by removing the club constraints and subsequently allowing us to use simplex algorithms. These algorithms are quicker and more capable at handling the large commodity and constraint space.

For ease of explanation, let us assume there is only two subtypes of merchants and consumers, that is, $i \in \{1, 2\}$ and $j \in \{1, 2\}$.

First we eliminate the club constraints. Recall equation (15), this constraint can be rewritten in matrices for each contract $d_T(N_A, N_B)$ as

$$\begin{bmatrix} \alpha_{A,1} & \alpha_{A,2} & 0 & 0 & -N_A \\ 0 & 0 & \alpha_{B,1} & \alpha_{B,2} & -N_B \end{bmatrix} \begin{bmatrix} x_{A,1}[d_A(N_A, N_B)] \\ x_{A,2}[d_A(N_A, N_B)] \\ x_{B,1}[d_B(N_A, N_B)] \\ x_{B,2}[d_B(N_A, N_B)] \\ y(N_A, N_B) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad (51)$$

Because, $x_{T,s}[d_T(N_A, N_B)]$ and $y(N_A, N_B)$ must be non-negative, with equation (51), let us define a polyhedral cone, with a single extreme point at the origin. Therefore, using the Resolution Theorem of Polyhedrons, the systems of equations can be represented as the set of all non-negative linear combinations of its extreme rays. Scaling such that each $y(N_A, N_B) = 1$, the extreme rays of this cone are:

$$\begin{aligned} & \left(\frac{N_A}{\alpha_{A,1}}, 0, \frac{N_B}{\alpha_{B,1}}, 0, 1 \right) \\ & \left(\frac{N_A}{\alpha_{A,1}}, 0, 0, \frac{N_B}{\alpha_{B,2}}, 1 \right) \\ & \left(0, \frac{N_A}{\alpha_{A,2}}, \frac{N_B}{\alpha_{B,1}}, 0, 1 \right) \\ & \left(0, \frac{N_A}{\alpha_{A,1}}, 0, \frac{N_B}{\alpha_{B,2}}, 1 \right) \end{aligned}$$

Let $y^{(i,j)}(N_A, N_B)$, the quantity of each ray, where i is the subtype A agent, j is the subtype B agent. Therefore, we can define the set of $\{x_{T,s}[d_T(N_A, N_B)], y(N_A, N_B)\}$ that satisfies (51) as:

$$\begin{aligned} \{x_{T,s}[d_T(N_A, N_B)], y(N_A, N_B)\} = & [y^{(1,1)}(N_A, N_B)] \left(\frac{N_A}{\alpha_{A,1}}, 0, \frac{N_B}{\alpha_{B,1}}, 0, 1 \right) + \\ & \dots + [y^{(2,2)}(N_A, N_B)] \left(0, \frac{N_A}{\alpha_{A,2}}, 0, \frac{N_B}{\alpha_{B,2}}, 1 \right) \end{aligned}$$

Where $y^{(i,j)}(N_A, N_B) \geq 0$, $i = 1, 2$ and $j = 1, 2$. Intuitively, each ray is a different composition of types of agents to fulfill the contract, for example $y^{(1,1)}(N_A, N_B)$ corresponds to the measure of platforms

which are fulfilled by agents $(A, 1)$ and $(B, 1)$. There are four extreme rays hence a linear combination of these four rays is able to replicate any combination of types of agents. In general, if there are I types of A and J types of B then there will be $I \times J$ extreme rays for each contract.

Furthermore, we have the following relations:

$$\begin{aligned} x_{A,i}[d_T(N_A, N_B)] &= \sum_j \frac{y^{(i,j)}}{\alpha_{A,i}} N_A \\ x_{B,j}[d_T(N_A, N_B)] &= \sum_i \frac{y^{(i,j)}}{\alpha_{B,j}} N_B \\ y(N_A, N_B) &= \sum_{i,i'} y^{(i,j)}(N_A, N_B) \end{aligned}$$

Hence, we are now ready to redefine the Pareto problem in terms of our new definitions that satisfy the matching constraints (to reduce notation and clarity, we normalize all $\alpha_{T,s} = 1$ for all subtypes).

$$\begin{aligned} \max_{y^{(i,j)}(N_A, N_B) \geq 0} \sum_i \lambda_{A,i} &\left[\sum_j \sum_{(N_A, N_B)} y^{(i,j)}(N_A, N_B) \times N_A \times U_A(N_A, N_B) \right] + \\ &+ \sum_j \lambda_{B,j} \left[\sum_i \sum_{(N_A, N_B)} y^{(i,j)}(N_A, N_B) \times N_B \times U_B(N_A, N_B) \right] \end{aligned}$$

Such that each agent is assigned to a platform with probability one (the counterpart to equation (14)).

$$\sum_j \sum_{(N_A, N_B)} y^{(i,j)}(N_A, N_B) N_A = 1 \quad \forall i, \quad (52)$$

$$\sum_i \sum_{(N_A, N_B)} y^{(i,j)}(N_A, N_B) N_B = 1 \quad \forall j \quad (53)$$

Such that the resource constraint is satisfied (the counterpart to equation(16)):

$$\sum_{(N_A, N_B)} \left[\sum_{i,j} y^{(i,j)}(N_A, N_B) \times C(N_A, N_B) \right] \leq \sum_{T,s} \kappa_{T,s} \quad (54)$$

The advantage of writing the Pareto problem in the above formulation is that it reduces the constraint set, in this example, there are only five constraints, however, the number of variables is very large and we can use a linear programming solver to compute the reformulated Pareto program.

To calculate the prices paid by each agent we use the shadow prices (duals) from the reformulated problem and the economy's budget constraint.

The first-order conditions for this reformulated problem are:

$$\gamma^\kappa C(N_A, N_B) + \gamma^{A,i} N_A + \gamma^{B,j} N_B \geq N_A \lambda_{A,i} U_A(N_A, N_B) + N_B \lambda_{B,j} U_A(N_A, N_B) \quad (55)$$

where equation (55) holds with equality for those platforms that exist in equilibrium. The variables, γ^κ , $\gamma^{A,i}$, and $\gamma^{B,j}$ are the Lagrange multipliers associated with the resource constraint (equation 54), and the matching constraints (equations 52 and 53) respectively.

Recall that for all platforms that exist (that is, $y(N_A, N_B) > 0$), then the sum of prices paid for the platform must equal the costs of producing the platform. That is,

$$C(N_A, N_B) = p[d_A(N_A, N_B)] * N_A + p[d_B(N_A, N_B)] * N_B \quad (56)$$

We can use equations (55) and (56) to solve for the price paid by each agent for all platforms that are created in equilibrium $y^{i,j}(N_A, N_B) > 0$. Solving this set of equations gives the following prices for each slot in a platform:

$$p^{T,s}(N_A, N_B) = \frac{\lambda_{T,s} U_T(N_A, N_B) - \gamma^{T,s}}{\gamma^\kappa} \quad (57)$$

Where $p^{T,s}(N_A, N_B)$ is the equilibrium price paid by an agent of subtype T, s to join a platform of size (N_A, N_B) . Notice that the price function varies by subtype, yet, in equilibrium, if agents of the same type, but different subtype, join the same platform they will still pay the same price. For example, if $x_{A,1}[d_A(N_A, N_B)] > 0$ and $x_{A,2}[d_A(N_A, N_B)] > 0$ for some $[d_A(N_A, N_B)]$ then:

$$p^{A,1}[d_A(N_A, N_B)] = p^{A,2}[d_A(N_A, N_B)] \iff (\lambda_{A,1} - \lambda_{A,2}) U_A(N_A, N_B) = (\gamma^{A,1} - \gamma^{A,2}) \quad (58)$$

Therefore, the difference in the weighted utility between the different subtypes must be equal to the difference in the Lagrangian multipliers associated with each agent's participation constraint.