

Anonymity and stability in large many-to-many markets

Karolina Vocke*

April 29, 2022

Abstract

In many-to-many matching markets, various stability concepts have been introduced. Not all of these stability concepts offer a clear interpretation. This paper argues that the differences between stability concepts reflect different implicit anonymity assumptions. Such anonymity assumptions can best be modeled in large markets, described in this paper with a continuum of agents. In such large markets, it is shown that various differences between stability concepts disappear. In particular, stability and weak setwise stability coincide. Stability is a better behaved solution concept; stability blocks do indeed lead to an improvement for all members of a blocking coalition, unlike in finite markets. Moreover, the relationship between anonymity and largeness of the market can be made explicit in natural non-cooperative foundations of stability concepts.

1 Introduction

In complex markets, where multiple contracts can be signed on both sides or agents interact along with complex networks that need not be two-sided, as in supply chains or trading networks, there is no obvious right notion of stability. While in two-sided one-to-one matching markets stability simply requires that there be no mutually desirable unsigned “blocking” contract between any pair of agents, extending this idea to many-to-many markets has proven to be nontrivial.

In many markets with buyers on the one side and sellers on the other, it can be natural to allow market participants to sign a whole set of contracts with the other side. Here, an outcome can be blocked by a coalition that could principally keep or drop subsets of existing contracts, while signing new contracts. Forming a coalition in such a many-to-many market requires complex decisions regarding the treatment of new and existing contracts for coalition members. While in one-to-one matching markets signing a new contract requires an agent to drop the existing one, in a more general setting with multiple contracts between multiple agents this is not necessarily the case. Moreover, blocking coalitions could also agree on signing sets of new contracts instead of a single contract. The option of agents to unilaterally drop contracts can restrict which set of new contracts a coalition is willing to sign. Various stability concepts have been introduced in the literature that reflect these differences. These stability concepts include pairwise stability (Gale and Shapley (1962)), stability (Roth, 1984; Hatfield and Milgrom, 2005), versions of setwise stability (Sotomayor, 1999; Echenique and Oviedo, 2006; Klaus and Walzl, 2009), group stability, (Konishi and Ünver, 2006), tree stability (Ostrovsky, 2008; Jagadeesan and Vocke, 2021) versions of the core, path-stability, chain stability (Ostrovsky, 2008; Hatfield, Kominers, Nichifor, Ostrovsky, and Westkamp, 2021), and versions of trail stability (Fleiner, Jankó, Tamura, and Teytelboym, 2018) among others. Given the variety of stability concepts and settings analysed, there is no global consensus regarding the choice of stability concepts.

*Department of Economics, University of Innsbruck; and Lincoln College, University of Oxford. Email: karolina.vocke@gmx.de.

This paper argues that the differences between stability concepts reflect different implicit anonymity assumptions, which can be best modeled in large markets. In a large anonymous market, many differences between stability concepts cease to matter. Since stability notions are defined by what counts as a block, we can classify stability concepts by the definition of blocks. Blocking coalitions in many-to-many markets can differ in three dimensions: The treatment of existing contracts, the treatment of new contracts, and the shape of a block. The main theorem shows, that several of the various solution concepts do indeed coincide in large markets; Particularly the dimension of how to treat existing contracts collapses. It is shown that weak setwise stability and stability coincide. In many-to-many matching markets dropping existing contracts within a coalition can impact other coalition members. As a consequence, in a finite market, it is possible that members of a blocking coalition can not jointly improve. This can never happen in a large market; stability blocks do indeed lead to an improvement for all agents. A novel non-cooperative foundation for different versions of stability is provided in the appendix to make the relation between stability and anonymity explicit.

Literature review. This paper draws heavily on the concepts and classification in Klaus and Walzl (2009), which introduces different versions of setwise stability in finite markets, and discusses the relationship between them in depth. Bando and Hirai (2021) recently analysed the relationship between, among other notions, setwise stability, stability, and efficiency. They put strong restrictions on the network topology and show the coincidence of many stability notions in acyclic networks. This paper, in contrast, allows for a general network but uses the largeness of markets to obtain the coincidence of stability and setwise stability concepts. Large markets are a natural way of modelling anonymity. A recent strand of the matching literature, starting with Azevedo and Leshno (2016), has used large markets to analyse the properties of stable matchings. The largeness of markets has particularly been useful to guarantee existence of stable outcomes. Azevedo and Hatfield (2018) shows the existence of stable outcomes in large, two-sided many-to-many markets when imposing a substitutability condition on one side. Che, Kim, and Kojima (2019) shows the existence of stable outcomes in large many-to-one markets with complementarities; Greinecker and Kah (2020) uses large markets to show existence of stable outcomes in two-sided one-to-one matching markets with externalities. Jagadeesan and Vocke (2021) shows existence of tree-stable outcomes in large many-to-many matching markets. To guarantee existence, it is necessary to impose restrictions on the network, blocking shapes of coalitions, or the preferences. This paper, in contrast, uses the largeness of the market for a better understanding of solution concepts, particularly for analysing relationships of stability notions. Thus, a general framework can be used. In terms of microfoundations, Konishi and Ünver (2006) also provided a microfoundation for pairwise stability and a stability concept called “credible group stability.”. The game in the appendix is most closely related to the offer game given in Jagadeesan and Vocke (2021) microfounding a stability concept called tree-stability.

The remainder of the paper is organized as follows. Section 2 explains the main results and contributions with examples. Section 3 introduces the formal model of matching in large networks. Section 4 defines solution concepts formally, and analyses the general relationship between them. Section 5 contains the main theorem, stating that weakly setwise stable and stable outcomes coincide in large markets, and sketches the idea of the proof.

2 Contribution

This section explains the main results and ideas with examples. Solution concepts in many-to-many markets differ along three dimensions. This paper focuses on how members of blocking coalitions treat existing contracts in the various solution concepts. In particular, the treatment of existing contracts with agents within the coalition (setwise stability versus stability) and the treatment of existing contracts with agents outside the coalition (core versus stability). Furthermore, solution concepts differ in how members of blocking coalitions treat new contracts (strong versus weak setwise stability) and the shapes of




Figure 1: Contracts in our examples. We let a, b, c denote the contracts between i_1 and i_2 .

feasible blocking structures (pairwise-, path-, tree- stability etc.).

Stability is a solution concept that allows for the independent dropping of existing contracts with agents inside and outside the coalition. A stability block consists of a coalition and a set of new contracts between the members of the coalition that are part of the best choice (given existing and new contracts) for every agent in the blocking coalition. Existing contracts can be independently cancelled even within the coalition, possibly leaving some members of the coalition worse off than before. This difference between cancelling contracts and adding contracts makes it hard to interpret stability in small markets, as can be seen in the following example.

Example 1. As depicted in Figure 6, there are two agents, and three contracts, and the agents' preferences are given by

$$\begin{aligned} i_1 : \{a, b\} &\succ \{a, c\} \succ b \succ \emptyset \\ i_2 : \{b, c\} &\succ \{a, c\} \succ b \succ \emptyset. \end{aligned}$$

Both agents like to sign contracts a and c . While b serves as a substitute to c for agent i_1 it serves as a substitute to contract a for agent i_2 . Both agents accept the single contract b .

- Both agents signing the contract $\{b\}$ is a stable outcome. In particular, the contracts $\{a, c\}$ will not be part of a block since $\{a, c\}$ is not part of a best choice given old and new contracts; i_1 would not sign c when a, b , and c are feasible.
- Both agents signing the contracts $\{a, c\}$ is not stable since it is blocked by the agents signing the new contract b . For both agents, there is a best choice given a, b and c that indeed includes b .

Strangely, the same agents within the coalition are treated differently when adding and cancelling contracts. An inefficient outcome (b) is stable because agents can 'foresee' that the possible block ($\{a, c\}$) will not be implemented. The efficient outcome ($\{a, c\}$) however is blocked with the contract b which cannot lead to any improvement for both agents.

Setwise stability in comparison treats existing contracts differently. A setwise block consists of a coalition and a set of new and old contracts within the coalition, such that the members of the coalition are better off when signing these contracts while keeping or dropping contracts with agents outside the coalition. Hence, setwise blocks in comparison to stability blocks necessarily lead to improvements for all members of the blocking coalition. In Example 1 the outcome where the contracts $\{a, c\}$ are signed is setwise stable since there is no set of contracts that is better for both agents. Both agents signing $\{b\}$, in comparison, is setwise blocked by the set of contracts $\{a, c\}$, which is better for both agents. In this example the sets of stable and setwise stable outcomes are disjoint. In setwise stability concepts, existing contracts within and outside the coalition are treated differently, as is illustrated in the following example:




Figure 2: Contracts in our examples. We let a, b, c denote the contracts between i_1 and i_2 and i_3 .

Example 2. As depicted in Figure 6 there are three agents and three contracts, and the agent's preferences are given by

$$\begin{aligned} i_1 : \{a, c\} &\succ \emptyset \\ i_2 : \{b, c\} &\succ \{a, c\} \succ b \succ \emptyset \\ i_3 : b &\succ \emptyset. \end{aligned}$$

Intuitively, i_1 has complementary preferences and only wants to sign $\{a, c\}$. Agent i_2 has the same preferences as in Example 1.

- Agents i_2 and i_3 signing contract $\{b\}$ is not a setwise stable outcome since it is blocked with the contracts $\{a, c\}$. On the contrary, this outcome is still stable since the contracts $\{a, c\}$ are not part of a best choice for i_2 given a, b , and c .
- Agents i_1 and i_2 signing contracts $\{a, c\}$ is neither a setwise stable outcome nor a stable outcome since signing the contract b (while i_2 drops the contract a) leads to an improvement for both coalition members i_2 and i_3 .

Comparing Example 1 and Example 2 shows that setwise stability can lead to different results in very similar settings. While i_2 is facing the same issue in both examples- her best choice is not compatible with the preferences of agent i_1 , in Example 1 the efficient outcome $\{a, c\}$ is setwise stable while in Example 2 it is not. Intuitively, agent i_2 is treating agent i_1 differently depending on whether agent i_1 is inside or outside the coalition. However, in both examples, starting from $\{a, c\}$, signing the new contract b and dropping the existing contract a will lead to a new outcome that is not acceptable for i_1 ; agent i_1 would in turn like to unilaterally drop the contract c .

Individual rationality is a solution concept that allows for dropping existing contracts unilaterally. In Examples 1 and 2 only the outcomes with $\emptyset$, b or $\{a, c\}$ being signed are individually rational. Individual rationality is a minimal requirement satisfied by most stability concepts, including stability and setwise stability.

The *core*, in contrast, is a solution concept in which members of the blocking coalition cannot keep any existing contracts with agents outside the coalition. A core block is a coalition that can improve by leaving the market and signing new contracts within the coalition; no contracts outside the coalition can be kept.¹ The only core outcome in both examples is the outcome where agents i_1 and i_2 sign the contracts $\{a, c\}$. While in Example 1 core and setwise stability coincide, Example 2 shows that setwise stability and the core are generally different concepts. Core outcomes are always efficient. However, outcomes in the core need not be individually rational. Thus, the core is not suitable in a matching market in which agents can independently cancel contracts.

¹More precisely one distinguishes between the weak core demanding strict improvement for all agents and the strict core demanding weak improvement for all agents.

Another dimension solution concepts differ in is the treatment of new contracts. This gives rise to different versions of setwise blocks and stability blocks. A *weak setwise block* consists of a coalition and a set of old and new contracts that is an improvement for all coalition members. On the other hand, a *strong setwise block* consists of a coalition and a set of old and new contracts that is an improvement and an optimal choice from the set of existing and new contracts for each coalition member. Analogously, there are two types of stability (*strong group stability* and stability).² In Example 1, both agents signing $\{b\}$ is not a strong setwise stable outcome because $\{a, c\}$ is part of a weak setwise block, but is weakly setwise stable because the contracts $\{a, c\}$ are not part of a strong block, since they are not part of the best possible choice given a, b and c . Similarly, both agents signing $\{b\}$ is a stable outcome but is not strong group stable; there exists a strong group block where the contracts $\{a, c\}$ are signed. On the other hand, both agents signing contracts $\{a, c\}$ is weakly and strongly setwise stable, but not stable and not strong group stable.

The third dimension in which the solution concepts discussed in this paper differ is the allowed shape of possible blocks. While a *pairwise block* consists of a single contract, *path-*, *chain-* and *tree-* stable outcomes rely exactly on path- chain- or tree- shaped blocks.

This paper argues that the different treatments of contracts that characterize the different solution concepts are best interpreted as reflecting different implicit anonymity assumptions. Going back to Example 1: Why is the outcome where both agents sign contract $\{b\}$ not blocked by the contracts $\{a, c\}$? The contracts $\{a, c\}$ are an improvement but not part of the best possible choice for both agents; agents take account of the fact that their counterparts will not sign both contracts a and c . Implicitly, the agents whom they add contracts with are treated as *non-anonymous* (similarly in weak setwise stability). Why then is $\{a, c\}$ blocked by b ? Why do agents agree to a block which cannot lead to an improvement? Because agents *do not* take account of the fact that their counterpart will not keep the contract a resp. c . Implicitly, they treat the agents whom they are cancelling contracts with as *anonymous* (in comparison to weak setwise stability). Therefore, stability requires that agents treat counterparties as anonymous when cancelling contracts and as non-anonymous when adding contracts.³ This seems strange if they cancel and add contracts with the same agent, as discussed in Example 1. No such problem would occur if all new contracts are signed between agents that have no existing contracts between them. The most natural setting for modelling anonymous interactions is a large market. But as this paper shows, it is precisely in such large markets that the problem just mentioned does not occur. Not only are the implicit anonymity assumptions explicitly modeled, but also stability blocks indeed lead to an improvement for all members of the blocking coalitions.

Economically, the relevant property of large markets is, that each agent is anonymous. In this paper, large markets are modelled by a continuum of agents of each type. Theorem 1 shows that weak setwise stability and stability do coincide in large markets. The idea is that whenever there is a blocking coalition in a large market, there are substitutes for the coalition members that do not have any existing contracts signed between them. This can be seen in Example 1: Assume there is a continuum of agents of each type. Starting from the outcome where contracts $\{a, c\}$ are signed by all agents, an agent of type i_1 can cancel the contract c and sign contract b with an independent agent of type i_2 who is cancelling a with another agent of type i_1 in turn. Hence we can find a new outcome that indeed leads to an improvement for both agents in the blocking coalition, showing that there is not only a stability block but also a weak setwise block in Example 1. On the other hand, signing contracts $\{a, c\}$ in Example 1 is no longer a setwise stable outcome in the large market analogously to Example 2. Thus, setwise stability is better

²While a strong group block only requires that new contracts can lead to an improvement, a stability block requires the set of new contracts being part of their best possible choice, given the existing and new contracts for all coalition members.

³More precisely: The different treatments of new contracts reflect the level of anonymity within the blocking coalition. While the core assumes no independent anonymous contracts outside the coalition that can be kept, it assumes a high level of anonymity inside the coalition. Stability assumes anonymity when cancelling and non-anonymity when adding contracts, weak setwise stability assumes anonymity outside and non-anonymity within the coalition.

behaved in large markets too.

The relation between stability and largeness of the market is made explicit in a noncooperative game-theoretic foundation of stability in large markets, which is provided in the appendix. The specific assumptions of stability require possible non-anonymous interactions within the coalition. Each outcome gives rise to a dynamic game with almost perfect information in which agents can propose blocking coalitions at some small cost. They cannot target specific agents but only types of agents, reflecting the anonymity in large markets. Theorem B.1 shows that an outcome is stable if and only if no agent participates in the game in any subgame perfect equilibrium. The anonymity assumptions of stability can be interpreted in a large market as follows: There still could arise some kind of non-anonymity within the blocking coalition, if during the deviation process coalition members can defect from the proposed set of contracts. Ex-ante, a single agent will be affected by the offer game with probability 0, but within the process of deviation involved agents are affected by the strategies of the other coalition members with positive probability. This explains why when adding contracts, agents within a coalition are treated as non-anonymous while when cancelling contracts, agents (outside the coalition) are treated as anonymous. Dropping the option of defection in the game provides a non-cooperative foundation for strong group stability. The different anonymity assumptions underlying the treatment of new contracts in these solution concepts are reflected in the option of defecting within the deviation process.⁴

3 Formal model

The basic framework is directly taken from Jagadeesan and Vocke (2021) which, in turn, extends the two-sided large market matching model in Azevedo and Hatfield's (2018) to networks. There are finitely many types with a continuum of agents of each type. Agents interact bilaterally with multiple counterparties via an exogenously specified finite set of contracts.

3.1 Agents, contracts, and preferences

There is a finite set I of types. For each type $i \in I$, there is a (homogenous) mass $\theta^i > 0$ of agents of type i , parameterized by an interval $[0, \theta^i]$. The set of agents of type i is $\Theta_i = \{i\} \times [0, \theta^i]$, and hence the set of all agents is

$$\Theta = \bigcup_{i \in I} \Theta_i.$$

Let μ denote the measure on Θ whose restriction to each space Θ_i is given by the Lebesgue measure. For each pair i, j of types, there is a finite set $X_{i,j} = X_{j,i}$ of contracts between an agent of type i and an agent of type j . Pairs of agents can sign multiple contracts between them (Fleiner, 2003; Ostrovsky, 2008; Kominers, 2012; Hatfield and Kominers, 2017). Contracts are type-specific: $X_{i,j} \cap X_{i',j'} = \emptyset$ when $\{i, j\} \neq \{i', j'\}$. The set of feasible contracts for an agent of type i is

$$X_i = \bigcup_{j \in I} X_{i,j}.$$

For each type i , there is an injective utility function $u^i : \mathcal{P}(X_i) \rightarrow \mathbb{R}$ defined over the sets of contracts that agents of type i can sign; u^i being injective means that preferences are strict. Given a set $X' \subseteq X_i$ of contracts, let

$$C^i(X') = \arg \max_{Y \subseteq X'} u^i(Y).$$

⁴ Anonymity assumptions can also be used to justify restrictions on the feasible shape of blocks (see Section 6 in Jagadeesan and Vocke (2021)).

3.2 Outcomes

An outcome specifies which set of contracts each agent signs; these sets must be compatible across agents. Formally, a *matched type* consists of a type $i \in I$ and a set $Y \subseteq X_i$ of contracts. An outcome consists of a (measurable) set $M_Y^i \subseteq \Theta_i$ of agents of type i that participate in set Y for each matched type (i, Y) such that (1) each agent is associated to exactly one set of contracts, and (2) contracts are signed by equal masses on either side.

Definition 1. An *outcome* M consists of a measurable subset $M_Y^i \subseteq \Theta_i$ for each matched type (i, Y) such that

- (feasibility) the sets M_Y^i are disjoint and satisfy

$$\bigcup_{Y \subseteq X_i} M_Y^i = \Theta_i$$

- (reciprocity) for each pair of distinct types $i \neq j$ and each contract $x \in X_{i,j}$, we have that

$$\mu \left(\bigcup_{x \in Y \subseteq X_i} M_Y^i \right) = \mu \left(\bigcup_{x \in Y \subseteq X_j} M_Y^j \right).$$

In Definition 1, feasibility requires that exactly one set of contracts be specified for each agent. In the reciprocity condition,

$$\bigcup_{x \in Y \subseteq X_i} M_Y^i$$

is the set of agents of type i that sign contract x .

Note that outcomes do not specify matches between agents; all interactions are mediated by contracts. For an agent $a \in \Theta$, let the set of contracts Y that she signs in the outcome M be $M(a)$, i.e. $M(a) = Y$ for some Y s.t. $a \in M_Y^i$.

4 Solution concepts

This section gives formal definitions of the stability concepts introduced in Section 2, defines relevant notation, and discusses the relationship between the concepts in depth.

Every stability concept is characterised by a blocking structure consisting of a coalition and a set of new contracts. In one-to-one matching markets, where agents can only sign a single contract between them, only pairwise blocks can arise in which two agents sign a new contract and cancel their old contract. However, in many-to-many markets more complex blocks are possible. In a one-to-one market, the counterpart of an agent that gets a contract cancelled by her cannot have any further impact on her, a property referred to as *anonymity* in this paper. In many-to-many markets, in comparison, she can have a further impact, due to the possibility of signing several contracts. Thus, whether agents in a coalition have an incentive to deviate from an outcome given a set of new contracts and how existing contracts are treated by coalition members can be interpreted in various ways, giving rise to various solution concepts.

First, we discuss stability concepts that allow agents in a blocking coalition to keep and drop existing contracts with agents inside and outside the coalition independently; these are individual rationality, strong group stability, stability, tree- and pairwise stability. Next, we discuss stability concepts that allow agents in a blocking coalition to keep and drop existing contracts with members outside the coalition independently but restrict the treatment of contracts with members inside the coalition.; these are the

different versions of setwise stability.⁵ All these solution concepts are adapted from the corresponding finite market concepts to large markets.

First, we define two notions of efficiency.

Definition 2. An outcome M is *weakly Pareto efficient* if there is no other outcome M' , such that almost every agent $a \in \Theta$ strictly prefers $M'(a)$ to $M(a)$. M is *strictly Pareto efficient* if there is no other outcome M' , such that almost every agent $a \in \Theta$ weakly prefers $M'(a)$ to $M(a)$ and a set of agents with positive measure strictly prefer the new outcome.

As in finite markets, efficient outcomes always exist. However, in matching markets an efficient outcome cannot always be maintained if agents are able to unilaterally cancel contracts, as will be seen in Example A.1.3.

The main difference between stability concepts, this paper focuses on, is the treatment of existing contracts. In some concepts, existing contracts can be independently dropped and only the new contracts are specified in a block. In comparison in setwise stability concepts both, new and existing contracts within the coalition need to be specified. To understand this need, we distinguish between a notion of preferability and desirability: A set W is *preferred* to a set Y if changing from Y to W is an improvement. In comparison, a set Z is *desired* from Y , if starting from Y adding the set of contracts Z can lead to an improvement. While a set W that is preferred to Y simply increases the utility, a *strongly preferred* set W is the best choice given W and Y . Formally, let $W, Y \subseteq X_i$. Then W is

- *preferred* to Y by type i if $u^i(W) > u^i(Y)$.
- *strongly preferred* to Y by type i if $W = C_i(Y \cup W)$.

Now, let $Z, Y \subseteq X_i$. We say Z is

- *desired* from Y by type i if there exists a set $W \subseteq Y \cup Z$ with $Z \subseteq W$ that is preferred to Y by i .
- *strongly desired* from Y by type i if there exists a set $W \subseteq Y \cup Z$ with $Z \subseteq W$ that is strongly preferred to Y by i .⁶

4.1 Stability

In this section stability concepts are introduced that allow agents in a blocking coalition to keep and drop existing contracts with agents inside and outside the coalition independently. A *block* consists of a coalition and a set of new contracts that is desired by all members of the blocking coalition from the existing set of contracts.

All stability concepts require *individual rationality*, i.e. the option to drop contracts unilaterally. Formally:

Definition 3. An outcome M is *individually rational* if $\mu(M_Y^i) = 0$ for each matched type (i, Y) such that $Y \notin C^i(Y)$.

Obviously, the outcome in which no contract is signed is always individually rational. The simplest stability concept is pairwise stability, due to Gale and Shapley (1962), which requires individual rationality and that there be no mutually desired unsigned “blocking” contract from an outcome between any pair of agents. Every stability concept refines pairwise stability by allowing additional blocking possibilities. Different versions of stability differ in the treatment of new contracts (strongly desirable/

⁵Moreover, in the appendix we discuss the core, a solution concept that doesn’t allow agents in a blocking coalition to keep any contracts with agents outside the coalition at all.

⁶Note that, if W is (strongly) preferred to Y then $W \setminus Y$ is (strongly) desired from Y .




Figure 3: Summary of the relationships between the stability concepts of this subsection. The arrows represent relationships between stability concepts.

desirable), and in the shape of blocks (pairwise- blocks/ blocks of any shape).

The definition of a block is taken from Jagadeesan and Vocke (2021).

Definition 4. Let ν be a graph with $n > 1$ vertices. A *block of shape ν* consists of

- a matched type (i_j, Y^j) for each $1 \leq j \leq n$, and
- a contract $x_{j,k} = x_{k,j} \in X_{i_j, i_k} \setminus (Y^j \cup Y^k)$ for each $\{j, k\} \in \nu$

for which

- [compatibility] for each index j , the contracts $(x_{j,k})_{\{j,k\} \in \nu}$ are distinct;
- [desirability] writing

$$Z^j = \{x_{j,k} \mid \{j, k\} \in \nu\},$$

we have that Z^j is strongly desired from Y^j by i_j .

A block *arises* at an outcome $(M_Y^i)_{i,Y}$ if $\mu(M_{Y^j}^{i_j}) > 0$ for all j .

Strong group blocks are defined analogously with contracts being desired instead of strongly desired.

Here, compatibility requires that no agent participates in the same contract multiple times during a block. We can now introduce the formal definitions of the stability concepts.

Definition 5. An outcome is:

- *pairwise-stable* if it is individually rational and no block of the shape of a graph with two nodes and one edge arises.
- *stable* if it is individually rational and no block of any shape arises.⁷
- *strong group stable* if it is individually rational and no strong group block of any shape arises.

Since every strongly desired set of contracts is also desired, strong group stability implies stability. Furthermore, each stable outcome is clearly pairwise stable, and each pairwise stable outcome is individually rational by assumption, see Figure 3. In general, a stable outcome need not be strong group stable, as seen in Example 1, and a pairwise stable outcome need not be stable, as will be seen in Example A.1.4. These relations hold also in finite markets and parallel the results on setwise stability concepts in Klaus and Walzl (2009).

While stable outcomes (and hence strong group stable outcomes) are not guaranteed to exist even in large markets, see Azevedo and Hatfield (2018), pairwise stable outcomes always exist as a corollary to the main existence result in Jagadeesan and Vocke (2021). These facts are summarized in the following proposition.

Proposition 1.

⁷While formally multiple contracts between pairs of agents in a block aren't allowed, this restriction does not affect the definition of stability in this large-market model as shown in Jagadeesan and Vocke (2021).

- a. Every strong group stable outcome is stable and every stable outcome is pairwise stable. None of these implications can be reversed.
- b. Stable outcomes may not exist even in large markets.
- c. Pairwise stable outcomes generally exist in large markets.

4.2 Setwise stability

In this section, setwise stability concepts are introduced, that allow agents in a blocking coalition to keep and drop existing contracts with members outside the coalition independently but restrict the treatment of existing contracts with members inside the coalition. Formalizing setwise stability concepts requires a richer notion of a block. In particular, a different new outcome is specified.

Definition 6. Let ν be a graph with $n > 1$ vertices. A (strong) *setwise block of shape ν* of an outcome M consists of:

- an outcome M' ,
- a coalition, given by: a type i_j and a set of agents $A(j) \in \Theta_{t(j)}$ for each $1 \leq j \leq n$ with $\mu(A(j)) = \kappa$ for some constant $\kappa > 0$,
- a contract $x_{j,k} = x_{k,j} \subset X_{i_j, i_k}$ for each $\{j, k\} \in \nu$. Let $x(j) = \cup_{\{j,k\} \in \nu} x_{j,k}$,

such that:

- for all agents $a \in \Theta$ with $M'(a) \not\subseteq M(a)$ it holds that $a \in A(j)$ for some $1 \leq j \leq n$ and $M(a) \cap x(j) = \emptyset$ and $M'(a) = Y(a) \cup x(j)$ for some $Y(a) \subseteq M(a)$,
- for all agents $a \in A(j)$ it holds that $M'(a)$ is (strongly) preferred to $M(a)$ by i_j ,
- for all agents $a \notin A$ holds $M'(a) \subseteq M(a)$.

Here, the first property requires that all new contracts are among the members of the blocking coalition only, and that the outcome M' contains no new contracts besides the contracts defined in the blocking coalition. The second property requires that all members of the blocking coalition receive a better set of contracts. The third property requires that agents outside the blocking coalition do not receive new contracts. However, it is allowed that some of the contracts of agents outside the blocking coalition are dropped by members of the blocking coalition.⁸

We can now introduce the formal definitions of the setwise stability concepts.

Definition 7. An outcome M is *weakly setwise stable*, if it is individually rational and no strong setwise block of the outcome M of any shape exists. It is *strongly setwise stable* if it is individually rational and no weak setwise block of the outcome M of any shape exists.

Since every strongly preferred set of contracts is also preferred, strong setwise stability implies weak setwise stability, as illustrated in Figure 4.

The main result of this paper shows that all stability concepts coincide with their setwise counterparts in large markets. Hence, a result exactly parallel to Proposition 1 holds for the setwise versions of the stability concepts. These relations between setwise stability concepts in finite markets have already been shown in Klaus and Walzl (2009).

⁸Intuitively, contracts are only dropped by agents of the blocking coalition, however, technically it doesn't make a difference whether other agents do drop independently other contracts, as long as the blocking coalition isn't affected.




Figure 4: Summary of the relationships between the stability concepts of this subsection. The arrows represent relationships between stability concepts.




Figure 5: Summary of the relationships between stability concepts. The regular arrows represent obvious relationships between stability concepts, the dashed arrows relationships that rely on markets being large.

Since setwise blocks not only consider the desirability of agents but also require a restriction on the treatment of existing contracts between agents within the blocking coalition, setwise blocks are a stronger requirement than stability blocks. Whenever a weak setwise block exists at an outcome, a block arises. Therefore, every stable outcome must be weakly setwise stable. The converse is only true in large markets but is highly nontrivial and will be discussed in Section 5. For finite markets, the converse fails even for pairwise-stability as seen in Example 1; the difference does not rely on the shape of blocks.

5 Role of anonymity for stability

This section contains the first main contribution of the paper, which is showing that stability and weak setwise stability coincide in large markets. In particular, whenever a block arises, there exists also a weak setwise block. Consequently, if an outcome is not stable it can be blocked by a coalition in a way that makes the coalition members better off. This is formalized in the following theorem and corollary.

Theorem 1. *Weak setwise stability and stability coincide and strong setwise stability and strong group stability coincide.*

As already discussed in section 4, the direction that every stable outcome is weakly setwise stable follows directly from the definition. The other direction is the nontrivial part of Theorem 1, we state it explicitly at the level of coalitions in the following corollary.

Corollary 1. *If a strong group block arises at an outcome, then there exists an outcome that is better for (a positive measure of) agents of the corresponding matched types.*

Since every block is a strong group block, Corollary 1 can also be applied to blocks. In Figure 5 the results of Theorem 1 are summarized together with the implication that holds in finite markets.

The logic of Theorem 1 and the proof idea can be illustrated with Example 1 from Section 2.

Example 1. [revisited]




Figure 6: Contracts in Example 1.

As depicted in Figure 6, there are two agents and three contracts. Types' preferences are given by

$$\begin{aligned} i_1 : \{a, b\} &\succ \{a, c\} \succ b \succ \emptyset \\ i_2 : \{b, c\} &\succ \{a, c\} \succ b \succ \emptyset. \end{aligned}$$

Both types like to sign contracts a and c . While b serves as a substitute to c for type i_1 it serves as a substitute to contract a for type i_2 . Both types accept the single contract b .

In a finite market in which there is only one agent of each type, the set of weakly setwise stable outcomes and the set of stable outcomes are different, as discussed in Section 2. In particular, both agents signing the contracts $\{a, c\}$ is a weakly setwise stable outcome but not stable, since it is blocked with contract b . Here both agents are part of a block though no improvement for both is possible. Intuitively, agents treat each other as anonymous when cancelling contracts. Conversely, in the outcome, where both agents sign contract $\{b\}$ no block exists, even though signing the contracts $\{a, c\}$ could lead to an improvement for both agents; they treat each other as non-anonymous when adding contracts. This asymmetry disappears in large markets.

In a large market with a continuum of agents, in comparison, every block gives rise to a weak setwise block. Intuitively, there are enough agents in a large market that a coalition can always be chosen so that there are no existing contracts between coalition members. Hence, existing contracts are only cancelled with agents outside the coalition; the cancellation of existing contracts of the coalition members has no impact on other coalition members. Consequently, only the new contract choices affect coalition members and all coalition members can be made better off. In other words, the agents treated as anonymous and the agents treated as non-anonymous by coalition members do not need to be the same and the agents treated as non-anonymous can be chosen to be only agents outside the coalition.

In this example, with a unit mass of agents of each type, all agents signing contracts a and c is setwise blocked. For example, half of the mass of the agents can form a weak setwise block by signing the new contract b while dropping the existing contracts a (resp. c) in the remaining set of agents outside the blocking coalition.

An example of an outcome M that blocks the outcome in which all agents sign contracts a and c is the outcome defined by:

$$\begin{aligned} M_{\{c\}}^{i_1} &= \{i_1\} \times \left[0, \frac{1}{2}\right) & M_{\{a,b\}}^{i_1} &= \{i_1\} \times \left[\frac{1}{2}, 1\right] \\ M_{\{a\}}^{i_2} &= \{i_2\} \times \left[0, \frac{1}{2}\right) & M_{\{b,c\}}^{i_2} &= \{i_2\} \times \left[\frac{1}{2}, 1\right]. \end{aligned}$$

Here the agents in $\{i_1\} \times [\frac{1}{2}, 1]$ of type 1 and the agents in $\{i_2\} \times [\frac{1}{2}, 1]$ of type 2 are the coalition members that do, indeed, strictly improve in the new outcome M .

References

- Azevedo, E. M. and J. W. Hatfield (2018). Existence of equilibrium in large matching markets with complementarities. Working paper.
- Azevedo, E. M. and J. D. Leshno (2016). A supply and demand framework for two-sided matching markets. *Journal of Political Economy* 124(5), 1235–1268.
- Bando, K. and T. Hirai (2021). Stability and venture structures in multilateral matching. *Journal of Economic Theory* 196(C).
- Che, Y.-K., J. Kim, and F. Kojima (2019). Stable matching in large economies. *Econometrica* 87(1), 65–110.
- Echenique, F. and J. Oviedo (2006). A theory of stability in many-to-many matching markets. *Theoretical Economics* 1, 233–273.
- Fleiner, T. (2003). A fixed-point approach to stable matchings and some applications. *Mathematics of Operations Research* 28(1), 103–126.
- Fleiner, T., Z. Jankó, A. Tamura, and A. Teytelboym (2018). Trading networks with bilateral contracts. Working paper.
- Gale, D. and L. S. Shapley (1962). College admissions and the stability of marriage. *American Mathematical Monthly* 69(1), 9–15.
- Greinecker, M. and C. Kah (2020). Pairwise stable matching in large economies. Working paper.
- Hatfield, J. W. and S. D. Kominers (2017). Contract design and stability in many-to-many matching. *Games and Economic Behavior* 101, 78–97.
- Hatfield, J. W., S. D. Kominers, A. Nichifor, M. Ostrovsky, and A. Westkamp (2021). Chain stability in trading networks. *Theoretical Economics* 16(1), 197–234.
- Hatfield, J. W. and P. R. Milgrom (2005). Matching with contracts. *American Economic Review* 95(4), 913–935.
- Jagadeesan, R. and K. Vocke (2021). Stability in large markets. Technical report.
- Klaus, B. and M. Walzl (2009). Stable many-to-many matchings with contracts. *Journal of Mathematical Economics* 45(7), 422–434.
- Kominers, S. D. (2012). On the correspondence of contracts to salaries in (many-to-many) matching. *Games and Economic Behavior* 75(2), 984–989.
- Konishi, H. and M. U. Ünver (2006). Credible group-stability in many-to-many matching problems. *Journal of Economic Theory* 129, 57–80.
- Ostrovsky, M. (2008). Stability in supply chain networks. *American Economic Review* 98(3), 897–923.
- Roth, A. E. (1984). Stability and polarization of interests in job matching. *Econometrica* 52(1), 47–57.
- Sotomayor, M. (1999). Three remarks on the many-to-many stable matching problem. *Mathematical Social Sciences* 38(1), 55–70.

A The Core

This section introduces the core and compares it to the stability concepts discussed in the paper. The definition of a core block is also taken from Jagadeesan and Vocke (2021).

Definition A.1. A *strict core block of shape ν* consists of

- a matched type (i_j, Y^j) for each $1 \leq j \leq n$, and
- a set $W_{j,k} = W_{k,j} \subseteq X_{i_j, i_k}$ of contracts for each $1 \leq j, k \leq n$

for which

- [compatibility] for index j , the sets $(W_{j,k})_{1 \leq k \leq n}$ are pairwise disjoint;
- [preferability] writing

$$W^j = \bigcup_{k=1}^n W_{j,k},$$

we have that W^j is preferred to Y^j

Analogously a *weak core block* is defined with the preferability condition that $u^{i_j}(W^j) > u^{i_j}(Y^j)$ for all $1 \leq j \leq n$ replaced by the condition that $u^{i_j}(W^j) \geq u^{i_j}(Y^j)$ for all $1 \leq j \leq n$ with strict inequality for some j .

Such a core block *arises* at an outcome $(M_Y^i)_{i,Y}$ if $\mu(M_{Y^j}^{i_j}) > 0$ for all j .

We can now introduce the formal definitions of the core concepts.

Definition A.2. An outcome is in the *weak (strict) core* if no strict (weak) core block arises.

It is clear that every strict core outcome lies in the weak core. In general, the inclusion is strict; see Example A.1.1. Outcomes in the strict core are Pareto efficient; inefficient outcomes can be blocked by the grand coalition. Similarly, outcomes in the weak core are weakly Pareto efficient. Outcomes in the core generally are not individually rational; see Example A.1.3. While in finite markets the weak core may be empty,⁹ in large markets a weak core outcome always exists; see Azevedo and Hatfield (2018).

Proposition A.1. *Strong group stable outcomes lie in the strict core and are, consequently, Pareto efficient.*

However, stable outcomes can be inefficient; see Example A.1.3. The relations between core concepts, efficiency, and stability concepts are summarized in Figure 8. The following examples show that no implication besides the ones in Figure 8 holds.

Example A.1. 1. (core versus efficiency) As depicted in Figure 7(a), there are three types and two contracts. Types' preferences are given by

$$i_1 : a \succ \emptyset \quad i_2 : \{a, b\} \succ a \succ \emptyset \quad i_3 : b \succ \emptyset.$$

Agents of type i_1 and i_2 signing contract $\{a\}$ is an individually rational outcome and is weakly efficient but not efficient, it is in the weak core but not in the strict core, and not (strong group) stable.

⁹See, for example, the roommates' problem in Azevedo and Hatfield (2018).




Figure 7: Contracts in Example A.1.




Figure 8: Summary of the relationships between stability, efficiency, and the core concepts. An arrow from A to B means that each outcome of type A is also of type B . No further implications hold in general.

2. (core versus efficiency) As depicted in Figure 7(a), there are two types, and two contracts as before. Now, types' preferences are given by

$$i_1 : a \succ \emptyset \quad i_2 : b \succ a \succ \emptyset \quad i_3 : b \succ \emptyset.$$

Agents of type i_1 and i_2 signing contract $\{a\}$ is an individually rational outcome, and is (weakly) efficient but is not in the (weak/strict) core, and not (strong group) stable.

3. (Individual rationality versus core) As depicted in Figure 7(b), there are three types, and two contracts. Types' preferences are given by

$$i_1 : a \succ \{a, b\} \succ \emptyset \quad i_2 : \{a, b\} \succ \emptyset.$$

Agents of type i_1 and i_2 signing contracts $\{a, b\}$ is not an individually rational outcome but is (weakly) efficient, and is in the (weak/strict) core.

Agents of type i_1 and i_2 signing no contract is an individually rational outcome but is not (weakly) efficient, is not in the (weak/strict) core and not strong group stable, but is stable.

4. (stable versus pairwise stable) As depicted in Figure 7(b), there are two types, and two contracts as before. Now, types' preferences are given by

$$i_1 : \{a, b\} \succ \emptyset \quad i_2 : \{a, b\} \succ \emptyset.$$

Agents of type i_1 and i_2 signing contracts $\{a, b\}$ is an individually rational outcome, is (weakly) efficient, and is in the (weak/strict) core.

Agents of type i_1 and i_2 signing no contract is an individually rational outcome but is not (weakly) efficient, is not in the (weak/strict) core, and not (strong group) stable but pairwise stable.

B Non-cooperative foundation of stability

This section provides a natural non-cooperative interpretation of stability in large markets. Each outcome canonically defines a game that models the process of organizing on a blocking coalition. An outcome is stable if and only if no block can be found under this process in some subgame-perfect equilibrium. If there are any sufficiently small costs to forming a coalition, this holds for every subgame-perfect equilibrium.

Starting from an initial outcome the deviation game describes the process of how a deviating coalition can be organized. First, a random agent is selected who can offer a set of contracts. Due to the anonymity of the market, offers are made not to specific agents but to types. Nature draws agents of corresponding types at random to receive contract offers. In the second round, agents who receive offers decide whether to participate or not. Participation is costly. If all contacted agents participate, they can decide in a defection Stage which of the offered contracts to sign and, in a final stage, which of their initial existing contracts to keep.

The dynamics of the game make it necessary for the first agent to anticipate the behavior of later agents. The option of defection after participation makes it necessary for the initial agent to make a proposal that leaves no incentives for later defections. Since the agent has to take into account how specifically other agents in the coalition will react, they are not anonymous to her.

The deviation game Let $i(a) \in I$ denote the type of agent $a \in \Theta$. Let $n > 1$ and ν be a graph with n nodes. We write $Z(\nu)$ for a *network of new contracts* of shape ν consisting of a type $t(j)$ for all $1 \leq j \leq n$ and a contract $x_{i,j} \in X_{i,j}$ for all edges $\{i, j\} \in \nu$.¹⁰ We write $O^j = \{x_{j,k}^j | \{j, k\} \in \nu\}$ for the set of contracts assigned to a node $1 \leq j \leq n$.

Definition B.1. For each outcome M and cost $\kappa \geq 0$, let $G_\kappa^*(M)$ be the following finite, dynamic game with almost perfect information and exogenous uncertainty, called a *deviation game*. Actions are observed immediately by all other players. The stages of the game are as follows:

- Round 1: Contracts are offered
 - Stage 1: Nature selects an agent $a_0 \in \Theta$ uniformly at random.
 - Stage 2: Agent a_0 decides whether to participate.
If she decides to participate she has to pay a participation cost κ , otherwise the game terminates immediately without a change of matches and the payoff of every agent $a \in \Theta$ is $u^{i(a)}(M(a))$, i.e. the utility the agent receives in the initial outcome.
 - Stage 3: If agent a_0 does decide to participate, she chooses a network $Z(\nu)$ of new contracts and a node j_0 for herself in the network.
 - Stage 4: Nature selects for each node $j_0 \neq j \in \{1, \dots, n\}$ an agent $a(j) \in \Theta_{(t(j))}$ uniformly at random and offers the set O^j to the agent $a(j)$.
- Round 2: Contracts are accepted
 - Stage 5: Each agent $a(j)$ for $j \neq j_0$ decides to participate and has to pay a participation cost κ . If an agent decides not to participate, the game ends without a change of matches and every agent receives her utility of the initial outcome ($u^{i(a)}(M(a))$); those agents that decided to participate additionally pay the participation cost κ .

¹⁰Note, that the network of new contracts is also part of the data of a block.

- Stage 5* (Defect): Each agent $a(j)$ decides which subset of the offered contracts $Z(j) \subseteq O^j$ to sign. Each contract $x_{i,j} \in Z(j) \cap Z(i)$ that is mutually signed is implemented. The set of all mutually signed contracts in which agent $a(j)$ participates is denoted by $Z_{a(j)}$.
- Stage 6: Each agent $a(j)$ decides which subset $Y_{a(j)} \subset M(a(j))$ of her existing contracts to keep. The payoff of agent $a(j)$ is

$$u^{t(j)}(Y_{a(j)} \cup Z_{a(j)}) - \kappa.$$

The payoff of agents that do not participate is simply the utility of the initial outcome.

With a continuum of agents, no agent is selected more than once by nature almost surely. We can therefore condition on nature never selecting any agent more than once.

We can now characterize stability in terms of subgame-perfect equilibria¹¹ of deviation games.

Theorem B.1. *The outcome M is stable if and only if for all costs $\kappa > 0$ and all subgame-perfect equilibria of the deviation game $G_\kappa^*(M)$, with probability one, no agent chooses to participate.*

If an outcome is stable, there always exists a subgame-perfect equilibrium in which no agent chooses to participate. However, if all costs are zero there exist subgame-perfect equilibria where agents participate; for example agents may participate in Stage 2 but not in Stage 4. The existence of such trivial equilibria is ruled out by imposing any strictly positive cost κ . On the other hand, in an instable outcome, for sufficiently small costs there always exists a subgame-perfect equilibrium where agents choose to participate.

Intuitively, since with probability one in Stage 6 actions cannot have any impact on other coalition members, agents can freely sign their optimal choice, given old and new contracts in the defection Stage 5*. Thus, with probability one in every Nash equilibrium of Stage 5*, agents only sign a set of mutually optimal choices. This reflects exactly the strong desirability condition in the definition of a block.

Implicitly, in the deviation game $G_\kappa^*(M)$ agents do not have to sign the contract right away when accepting in Stage 5. If in comparison, agents have to sign the contracts immediately when agreeing to participate, Stage 5* becomes superfluous, and might as well be dropped. The subgame-perfect equilibria of the modified game $G_\kappa(M)$ without Stage 5*, allow for more commitment and make it therefore easier to agree on a deviation. In particular, this modification allows for a non-cooperative foundation of strong group stability in large markets.

Definition B.2. For each outcome M and cost $\kappa > 0$ a game $G_\kappa(M)$ is defined as $G_\kappa^*(M)$ without Stage 5* and letting $Z_{a(j)} = O^j$.

Here, the identification $Z_{a(j)} = O^j$ means that if an agent participates she has to sign all contracts in Stage 5. As a consequence, she can ignore the decisions of other agents, and therefore treat even other agents inside the blocking coalition as anonymous. This reflects exactly the implicit anonymity assumption of strong group stability.

Theorem B.2. *The outcome M is strong group stable if and only if for all costs $\kappa > 0$ and all subgame-perfect equilibria of the offer game $G_\kappa(M)$, with probability one, no agent chooses to participate.*

Intuitively, the only decision an agent in Round 2 can make is to participate or not. The incentive to participate depends only on the desirability of the new set of contracts. This entails exactly the properties of a strong group block.

¹¹Technically, as there is a continuum of players but only finitely many players make moves in each history, we restrict to strategy profiles that are measurable.

Similarly, restricting the allowed shape of new contracts in Stage 3 of the deviation game allows one to give non-cooperative foundations of stability concepts that restrict the shape of possible blocks, as for example in pairwise-stability.¹²

C Proofs

C.1 Proof of Proposition A.1

Assume an outcome M is not in the strict core. Consider the minimal connected graph that comprises a weak core block of M . We now construct a strong group block by restricting to the nodes, where the corresponding matched types can indeed strictly improve. If an agent is indifferent, the agent has to keep the old contracts, as preferences are strict.

C.2 Proof of Theorem 1

As already discussed in section 5 stability trivially implies setwise stability. This is also true in finite markets.

To prove the converse, starting from an outcome M that is blocked by a block of shape ν , we use the largeness of the market to construct an outcome M' and choose a coalition A such that the outcome M is setwise blocked. For the setwise block we use the shape ν and the new contracts from the initial block of shape ν . Then we define the outcome M' and the coalition A in the following: Intuitively, we divide the set of agents in subsets such that the blocking coalition is not affected by dropped contracts. As blocks are finite and the market is large this is possible. Hence, the blocking set of contracts can be implemented for a small enough measure κ , such that the contracts that agents within the blocking coalition want to drop do not affect the other agents of the blocking coalition. In particular, we divide the set of agents into three disjoint sets which are either part of the blocking coalition (A) or outside of the blocking coalition but affected by the implementation (hence the set of agents that get contracts cancelled) (C) and the set of agents that remain unchanged (O) in the new outcome.

Formally, let the minimal positive measure of a matched type that occurs in M be $\epsilon_M := \min\{\mu(M_Y^i) | \mu(M_Y^i) > 0\}$. Then choose $0 < \kappa < \frac{\epsilon}{2n}$. Now, for each matched type (i, Y) we choose an arbitrary measurable partition of M_Y^i into two sets $M_Y^i(C)$ and $M_Y^i(A)$ of equal measure.

$$M_Y^i = M_Y^i(A) \cup M_Y^i(C)$$

We choose the coalition A to be a subset of $M(A) = \bigcup_{i,Y} M_Y^i(K)$.

- Defining the blocking coalition A :

We choose for all nodes j of the graph ν arbitrary but disjoint measurable sets $A(j)$ of agents of measure κ within $M_{Y_j}^{i_j}(K)$ and define $A := \bigcup_j A(j)$ to be the blocking coalition.

- Defining the set C of agents that are not part of the coalition but get contracts cancelled:

For each node j there is a matched type (i_j, Y_j) and a set of contracts corresponding agents want

¹²In Jagadeesan and Vocke (2021) a game is defined, where agents can iteratively offer contracts and eventually detect tree-blocks (this iterative approach is inspired by the heuristics introduced in (Fleiner, Jankó, Tamura, and Teytelboym, 2018) and (Hatfield et al., 2021)). While in Jagadeesan and Vocke (2021) an agent can offer a set of contracts that she wants to sign, in the deviation game of this paper, an agent can offer sets of contracts including offers she is personally not involved in. Intuitively, agents can suggest cooperations that include contracts between other agents. Consequently, different acceptability notions matter for the strategic behavior of agents. This particularly reflects the anonymity assumptions in stability concepts, making it possible to give non-cooperative foundations of different versions of stability.

to cancel, which we call $C(j) := W_j \setminus Y_j$. Hence, for each type k every contract $x_{i,j,k}$ has to be cancelled for a set of agents of type k with measure

$$c(x_{i,j,k}, k) = \kappa \cdot \sum_{r \leq n: i_r = i} 1\delta_{x \in C(r)}.$$

Now, for each contract $x_{i,j,k}$ and type k we choose an arbitrary measurable set

$$C(x_{i,j,k}, k) \subseteq \left(\bigcup_{i,Y} M_Y^i(C) \right) \cap \left(\bigcup_{Y:x \in Y} M_Y^k \right)$$

of measure $c(x_{i,j,k}, k)$. We define the set of all agents that get contracts cancelled to be

$$C = \bigcup_{(x_{i,j,k}, k): 1 \leq j \leq n} C(x_{i,j,k}, k).$$

- Defining the set of agents without changes:

Simply let $O = M \setminus (C \cup A)$.

Consider a set of contracts $W \subseteq X$. Let $W(A)$ be the set of agents in the coalition A that changes from a set of contracts in the initial outcome M to the set W in the outcome M' . Formally, this is the set of agents $a \in A$ s.t. $a \in M_{Y_j}^{i_j}$ for some node j and $W_j = W$.

Let $W(C)$ be the set of agents in the set C that changes from a set of contracts in the initial outcome M to the set W in M' : Formally, this is the set of agents $a \in C$ s.t. $a \in M_Y^i$ for some (i, Y) and $x_{i,k} \in Y \setminus W$ if and only if $a \in C((x_{i,k}, i))$. Intuitively, agent a changes from some set of contracts Y to a set of contracts W if for each contract x in $Y \setminus W$ agent a gets the contract x cancelled, hence is in the corresponding cancelled contract set $C(x_{i,k}, i)$.

Now we are ready to define a new outcome M' in the following way:

$$M_W^i = W(A) \cup W(C) \cup (M_W^i \cap O).$$

Agents who sign W are either agents in the coalition (A) who change to W or in the set of agents who get contracts cancelled (C) who change to W or in the unchanged set (O) who have signed the set W before.

We can easily see from the construction of the new outcome M' and the coalition A that M' is an outcome and that M is blocked by the setwise block of shape ν defined above.

C.3 Proof of Theorem B.1

$\Rightarrow$ For the first direction, we show that if there exists a subgame-perfect equilibrium in $G_\kappa^*(M)$ with cost $\kappa > 0$ such that with positive probability an agent chooses to participate, then the outcome M is not stable.

Suppose a positive measure chooses to participate, then a positive measure of agents must already choose to participate in Stage 2. The expected payoff of the agents when participating in Stage 2 must be at least as high as if they were to choose not to participate in Stage 2, in which case their payoff would be the utility of their current match. Since agents experience a cost κ of participation, each participating agent must have a positive probability of changing her match if she is to participate in Stage 2. We divide into cases based on whether a positive measure of agents choose an empty or non-empty network

of new contracts (in Stage 3) to show that M cannot be stable.

Case 1: A positive measure of agents choose an empty set of new contracts in Stage 3. Let A' be the set of agents that offer a network of contracts. Let (i, Y) be a matched type such that $M_Y^i \cap A'$ has positive measure. Then, each agent in $M_Y^i \cap A'$ must obtain strictly higher utility by dropping some of their existing contracts than by retaining them all—i.e., we must have that $Y \notin C^i(Y)$. Hence, M is not individually rational.

Case 2: A positive measure of agents choose a non-empty network of contracts in Stage 3. In every subgame-perfect equilibrium, agents play a Nash equilibrium in Stage 5*, since Stage 6 doesn't impact other participating agents with probability one. In any Nash equilibrium of Stage 5*, mutually signed contracts correspond exactly to contracts that are best choices given the set of existing and new contracts ($Z_{a(j)} \subseteq C(Y \cup Z_{a(j)})$). Thus, the corresponding agents form a block. As counterparties are selected uniformly at random from the agents of each type, we have that, almost surely, the block must arise at M .

$\Leftarrow$ For the other direction, we show that if the outcome M is not stable, then there exists a subgame-perfect equilibrium σ in $G_\kappa^*(M)$ and costs $\kappa > 0$ such that with positive probability, some agent chooses to participate.

Case 1: Suppose that M is not individually rational. Then, a matched type (i, Y) arises such that $Y \notin C^i(Y)$. If an agent of matched type (i, Y) is selected by nature in Stage 1, then she would receive a strictly higher expected payoff than the utility of her existing match by choosing to participate and choosing no new contracts for κ small enough. Hence, such agents must choose actions such that a change of match occurs with positive probability in every subgame-perfect equilibrium of $G_\kappa(M)$.

Case 2: Suppose that M is individually rational and a block of shape ν arises. Let $\epsilon_\nu > 0$ be the minimum measure of any matched type (i_j, Y_j) corresponding to some j with $1 \leq j \leq n$. Let c_ν be the minimal utility improvement a matched type receives in the block. Let κ be smaller than $c_\nu \cdot \frac{\epsilon_\nu}{(\mu(\Theta))^n}$. We can construct a subgame-perfect equilibrium in the following way: Consider first an agent a of a matched type (i, Y) that corresponds to a node of the block. Agent a participates in Round 2 if and only if the network of shape ν corresponding to the block was offered in Round 1 and the set of offers yields indeed a best choice for her. If she participates in Stage 5, she chooses the set of utility maximizing contracts in Stage 5*, which is, by construction, exactly the set of contracts offered. If she is selected by nature in Round 1, then she offers the corresponding network while choosing the best possible node j_0 for herself. All agents of other matched types do not participate. This is a subgame-perfect equilibrium, since in Round 2 participation is a best response if and only if other agents participate and an improvement is possible, hence in particular, only if the corresponding network was offered. On the other hand, offering contracts can only lead to an improvement if agents in Round 2 participate. Participation in Round 1 for agent a is part of a subgame-perfect equilibrium since costs are smaller than the expected improvement.

C.4 Proof of Theorem B.2

Note that the only part that needs to be changed in the proof of Theorem B.1 for proving Theorem B.2 is Case 2 in the second direction " $\Leftarrow$ ": It remains to show that if an outcome M is not strong group stable, there exists a subgame-perfect equilibrium in $G_\kappa(M)$ and costs $\kappa > 0$ such that with positive probability an agent chooses to participate:

Suppose that M is individually rational and a strong group block of shape ν arises. Let $\epsilon_\nu > 0$ be the minimum measure of matched types corresponding to the block in the outcome. Let c_ν be the

minimal utility improvement a matched type receives in the strong group block. Let κ be smaller than $c_\nu \times \frac{\epsilon_\nu}{\mu(\Theta)}$. We can construct a subgame-perfect equilibrium in the following way: Consider first an agent a of a matched type (i, Y) that corresponds to a node of the strong group block. Agent a participates in Round 2 if and only if the network of shape ν corresponding to the strong group block was offered in Round 1 and the set of offers yields indeed an improvement for her. She offers the corresponding network in Round 1 while choosing the best possible node j_0 for herself. All agents of other matched types do not participate. This is a subgame-perfect equilibrium, since in Round 2 participation is a best response if and only if other agents participate and an improvement is possible, hence in particular, only if the corresponding network was offered. On the other hand, offering contracts can only lead to an improvement if agents in Round 2 participate. Participation in Round 1 for agent a is part of a subgame-perfect equilibrium since costs are smaller than the expected improvement.