Decentralized Foundation for Stability of Supply Chain Networks

Kirill Rudov*

November 12, 2022

Click here for the latest version

Abstract

This paper proposes simple dynamics generating a stable supply chain network. We prove that for any unstable network, there exists a finite sequence of successive myopic blocking chains leading to a stable network. Our proof suggests an algorithm for finding a stable network that generalizes the classical Gale and Shapley (1962)'s deferred acceptance algorithm.

*Department of Economics, Princeton University (email: krudov@princeton.edu). I am grateful to Alessandro Lizzeri for insightful discussions that inspired this project as well as Leeat Yariv and Wolfgang Pesendorfer for their guidance. I thankfully acknowledge financial support from the William S. Dietrich II Economic Theory Center at Princeton University.

1. Introduction

Various natural resource industries—oil, coal, and natural gas—and manufacturing sectors—clothing, furniture, and electronics—operate through supply chains. Despite their fundamental role in the modern economy, few theoretical models predict the architecture of supply chain networks. A notable exception is the paper by Ostrovsky (2008) that extends the study of stability developed for two-sided matching markets to supply chain networks.[1] In practice, however, many supply chain networks lack a central authority enforcing stability and instead are formed by decentralized decisions of self-interested entities. Can their uncoordinated interactions lead to a stable network?

This paper introduces simple dynamics generating a stable supply chain network. We prove that for any unstable network, there exists a finite sequence of successive myopic *blocking chains*—each represented by a downstream sequence of entities that benefit from re-contracting with each other—that leads to a stable network. This result implies that a decentralized process, in which a randomly-chosen blocking chain re-contracts each period, converges to a stable supply chain network with certainty. An advantage of these dynamics is that convergence to stability is ensured even when decision makers have limited sophistication and only local information about a network's evolving topology and others' preferences, features that may be inherent in large, complex networks.

As by-products, we obtain a new proof for the existence of stable supply chain networks and a simple algorithm to find such a network. When applied to an empty network with no links between entities, our algorithm extends the famous Gale and Shapley (1962)'s deferred acceptance algorithm that underlies a number of centralized matching clearinghouses.

We generalize the analogous stabilization theorem of Roth and Vande Vate (1990) for

[1] See Roth and Sotomayor (1992) for a comprehensive introduction to two-sided matching. Hatfield et al. (2013) introduce trading networks that generalize supply chain networks. For recent developments, see Hatfield et al. (2021). For alternative models of supply chain networks, see Amelkin and Vohra (2020), Elliott, Golub, and Leduc (2022), and references there.

one-to-one matching markets.[2,3] Nonetheless, our proof differs substantially from theirs. Supply chain networks pose an additional challenge: by re-contracting, a blocking chain of entities might affect many other entities in the network and cause significant changes in the network architecture. To deal with this new issue, this paper builds on recent techniques of Ackermann et al. (2011) from computer science.

2. The Model

Consider a finite set A of agents (firms, countries, or other entities). Let $\triangleright$ be an "upstream–downstream" strict partial order on this set, so that $a \triangleright b$ implies b is a downstream partner for a. This means that a can potentially supply an indivisible good or service to b, or equivalently, b can potentially demand it from a. If $a \ntriangleright b$ and $b \ntriangleright a$, then a can neither supply to b nor demand from b. Since $\triangleright$ is transitive, there are no loops: no agent can supply—directly or indirectly—to any of her upstream partners.

Let S be the set of agents that can only supply; refer to such agents as *suppliers*. Also, let C denote the set of agents called *consumers* that can only demand. All other agents $I = A \setminus (S \cup C)$ are called *intermediaries*: they can both supply and demand.

2.1. Preferences. Suppose that all agents have unit capacities: each agent can be matched with at most one downstream partner and at most one upstream partner. For each $a \in A$, let $U(a)$ be the set of upstream agents that can supply to a. Similarly, $D(a)$ is the set of downstream agents that can demand from a.

Each supplier $a \in S$ has strict preferences $\succ_a$ over possible downstream partners $D^+(a) \equiv D(a) \cup \{\varnothing\}$, where $\{\varnothing\}$ denotes the outside option of having no downstream partner. A

[2]Recently, Rudov (2022) characterized conditions, under which the rematching dynamics can generate any stable outcome. However, the focus of that paper is rather on the fragility of stable outcomes.

[3]Similar stabilization results were later obtained for roommate markets (Chung, 2000; Diamantoudi, Miyagawa, and Xue, 2004; Inarra, Larrea, and Molis, 2008), many-to-one markets with couples (Klaus and Klijn, 2007), many-to-many markets without contracts (Kojima and Ünver, 2008) and with contracts (Millán and Pepa Risma, 2018), and matching markets with incomplete information (Lazarova and Dimitrov, 2017; Chen and Hu, 2020).

downstream partner $b \in D(a)$ is *acceptable* if $b \succ_a \varnothing$.

Also, each consumer $a \in C$ has strict preferences $\succ_a$ over all potential upstream partners $U^+(a) \equiv U(a) \cup \{\varnothing\}$ with $\{\varnothing\}$ being the outside option of having no upstream partner. An upstream partner $b \in U(a)$ is acceptable if $b \succ_a \varnothing$.

Finally, each intermediary $a \in I$ has strict preferences $\succ_a$ over all its possible upstream-downstream pairs $U^+(a) \times D^+(a)$. Assume that any such intermediary a has two separate strict rankings, $\succ_a^U$ over $U^+(a)$ and $\succ_a^D$ over $D^+(a)$, respectively. Furthermore, these rankings are linked to $\succ_a$ in such a way that (i) for any $(b, c) \in U^+(a) \times D^+(a)$, we have $(b, c) \succ_a (\varnothing, \varnothing)$ if and only if $b \in U(a)$ and $c \in D(a)$ are acceptable with respect to $\succ_a^U$ and $\succ_a^D$, respectively; (ii) $\succ_a$ over two different acceptable upstream-downstream pairs $(b, c), (b, c') \succ_a (\varnothing, \varnothing)$ agrees with $\succ_a^D$ over c, c'; and (iii) $\succ_a$ over $(b, c), (b', c) \succ_a (\varnothing, \varnothing)$ agrees with $\succ_a^U$ over b, b'.

Intuitively, condition (i) implies that an intermediary is willing to supply an output good only if she can also demand an input good required to produce that output, and vice versa. In other words, she prefers not to supply (demand) at all if she has no one to demand from (respectively, supply to). Conditions (ii) and (iii) mean that her overall preferences are responsive with respect to her rankings of all possible upstream and downstream partners. Namely, (ii) says that an intermediary prefers a downstream partner that is higher in her ranking of downstream partners, given her fixed upstream partner; (iii) is analogous.

2.2. Stable Networks. We employ a graph representation to define a network. Each agent $a \in A$ is viewed as a node. A *partnership* between agents $a, b \in A$ is an edge between nodes a and b, denoted by (a, b) or (b, a); we consider undirected edges because the predetermined vertical ordering $\triangleright$ itself defines a direction of exchange. A *graph* is a set of partnerships, or simply a set of unordered pairs of agents.

A *network* μ is a graph such that for every $a, b, c \in A$, (i) $(a, b) \in \mu$ implies $b \in U(a) \cup D(a)$; (ii) $(a, b), (a, c) \in \mu$, $b \neq c$, $b \in U(a)$ implies $c \in D(a)$; and (iii) $(a, b), (a, c) \in \mu$, $b \neq c$, $b \in D(a)$ implies $c \in U(a)$. That is, a network is a collection of supply-demand relationships between agents such that every agent has at most one upstream and at most one downstream partner. Let $\mu(a)$ denote the set of partners for agent $a \in A$. In addition, let $\mu^U(a)$ and $\mu^D(a)$ be her upstream and downstream partner sets under μ, respectively.[4] Call an agent $a \in A$ *matched* if her partner set $\mu(a)$ is non-empty.

An agent $a \in A$ is said to *block* network μ if she prefers to unilaterally drop any of her partners from $\mu(a)$. A network is *individually rational* if no agent blocks it.

A *chain* is a sequence of agents, $(a_1, a_2, \ldots, a_n) \in A^n$, such that for any $i < n$, $a_i \in U(a_{i+1})$, i.e., a_i can potentially supply to a_{i+1}. The *length* of the chain is n, the number of agents involved in the chain. Also, the chain does not have to start with a supplier and end with a consumer; for instance, it can involve only intermediaries.

We say that a chain $(a_1, a_2, \ldots, a_n) \in A^n$ is a *blocking chain* for network μ if (i) for any $i < n$, $(a_i, a_{i+1}) \notin \mu$; (ii) a_1 gets better off by replacing $\mu^D(a_1)$ with $\{a_2\}$; (iii) a_n gets better off by replacing $\mu^U(a_n)$ with $\{a_{n-1}\}$; (iv) and for any $1 < i < n$, a_i gets better off by replacing $\mu(a_i)$ with $\{a_{i-1}, a_{i+1}\}$. In other words, a blocking chain for network μ is a downstream sequence of agents in which any consecutive pair of agents does not form a partnership under μ but every agent would benefit from forming partnerships with her neighbor(s) in the sequence by potentially dropping her current partner(s) under μ. A network is *chain stable* if it is individually rational and has no blocking chains.

In fact, our setup is a special case of a more general setting of supply chain networks with contracts introduced by Ostrovsky (2008). In his setting, each agent can be involved in multiple contracts and has fully substitutable preferences over all possible contracts involving her. Hatfield and Kominers (2012) showed that for this setting, blocking chains are the essential blocking sets in the sense that any chain stable network is also stable. In other words, a chain stable network is robust with respect to blocks by arbitrary coalitions of agents, not only chains. Given this equivalence, for the remainder of the paper, we refer to chain stability as *stability*.

[4]Each of $\mu^U(a)$ and $\mu^D(a)$ is either a singleton or an empty set.

Supply chain networks naturally extend standard two-sided matching markets. Indeed, any two-sided market can be considered as a supply chain with no intermediaries, $I = \varnothing$, in which agents on one side can be viewed as suppliers and agents on the other side can be regarded as consumers. Furthermore, for two-sided markets, chain stability reduces to pairwise stability, that is known to be equivalent to overall stability.

2.3. Rematching Dynamics. Let network μ be unstable. For any blocking agent $a \in A$ of μ, a network μ' is *obtained from μ by satisfying a*, if

$$\mu' = \mu - \{(a,b) \mid b \in \mu(a)\},$$

where $\{(a,b) \mid b \in \mu(a)\}$ is the a's set of dropped partners.[5]

For any blocking chain $(a_1, a_2, \ldots, a_n) \in A^n$, a network μ' is *obtained from μ by satisfying* $(a_1, a_2, \ldots, a_n)$, if

$$\mu' = \mu - \{(a_1,b) \mid b \in \mu^D(a_1)\} - \bigcup_{i=2}^{n-1}\{(a_i,b) \mid b \in \mu(a_i)\} - \{(a_n,b) \mid b \in \mu^U(a_n)\} \cup \bigcup_{i=1}^{n-1}(a_i, a_{i+1}).$$

That is, each agent in the chain forms partnerships with her chain neighbor(s), $\bigcup_{i=1}^{n-1}(a_i, a_{i+1})$, by potentially dropping her corresponding current partner(s) under μ.

We say that there is a *path* from network λ to network μ if there exists a sequence of networks $\lambda_1, \lambda_2, \ldots, \lambda_k$, such that $\lambda = \lambda_1$, $\mu = \lambda_n$, and for each $i < k$, λ_{i+1} is obtained from λ_i by satisfying a blocking agent or a blocking chain. In that case, we say network λ can *reach* or *attain* network μ.

From a theoretical perspective, our rematching process parallels the one of Roth and Vande Vate (1990), which is the leading dynamics for the standard two-sided matching environment. It is appealing for several reasons. Agents do not need to hold precise information about the network topology and all agents' preferences—they can use only local information about their neighbors. Also, these dynamics impose low sophistication requirements on agents; this feature is particularly relevant for large, complex networks.

[5] By our preference assumptions, if intermediary $a \in I$ blocks network μ, she wants to drop all her partners.

In practice, it is easy for agents to identify and implement blocking chains. For example, a customer could call a potential supplier asking her whether she would form a partnership. Then, that supplier could call one of her potential suppliers, and so forth until a blocking chain is identified; for an extended discussion, see Section 4 in Ostrovsky (2008).

3. Convergence to Stability

In this section, we show that for any unstable network, there exists a finite sequence of myopic blockings that leads to a stable network. In other words, stable networks can always be attained by means of decentralized decision making.

In fact, our stabilization result is stronger—it holds under an additional natural restriction on blocking chains. Consider a blocking chain $(a_1, a_2, \ldots, a_n) \in A^n$ for network μ. In that case, agent a_1 is said to *start* the blocking chain $(a_1, a_2, \ldots, a_n)$. We say that this chain, $(a_1, a_2, \ldots, a_n)$, is a *downstream-best blocking chain* if for any $i \leq n$, it is the best blocking chain for a_i among all blocking chains starting with a_1 and going through agents $a_2, a_3, \ldots, a_i$. Intuitively, such blocking chains provide more upstream potential suppliers with more bargaining power when coordinating on a blocking chain to be implemented. For the theorem below, restrict attention to downstream-best blocking chains.[6]

Theorem 1. *For any unstable network, there is a finite sequence of blocking agents and downstream-best blocking chains that leads to a stable network.*

This stabilization theorem generalizes the one of Roth and Vande Vate (1990) for one-to-one matching markets. However, our proof differs significantly from theirs. First, they do not put any additional restrictions on blocking pairs, an analogue of blocking chains for two-sided markets. Second, their proof constructs a sequence of blocking pairs in such a way that it induces a monotonically expanding sequence of agents' sets containing no blocking pairs—also often called internally stable sets—until it reaches a set of all agents. Such construction

[6]The analysis is symmetric for upstream-best blocking chains, defined analogously, in which more downstream potential demanders are endowed with more bargaining power.

is difficult in our setting. Indeed, since blocking chains might involve many agents, each of those may also have her own partners, maintaining internal stability is problematic. We instead employ and modify recent computer science techniques of Ackermann et al. (2011) to prove the theorem. In fact, our rematching process coincides with their "best response" dynamics for one-to-one matching markets. Nonetheless, our analysis is considerably more general and also applies to supply chain networks.

As by-products, we obtain a novel proof for the existence of stable networks and a simple algorithm to compute one such network; in contrast to Ostrovsky (2008) and Hatfield and Kominers (2012), our analysis does not rely on fixed-point arguments. In particular, when applied to an empty network, the proposed algorithm naturally generalizes the celebrated deferred acceptance algorithm (Gale and Shapley, 1962).

Before proving the theorem, consider the *random* rematching dynamics that for any given network, satisfies every (downstream-best) blocking chain with positive probability that depends only on the network; for convenience, suppose that a blocking agent is a blocking chain involving herself. Our result implies that for an arbitrary initial network, this process converges to a stable network with certainty.

In what follows, we prove the theorem in a series of lemmas. All proofs are relegated to the appendix.

We can first sequentially satisfy one agent at a time to obtain an individually rational network, starting from an arbitrary given unstable network. Since at every step of this procedure, agents only drop their partners, this process necessarily terminates after finitely many steps; see Lemma 1.

Lemma 1. *For any unstable network, there is a finite sequence of blocking agents that leads to an individually rational network.*

Next, we proceed in two stages to generate a stable network. In the first stage, we construct a finite sequence of blockings leading to an individually rational network in which no matched agent can start a blocking chain (Lemma 3). In the second stage, we iteratively consider one unmatched supplier at a time, satisfy her downstream-best blocking chain, and appropriately compensate affected intermediaries within each iteration in order to maintain the property that no matched agent can start a blocking chain (Lemma 5). This stage terminates after finitely many steps and culminates in a network in which neither matched agent nor unmatched supplier can start a blocking chain (Lemma 6). The resulting network is stable since no unmatched intermediary can start a blocking chain either (Lemma 4).

In order to describe these stages, we employ a more refined partition of agents into multiple tiers. *Suppose that the longest chain of agents has length n.* Lemma 2 states that we can classify all agents into n tiers such that any agent from a higher tier—with a lower index—can potentially supply only to those from tiers below. In our terminology, suppliers constitute the highest tier. Similarly, the lowest tier consists of consumers. Intermediaries belong to other tiers.

Lemma 2. *It is possible to partition A into n disjoint subsets A_1, A_2, ..., A_n, such that for any $i < n$ and $a \in A_i$ with $a \in U(b)$, we have that $b \in A_{i+1} \cup A_{i+2} \cup \ldots \cup A_n$. In particular, $S = A_1$, $C = A_n$, and $I = A_2 \cup A_3 \cup \ldots \cup A_{n-1}$.*

Based on the tier structure, we describe the first stage by using the *rematch* operation. It is defined as follows:

Rematch. This operation can be modelled as a finite-state machine. It consists of n states, starts in state 1, and transitions between states as described below.

For any $k < n$, define

State k. Consider matched $a \in A_k$ that can start a blocking chain. If there is no such agent, proceed to state $k + 1$. Otherwise, satisfy the downstream-best blocking chain starting with a, restore individual rationality, and proceed to state 1.

State n is a termination state.

Lemma 3. *For any individually rational network, the rematch operation induces a finite sequence of blocking agents and downstream-best blocking chains that leads to an individually rational network in which no matched agent can start a blocking chain.*

Intuitively, focus on matched suppliers from $S = A_1$. In state 1, when a matched supplier satisfies her downstream-best blocking chain, she becomes better off and no other matched supplier—that remains to be matched after satisfying the chain—is worse off. Since the set of matched suppliers can only shrink and one of those suppliers becomes better off, this can happen only finitely many times. Between any such two visits of state 1, matched intermediaries from A_2 can satisfy their downstream-best blocking chains in state 2 only finitely many times, and so forth. By iteratively using this argument, we conclude that the rematch operation necessarily terminates. This ends the first stage.

If after the first stage, no unmatched supplier can start a blocking chain, the resulting network is stable, thus concluding the proof (see Lemma 4 below). Indeed, in that case, no unmatched intermediary can start a blocking chain either—recall that an intermediary prefers to supply only if she can also demand, and vice versa.

Lemma 4. *Any individually rational network, in which neither matched agent nor unmatched supplier can start a blocking chain, is stable.*

Thus, for the remainder of the analysis, assume that there is an unmatched supplier $a \in S = A_1$ that can start a blocking chain. Next, proceed to the second stage.

In the second stage, we iteratively consider one unmatched supplier at a time, $a \in S = A_1$, and apply the *match* operation for this supplier, denoted by *match*(a). Essentially, this operation first satisfies the a's downstream-best blocking chain. Then, the operation recursively compensates intermediaries affected along the rematching process in order to maintain the property that no matched agent can start a blocking chain; see Lemma 5.

Match(a). Satisfy the downstream-best blocking chain starting with supplier a. Suppose that it ends at agent b. The rest can be represented as a finite-state machine with states $\theta \in \Theta$ encoded by intermediaries and consumers, $\Theta \equiv I \cup C$, including an additional termination state. It starts in state $\theta = b$ and transitions between states as described below.

State θ. Let c be θ's previous—that is, before satisfying the most recent block that can be either a blocking chain or a blocking agent —upstream partner, if any. If θ was previously unmatched or $c \in S = A_1$, terminate. Otherwise, verify if c can start a blocking chain. If not, satisfy blocking agent c and proceed to state c. If yes, satisfy the downstream-best blocking chain starting with c. Suppose it ends at agent d. Proceed to state d.

Lemma 5. *Consider any individually rational network λ in which no matched agent can start a blocking chain. Suppose an unmatched supplier $a \in S = A_1$ can start a blocking chain. Then, the match(a) operation induces a finite sequence of blocking agents and downstream-best blocking chains that leads to an individually rational network μ in which no matched agent can start a blocking chain. In addition, there exists $i \leq n$ such that:*

(i) no agent from A_i is worse off and at least one of them is better off in the new network μ compared to the original network λ;[7]

(ii) for any $j > i$, all agents from A_j have identical partners sets under both μ and λ.

To glean some intuition for the lemma's first part, satisfy the a's downstream-best blocking chain and suppose it ends at agent b. Follow the instructions for state b. Let c be b's previous upstream partner, if any. If b was previously unmatched, then $b \in C = A_n$ must be a consumer. In that case and when $c \in S = A_1$ is a supplier, no matched agent can start a blocking chain in the obtained network by construction; the match operation terminates and the result follows immediately. Otherwise, $c \in I$ is an intermediary.

By construction, only c is worse off compared to the previous network and needs to be compensated in order to maintain the property that no matched agent can start a blocking

[7]It means that agents from A_i Pareto-improve upon the starting network.

chain. If she can start a blocking chain, we satisfy her downstream-best blocking chain. Assuming it ends at agent d, we transition to state d, compensate d's previous upstream partner e, if any, and then proceed recursively. Else, if c cannot start a blocking chain, we satisfy blocking agent c herself; namely, c drops her upstream partner f. Then, we switch to state c, compensate f, if needed, and proceed iteratively.

This iterative process results in a compensation chain that terminates in a finite number of steps. Indeed, focus on consumers from A_n. At any step, when a satisfied blocking chain ends at a consumer, she becomes better off. Since consumers cannot improve indefinitely, such steps can appear only finitely many times. Between any such two steps, intermediaries from A_{n-1} can be endpoints of satisfied blocking chains only finitely many times, and so on. By iteratively employing this reasoning, we prove that the match operation terminates.

Furthermore, as concerns the lemma's second part, i corresponds to the largest tier of an endpoint among all endpoints of satisfied blocking chains; see details in the appendix.

Finally, Lemma 6 concludes the second stage: it asserts that the iterative application of the match operation generates a stable network, as required.

Lemma 6. *For any individually rational network, in which no matched agent can start a blocking chain, the iterative application of the match operation induces a finite sequence of blocking agents and downstream-best blocking chains that leads to a stable network.*

The idea is to apply the match operation until there is no unmatched supplier that can start a blocking chain and then use Lemma 4 to conclude the proof. Based on the Lemma 5's second part, this procedure terminates in a finite number of iterations; it follows by similar arguments to those employed to prove termination in the previous lemma.

To sum up, we can combine all presented lemmas to prove the theorem. Figure 1 explains graphically the algorithm used to obtain a stable network starting from an arbitrary network.

Proof of Theorem 1. The result follows immediately from Lemmas 1, 3, and 6. ∎

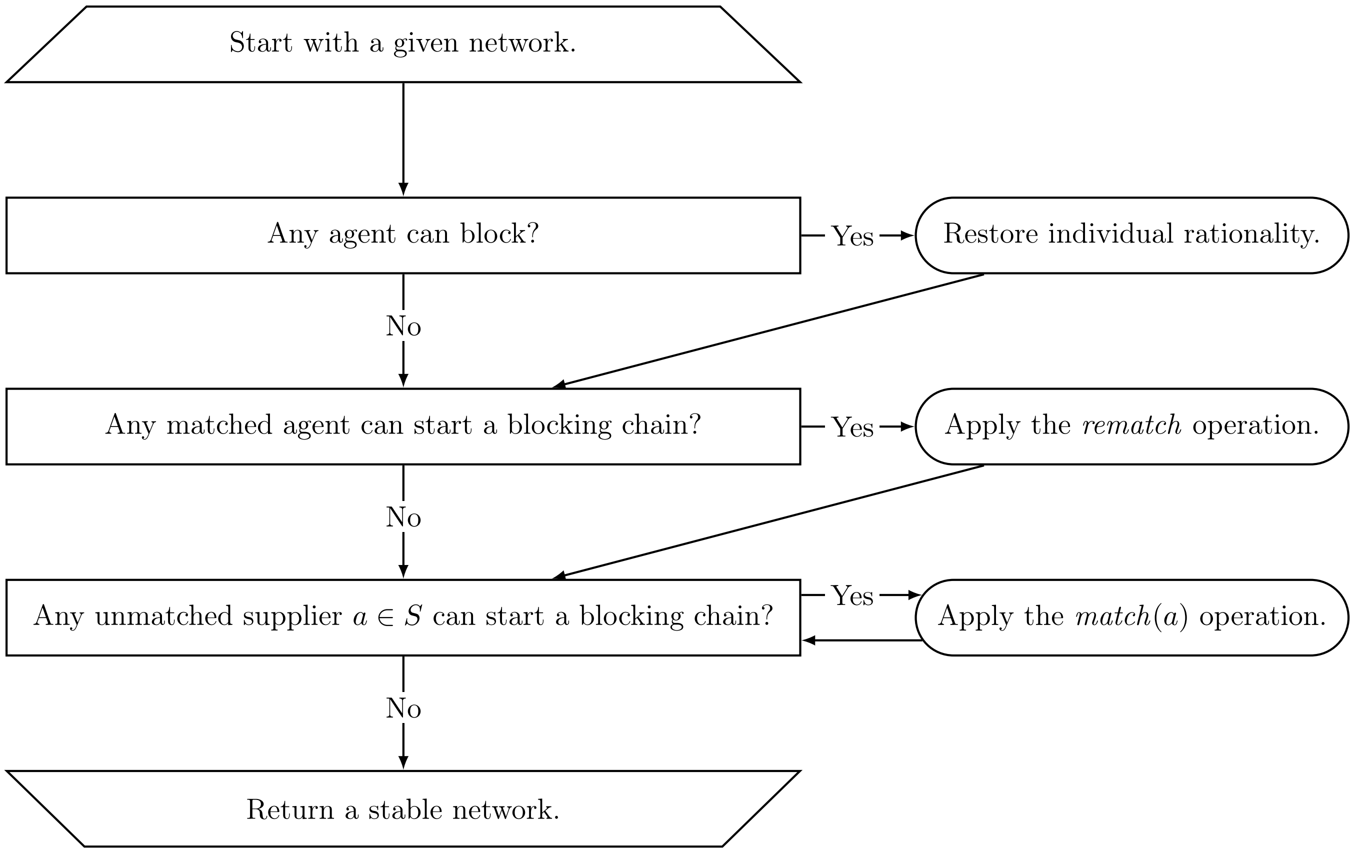

Figure 1: Algorithm to obtain a stable network starting from an arbitrary network

The following simple example illustrates how the algorithm works.

Example 1. Consider set $A = S \cup I \cup C$ of agents with two suppliers $S = \{s_1, s_2\}$, three intermediaries $I = \{i_1, i_2, i_3\}$, and two consumers $C = \{c_1, c_2\}$. Below are preferences for suppliers

$$s_1 : i_1 \succ i_2 \succ i_3,$$

$$s_2 : i_2 \succ i_1 \succ i_3,$$

consumers

$$c_1 : i_2 \succ i_3 \succ i_1,$$

$$c_2 : i_3 \succ i_2 \succ i_1,$$

and intermediaries

$$i_1 : s_2 \succ^U s_1 \text{ and } c_2 \succ^D c_1,$$

$$i_2 : s_1 \succ^U s_2 \text{ and } c_2 \succ^D c_1,$$

$$i_3 : s_1 \succ^U s_2 \text{ and } c_1 \succ^D c_2,$$

respectively. The preferences imply that suppliers cannot supply directly to consumers: supply chains have to involve intermediaries.

Consider the initial network

$\lambda = \lambda_1$

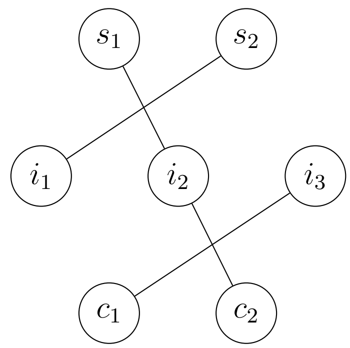

in which intermediary i_2 has both an upstream partner, supplier s_1, and a downstream partner, consumer c_2; while each of intermediaries i_1 and i_3 has only a single partner, supplier s_2 for i_1 and consumer c_1 for i_3.

It is not individually rational. In particular, i_1 and i_3 prefer to drop s_2 and c_1, respectively. By satisfying the corresponding blocking agents, we attain the individually rational network

λ_2

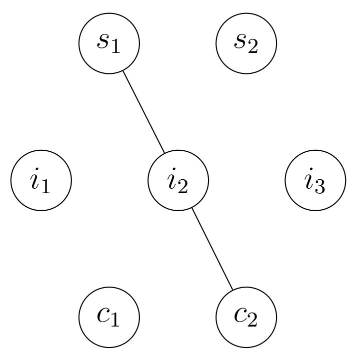

Next, we apply the *rematch* operation. We first satisfy the downstream-best blocking chain (s_1, i_1, c_1) starting with supplier s_1, the only matched supplier, and restore individual rationality—by severing the link between i_2 and c_2—to get

λ_3

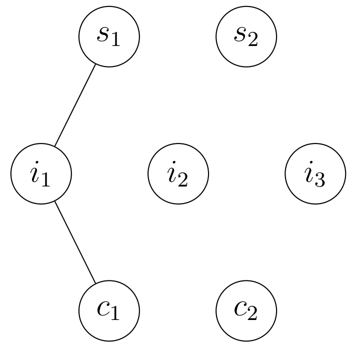

Then, we satisfy the downstream-best blocking chain (i_1, c_2) starting with intermediary i_1 to obtain the network

λ_4

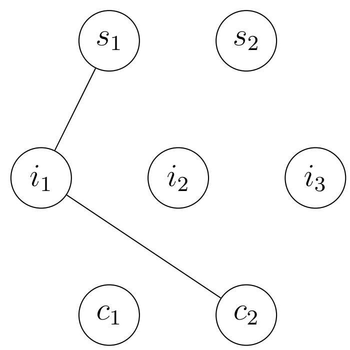

in which no matched agent can start a blocking chain.

Finally, we apply the *match*(s_2) operation for supplier s_2, the only unmatched supplier. We start by satisfying the downstream-best blocking chain (s_2, i_2, c_2) starting with supplier s_2 to attain

λ_5

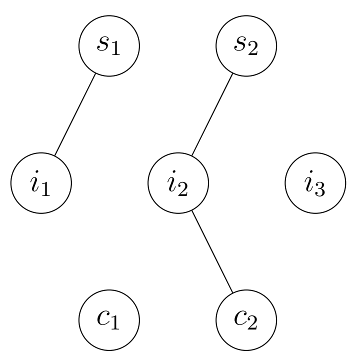

We then proceed by compensating intermediary i_1, c_2's previous upstream partner, and satisfying her downstream-best blocking chain (i_1, c_1) in order to attain the final network

$\mu = \lambda_6$

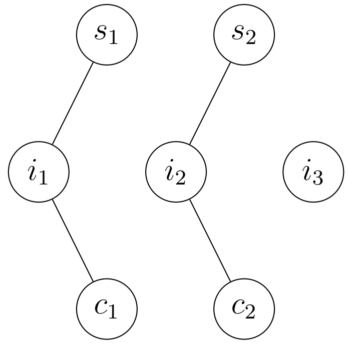

that is a stable network, as desired. △

When applied to an empty network, the introduced algorithm naturally extends the deferred acceptance algorithm (Gale and Shapley, 1962) that is at the basis of many centralized matching clearinghouses. Indeed, for an empty network, our algorithm reduces to the iterative application of the match operation. This procedure in turn coincides with the deferred acceptance algorithm for one-to-one matching markets.

We conclude by making several remarks on directions for future research. Recently, Rudov (2022) used a similar framework to examine the fragility of stable outcomes in two-sided markets. In particular, that paper shows that, under mild conditions, starting from any unstable outcome, the rematching dynamics may lead to any stable outcome. Analyzing the fragility of supply chain networks and ways to make them more resilient may be a promising direction of research. From a more technical perspective, we believe that analogous methods can be used to extend our main result to the model with contracts in which each agent can be involved in at most one contract as a supplier and at most one contract as a demander and has fully substitutable preferences over all possible contracts involving her.[8]

[8] Ostrovsky (2008) allows each agent to be involved in multiple contracts. In that case, one might need to devise a more sophisticated algorithm to prove the stabilization result. In fact, this result has not yet been established even for two-sided many-to-many matching markets in which agents from *both* sides have substitutable preferences (see Kojima and Ünver, 2008; Millán and Pepa Risma, 2018).

Appendix – Proofs

Proof of Lemma 2. For each $a \in A$, put a in A_i, where i is the length of the longest chain ending at a. Consider any $a \in A_i$ with $a \in U(b)$. Since there is a chain $(a_1, a_2, \ldots, a_i)$ of length i ending at $a_i = a$, we can obtain the chain $(a_1, a_2, \ldots, a_i, a_{i+1})$ of length $i+1$ ending at $a_{i+1} = b$. By construction, $b \in A_{i+1} \cup A_{i+2} \cup \ldots \cup A_n$, as stated. ∎

Proof of Lemma 3. Extend the *rematch* operation by printing the tier number $i < n$ after any time when we satisfy the downstream-best blocking chain starting with $a \in A_i$ and restore individual rationality. It is sufficient to show that no printed sequence can be infinitely long.

Suppose by contradiction that we have a sequence of sufficiently large length, to be defined in the course of the proof.

Note first that we can print 1 only finitely many times. Indeed, in the algorithm, we can only shrink the set of matched agents from A_1. In addition, when we print 1, one matched agent from A_1 strictly improves and other remaining matched agents from A_1 are not affected. Finally, when we print $i > 1$, no remaining matched agent from A_1 is affected. Therefore, we must have a sufficiently long consecutive (sub)sequence of printed numbers from $\{2, 3, 4, \ldots, n-1\}$.

Next, this new (sub)sequence can have only finitely many 2s. Indeed, since it contains no 1s, we can only shrink the set of matched agents from A_2. Similar to the previous paragraph, when we print 2, one matched agent from A_2 strictly improves and all other remaining matched agents from A_2 are not affected. Also, when we print $i > 2$, no remaining matched agent from A_2 is affected either. Hence, there must be a sufficiently long consecutive (sub)sequence of printed numbers from $\{3, 4, \ldots, n-1\}$.

By proceeding iteratively, we must find a sufficiently long consecutive (sub)sequence of printed numbers from $\{n-1\}$, leading to a contradiction. Indeed, when we print $n-1$, we can only shrink the set of matched agents from A_{n-1} and one such agent becomes better off;

thus, it is impossible to print $n - 1$ too many times in a row. ∎

Proof of Lemma 4. By the preference assumptions, no unmatched intermediary can start a blocking chain. As a result, no agent at all, matched or unmatched, can start a blocking chain. Hence, stability follows. ∎

Proof of Lemma 5. Consider λ and $a \in S = A_1$ from the lemma's statement. Satisfy the downstream-best blocking chain starting with supplier a. Suppose it ends at agent b.

Notice that no agent, if any, that is involved in the above blocking chain and distinct from a and b, was previously matched; otherwise, that intermediary—she has responsive preferences with respect to her rankings over her upstream and downstream partners—could start a blocking chain, contradicting the premise. Therefore, only b's previous upstream partner, if any, becomes worse off.

Now, consider state b and follow the instructions. Let c be b's previous upstream partner, if any. If b was previously unmatched or $c \in S = A_1$, then no intermediary/consumer is worse off compared to the previous network. Therefore, the obtained network satisfies the conclusion for i being defined as follows: i corresponds to the tier number that agent $b \in A_i$ belongs to. For instance, if b was previously unmatched, then $b \in C = A_n$ must be a consumer, and thus $i = n$.

Otherwise, $c \in I$ is an intermediary. Notice that no intermediary/consumer except c is worse off compared to the previous network, so that the obtained network has the following property, denoted by $(c.\star)$: no matched agent—including c herself—can start a blocking chain that does not involve c.

Suppose first that $c \in I$ can start a blocking chain. Then, satisfy the downstream-best blocking chain starting with c. Suppose it ends at agent d. As before, no agent, if any, that is involved in the previous blocking chain and distinct from c and d, was previously matched. Therefore, only d's previous upstream partner, if any, becomes worse off, compared to the previous network. Now, consider state d and let e be d's previous upstream partner, if any. By construction, no matched agent—except possibly e—can start a blocking chain involving

c; otherwise, in the network before the previous one, some matched agent could start a blocking chain. As a consequence, if d was previously unmatched or $e \in S = A_1$, then no intermediary/consumer is worse off compared to the previous network. Hence, the obtained network satisfies the conclusion for i defined as follows: $i = \max(i_b, i_d)$ corresponds to the largest tier number between agents $b \in A_{i_b}$ and $d \in A_{i_d}$. If instead $e \in I$, the obtained network satisfies the property $(e.\star)$, so that we can proceed recursively by verifying whether e can start a blocking chain or not.

If $c \in I$ cannot start a blocking chain, satisfy blocking agent c. Therefore, only c's previous upstream partner—in this case, c definitely had an upstream partner in the previous network—becomes worse off, compared to the previous network. Proceed to state c and let f be c's previous upstream partner. By construction, no matched agent can start a blocking chain involving c; otherwise, c could start a blocking chain as well. Hence, if $f \in S = A_1$, then no intermediary/consumer is worse off compared to the previous network, so the result follows for i corresponding to the tier number of agent $b \in A_i$. If instead $f \in I$, the obtained network satisfies the property $(f.\star)$, so that we can proceed iteratively by checking if f can start a blocking chain or not.

Thus, if the the process has not yet terminated, we can proceed recursively in either case. Why does the process terminate in a finite number of iterations?

Extend the match operation by printing the tier number $i > 1$ after any time when we satisfy the downstream-best blocking chain *ending* at $a \in A_i$. Similar to the proof of Lemma 3, it is sufficient to show that no printed sequence can be infinitely long.[9]

Suppose by contradiction that we have a sequence—of numbers $\{3, 4, \ldots, n-1, n\}$—of sufficiently large length, to be defined in the course of the proof. We do not consider 2s since in that case the process terminates necessarily. For the same reason, no printed n can correspond to a previously unmatched consumer.

Then, we can print n only finitely many times. Indeed, agents from A_n cannot improve

[9] Note that one can satisfy at most $n-3$ blocking agents in a row without the process being terminated.

indefinitely. Given that, we should have a sufficiently long consecutive (sub)sequence of printed numbers from $\{3, 4, \ldots, n-1\}$. Since along this (sub)sequence, no agent from A_{n-1} can become worse off and one of them becomes better off when we print $n-1$, we can print $n-1$ only finitely many times. By proceeding iteratively, we can obtain a sufficiently long consecutive (sub)sequence of printed numbers from $\{3\}$. This leads to a contradiction.

To sum up, the process necessarily terminates. By construction, the final network is individually rational and no matched agent can start a blocking chain. Furthermore, it satisfies conditions (i) and (ii) for i defined as follows: i corresponds to the largest printed number in the above defined sequence. ∎

Proof of Lemma 6. By Lemmas 4 and 5, it suffices to show that we cannot apply the match operation indefinitely.

To see this, after any match operation, print its corresponding number i from the Lemma 5's statement. By the same argument as in the proof of that lemma, any printed sequence can be only finitely long, thus concluding the proof. ∎

References

Ackermann, Heiner, Paul W Goldberg, Vahab S Mirrokni, Heiko Röglin, and Berthold Vöcking, "Uncoordinated Two-Sided Matching Markets," *SIAM Journal on Computing*, 2011, *40* (1), 92–106.

Amelkin, Victor and Rakesh Vohra, "Strategic Formation and Reliability of Supply Chain Networks," *mimeo*, 2020.

Chen, Yi-Chun and Gaoji Hu, "Learning by Matching," *Theoretical Economics*, 2020, *15* (1), 29–56.

Chung, Kim-Sau, "On the Existence of Stable Roommate Matchings," *Games and Economic Behavior*, 2000, *33* (2), 206–230.

Diamantoudi, Effrosyni, Eiichi Miyagawa, and Licun Xue, "Random Paths to Stability in the Roommate Problem," *Games and Economic Behavior*, 2004, *48* (1), 18–28.

Elliott, Matthew, Benjamin Golub, and Matthew V Leduc, "Supply Network Formation and Fragility," *American Economic Review*, 2022, *112* (8), 2701–47.

Gale, David and Lloyd S Shapley, "College Admissions and the Stability of Marriage," *The American Mathematical Monthly*, 1962, *69* (1), 9–15.

Hatfield, John William and Scott Duke Kominers, "Matching in Networks with Bilateral Contracts," *American Economic Journal. Microeconomics*, 2012, *4* (1), 176.

_ , _ , Alexandru Nichifor, Michael Ostrovsky, and Alexander Westkamp, "Stability and Competitive Equilibrium in Trading Networks," *Journal of Political Economy*, 2013, *121* (5), 966–1005.

_ , _ , _ , _ , and _ , "Chain Stability in Trading Networks," *Theoretical Economics*, 2021, *16* (1), 197–234.

Inarra, Elena, Concepcion Larrea, and Elena Molis, "Random Paths to P-Stability in the Roommate Problem," *International Journal of Game Theory*, 2008, *36* (3), 461–471.

Klaus, Bettina and Flip Klijn, "Paths to Stability for Matching Markets with Couples," *Games and Economic Behavior*, 2007, *58* (1), 154–171.

Kojima, Fuhito and M Utku Ünver, "Random Paths to Pairwise Stability in Many-to-Many Matching Problems: a Study on Market Equilibration," *International Journal of Game Theory*, 2008, *36* (3), 473–488.

Lazarova, Emiliya and Dinko Dimitrov, "Paths to Stability in Two-Sided Matching under Uncertainty," *International Journal of Game Theory*, 2017, *46* (1), 29–49.

Millán, Beatriz and Eliana Pepa Risma, "Random Path to Stability in a Decentralized Market with Contracts," *Social Choice and Welfare*, 2018, *51* (1), 79–103.

Ostrovsky, Michael, "Stability in Supply Chain Networks," *American Economic Review*, 2008, *98* (3), 897–923.

Roth, Alvin E. and John H. Vande Vate, "Random Paths to Stability in Two-Sided Matching," *Econometrica*, 1990, *58* (6), 1475–1480.

_ **and Marilda A. Oliveira Sotomayor**, *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis* Econometric Society Monographs, Cambridge: Cambridge University Press, 1992.

Rudov, Kirill, "Fragile Stable Matchings," *mimeo*, 2022.