Pairwise stable matching in large economies*

Michael Greinecker⋆ and Christopher Kah‡

April 18, 2018

Abstract

We formulate a general model and stability notion for two-sided pairwise matching problems with individually insignificant agents. Matchings are formulated as joint distributions over the characteristics of the populations to be matched. These characteristics can be high-dimensional and need not be included in compact spaces. Stable matchings exist with and without transfers and stable matchings correspond exactly to limits of stable matchings for finite agent models. We can embed existing continuum matching models and stability notions with transferable utility as special cases of our model and stability notion. In contrast to finite agent matching models, stable matchings exist under a general class of externalities. This might pave the way for integrating matching problems in other economic models.

JEL-Classification: C62, C71, C78, D47.

Keywords: Stable matching; economies in distributional form; large markets.

1 Introduction

This paper provides a theoretical model of stable pairwise matchings in two-sided matching markets based on the joint statistical distribution of the characteristics of the agents involved. Stable matchings exist in full generality with and without transfers between agents and even in the presence of externalities. The stable matchings in our model exactly capture the limit behavior of stable matchings in large finite matching markets in terms of the joint distribution of characteristics. We also show that the model could be reformulated in terms of individual agents.

In the traditional theory of stable matching, agents can be split into two sides. A matching, in the simplest setting, specifies which agent on one side is matched to which agent on the other side, if at all. The matching is stable if no two agents on opposing sides would rather be matched to each other than to their current matches, and no agent

*We are grateful for helpful discussions with Lars Ehlers, Christoph Kuzmics, Jacob Leshno, Martin Meier, Idione Meneghel, Georg Nöldeke, Konrad Podczeck, Karolina Vocke, Rakesh Vohra, Dov Samet, Larry Samuelson, Alexander Teytelboym, Markus Walzl, Alexander Westkamp, seminar participants at the University of Zürich, at the University of Warwick, and at the Warsaw School of Economics, participants at the 2015 SAET conference in Cambridge, and participants at the 2017 Match-Up Conference in Cambridge MA.

⋆Department of Economics, University of Graz, Austria. Email: michael.greinecker@uni-graz.at

‡Department of Economics, University of Innsbruck, Austria. Email: christopher.kah@uibk.ac.at

would rather be alone. This is, in essence, the "marriage model" of Gale and Shapley (1962). Extensions allow for transfers between the matched agents as in Shapley and Shubik (1971) and Demange and Gale (1985). Stable matching models provide the natural frictionless benchmark for the analysis of markets in which interacting participants can neatly be divided into two groups, characteristics of participants matter, and all interaction is between matched agents. One can match workers to firms, students to schools, and medical residents to hospitals. None of these markets can be separated from the rest of the economy, but integrating matching theory with other economic models poses serious technical problems due to matching theory being fundamentally discrete. Wine, milk, and hours worked are divisible, people are not. At each point in time, a person can only be in one place and traveling is costly. From the perspective of matching theory, interactions with the rest of the economy amount to externalities within the matching market, but the combination of indivisibilities and externalities in matching problems is in general not compatible with the existence of stable matchings. We deal with the problem the way most economists deal with the molecular indivisibilities found in wine and milk: We scale the analysis so that even people look divisible. Matching markets are taken to be so large that individuals are negligible relative to market size. The economy is large, people are not. Our approach is guided by the following desiderata:

Existence Stable matchings should exist. Otherwise, our model models nothing.

Generality We want our model to provide a general, unifying framework. As such we allow for non-transferable utility, transferable utility, and everything in between. Characteristics need not lie in a finite, one-dimensional, or compact set.

Approximability A model in which agents are negligible relative to market size is at best an approximation to a model in which finitely many agents have little influence. The large market model should be interpretable in terms of the limit behavior of finite agent models.

Embeddability Our model should be combinable with other economic models and matching markets should not just be seen in isolation. For this purpose, we need to allow for general externalities as a first step.

Compatibility There is an existing literature on stable matching in large markets under transferable utility. We do not want to provide an alternative to this literature, we want to nest it.

We show that these desiderata can be satisfied by choosing the right model and the right stability notion. The most popular approach to modeling large economies is the individualistic approach introduced by Aumann (1964) in general equilibrium theory and by Schmeidler (1973) in non-cooperative game theory. In this framework, the economy is represented by a nonatomic probability space of agents and a function that maps agents to their characteristics. We show below in Examples 1 and 2 that naively adopting this approach to matching theory gives us a model in which stable matchings need not exist. Such an individualistic model is also hard to relate to finite agent models. If the number of agents changes, the dimension of the model changes and there is no common space for comparisons.

These problems can be overcome by adopting the distributional approach to modeling large economies. The distributional approach was introduced by Hart, Hildenbrand, and Kohlberg (1974) and Hildenbrand (1975) in general equilibrium theory and by Mas-Colell (1984) in non-cooperative game theory. The distributional approach disposes of the agents completely and only uses the distribution of their characteristics as the data of the model. Since one can compare distributions of characteristics independently of the set of agents having the characteristics, this allows us to relate the limit model and its stable matchings to finite agent matching models and their stable matchings. The distributional approach has already proven useful in matching with transferable utilities and even in unifying the econometric treatment of stable matchings, see Chiappori and Salanié (2016).[1] As we show, one can go much further.

We also get new problems with the distributional approach. Once we have disposed of the agents, we have disposed of blocking pairs of agents too and therefore of the usual stability notion. In a finite agent model, ignoring individual rationality constraints and unbalanced markets, a matching is stable if and only if there is no blocking pair of agents and this means that one cannot find a blocking pair of characteristics with positive mass in the distribution of joint characteristics of matched pairs. With a continuum of characteristics, which we allow for, there may be no pair of characteristics with positive mass at all in a matching, even though such a matching may be intuitively far from stable; see Example 3 below. With transferable utility, there are known equivalences of stability that can still be applied, but they are of no use without transfers.

We show that there is a way to detect the hidden instabilities: random sampling. If we randomly select two pairs of matched characteristics types, the probability that blocking is possible among these pairs must be zero in a stable matching. This definition of stability coincides with the usual one for finite matching models but is applicable in much wider settings. Stable matchings exist for this stability notion under extremely weak conditions (Theorem 1) and exist even when we allow for continuously varying but otherwise arbitrary externalities (Theorem 7). Moreover, stable matchings in the large economy correspond exactly to limits of stable matchings in finite matching models that approximate the large economy under the topology of weak convergence of measures (Theorem 2).

In our model, agents of the same type may not end up equally well off in a stable matching; see Example 4 below. In particular, there is in general no function assigning types to payoffs for a stable matching. In the transferable utility context, such functions have traditionally been used to define stability. In Section 7 we present a number of results that show that if some (not necessarily perfect) transfers are possible, such functions always exist and that, using these functions, our stability notion coincides with the one used in the transferable utility context.

Our model has no agents on the formal level, but we will talk a lot about what agents do in our model for heuristic purposes. For readers worried by our talk about fictional entities, we show in Theorem 9 that our model could be reinterpreted in terms of actual

[1] We point out some problems in using our approach for econometric purposes in the conclusions.

agents; our distributional model can be represented by an individualistic one.

Some of the simpler and more instructive proofs are found within the text, the other proofs are in a section of their own, Section 11. Mathematically, we rely heavily on concepts from the topology of metrizable spaces and weak convergence of measures. A mathematical appendix at the end contains the essentials.

Related Literature

Much of the literature on stable matchings in large economies focuses on the asymptotic behavior of large finite matching models without ever using a limit model. We refer to Kojima (2017) for an overview of this approach; here we focus on limit matching models for which existence results are available. Our paper is the first to prove existence of stable matchings in a large economy framework that allows for (but does not require) non-transferable utility with general spaces of characteristics. Other large economy models of stable matching fall into three categories:

1. Finite-type models. Baïou and Balinski (2002), Echenique, Lee, and Shum (2010) and Echenique, Lee, Shum, and Yenmez (2013) prove, among other things, the existence of stable matchings in a distributional marriage model with only finitely many types. Azevedo and Hatfield (2015) prove existence of stable matchings and nonemptiness of the core for a many-to-many matching model with finitely many contracts and types. In comparison to finite agent models, their model allows for complementarities. Galichon, Kominers, and Weber (2016) prove existence of stable matchings in a general two-sided finite type model that allows for both non-transferable utility and (imperfectly) transferable utility and can be used for the econometric estimation of matching models.

2. One-sided continuum models. Azevedo and Leshno (2016) have an individualistic model in which a continuum of students is matched to a finite number of colleges and use it to prove that there is generically a unique stable matching under a richness condition on preferences. Che, Kim, and Kojima (2015) study a similar setting in a distributional framework and show using fixed-point methods from nonlinear functional analysis that stable matchings exist even with complementarities. Both papers require the characteristics of agents to lie in a compact set, which rules out wealth, for example, as a match relevant characteristic when the distribution of wealth is not bounded. The definition of stability is not an issue if there are only finitely many agents on one side, so we have to deal with additional conceptual problems.

3. Two-sided continuum models with transfers. In the case of (perfectly) transferable utility and finitely many agents, stable matchings can be identified with solutions to the dual of a linear programming problem. There exists an infinite dimensional version of this linear programming problem in which one optimizes over spaces of measures, the optimal transport problem of Kantorovich.[2] Using a duality result for the Kantorovich optimal

[2] An introduction to optimal transport theory geared towards economists and econometricians is given by Galichon (2016). For advanced material on the mathematics of optimal transport, see Villani (2003) and Villani (2009).

transport problem, Gretsky, Ostroy, and Zame (1992, 1999) develop a distributional approach to stable matching under transferable utility for general compact metric type spaces. Further results and the generalization to separable, completely metrizable type spaces are provided by Chiappori, McCann, and Nesheim (2010).

An individualistic model of cooperative games that subsumes stable matching with imperfectly transferable utility is given by Kaneko and Wooders (1986). An existence result for the setting under the assumption that characteristics lie in a compact metric space is given by Kaneko and Wooders (1996). Feasible payoffs in the model may only be approximately realizable though, so the model is best interpreted as a model of approximately stable matchings (or cores, in the more general setting).

Closest to our model is the distributional model of stable matching with imperfectly transferable utility and compact metric type spaces by Nöldeke and Samuelson (2017).[3] Nöldeke and Samuelson develop a general nonlinear duality theory and apply it to contract theory and matching theory. Their matching model is less general than ours, but the additional structure of their model allows them to obtain results on the lattice structure of stable matchings.

With transferable and imperfectly transferable utility, one traditionally defines stable matchings in terms of the payoffs a type gets in a matching and this is how all papers mentioned above define stability. This requires that payoffs in a matching depend on types only, there is equal treatment of types and this is assumed in the papers just mentioned. Without transfers, this is generally not possible as we show in Example 4. However, we can actually prove in our general model that equal treatment and a stronger form of equal treatment is a consequence of stability and transfers, see Section 7. This allows us to show that our stability notion is equivalent to the one used by Gretsky, Ostroy, and Zame (1992, 1999), Chiappori, McCann, and Nesheim (2010), and Nöldeke and Samuelson (2017).[4]

Another literature deserves mention: The literature on ex-ante investments in competitive matching markets. Many match-relevant investments are made before agents join matching markets. People go to university before they know which firm they are going to work for eventually. Under imperfect competition, there will be a hold-up problem and resulting inefficiencies. To isolate whether other inefficiencies are possible, one needs a competitive matching model in which the classic hold-up problem cannot occur, and perfect competition rarely works with finitely many agents. A number of papers has carried out this program, the following list is not complete. Cole, Mailath, and Postlewaite (2001) and Iyigun and Walsh (2007) study investments under transferable utility in the optimal transport framework. Peters and Siow (2002) study investment in the non-transferable utility context in which characteristics are one-dimensional and all agents rank agents on the other side the same way. Matchings are assumed to be assortative, but no explicit stability argument is given. We see in Example 3 that our model and stability notion provide appropriate foundations. Nöldeke and Samuelson (2015) study investment under imperfectly transferable utility. The present paper provides a unified competitive matching framework for this literature.

[3]We became aware of each others' work only in July 2016 at the Game Theory Society World Congress.

[4]In the optimal transport context, we only nest those models in which the surplus function is continuous.

2 Motivating Examples

Our first two examples illustrate what can go wrong if one simply transfers standard definitions from the matching literature to an individualistic model with a continuum of agents: There may be no stable matching. The first example illustrates the problem in a setting without transfers, the second example illustrates the problem in a setting with transfers.[5] In both examples, a stable matching ought to be perfectly assortative, but the indivisibility of agents stands in the way of a fully assortative matching. This is no problem for an approach in which agents are replaced by types and matchings are joint distributions of types (and surplus, in the case of transfers.) Indeed, we cannot know how many agents lie behind a given type in a distributional framework, so types can be freely split, they are perfectly divisible.

Example 1. Let the set of agents on one side of the market be $A_W = [0, 1]$ and the set of agents on the other side of the market be $A_M = [0, 1/2] \times \{1, 2\}$. All agents in A_W have the same preferences and all agents in A_M have the same preferences. Every agent in A_W is indifferent between $(x, 1)$ and $(x, 2)$, but prefers $(x, 1)$ to $(y, 1)$ if $x > y$. Agents in A_M prefer x to y if $x > y$. A matching, as traditionally defined, pairs off agents and is therefore in the present context simply a bijection $f : A_W \to A_M$. Since stable matchings will not exist in this example, we refrain from putting even more restrictions on a matching. The matching f is stable if there is no blocking pair, that is, there are no $a \in A_W$ and $b \in A_M$ such that a prefers b to $f(a)$ and b prefers a to $f^{-1}(b)$.

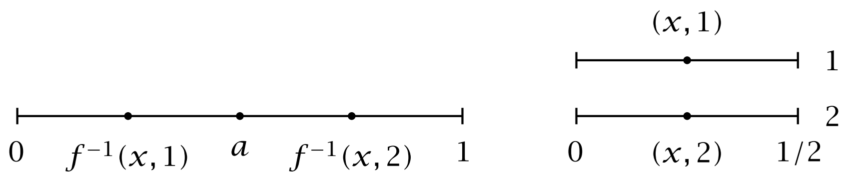

Take a look at the picture above. We let f be any matching; we show it is not stable. Pick any $x \in [0, 1/2]$. Without loss of generality, assume that $f^{-1}(x, 1) < f^{-1}(x, 2)$. Let $a \in A_W$ satisfy $f^{-1}(x, 1) < a < f^{-1}(x, 2)$ and let $(y, m) = f(a)$. We must either have $y > x$ or $y < x$. If $y > x$, then $f^{-1}(x, 2)$ and (y, m) can block the matching. If $y < x$, then a and $(x, 1)$ can block the matching.

Example 2. Let again the set of agents on one side of the market be $A_W = [0, 1]$ and the set of agents on the other side of the market be $A_M = [0, 1/2] \times \{1, 2\}$. But now we consider a transferable utility model in which the surplus generated by $a \in A_W$ and $(x, m) \in A_M$ when matched is ax. Matched agents are free to divide the surplus between them any way they see fit. A matching can now be represented by a bijection $f : A_W \to A_M$ and functions $V_W : A_W \to \mathbb{R}_+$ and $V_M : A_M \to \mathbb{R}_+$ such that $V_M(x, m) + V_W(f^{-1}(x, m)) = x f^{-1}(x, m)$ for all $(x, m) \in A_M$. The matching f is stable if there is no blocking pair, that is, there are no $a \in A_W$ and $(x, m) \in A_M$ such that $V_W(a) + V_M(x, m) < ax$.

[5]Other examples of non-existence of “individualistic” matchings under transferable utility are Example 5 in Gretsky, Ostroy, and Zame (1992) and Example 8 in Chiappori, McCann, and Nesheim (2010). The example given here is somewhat simpler. Our examples rely only on the order structure of the unit interval and make no use of measure theory. They could be adapted to any order dense sets of agents, say the rationals. If preferences are dually well-ordered, stable matchings would exist even for infinite populations by Theorem 2 of Fleiner (2003).

Let f, V_W, and V_M be any matching; we show it is not stable. Pick any $x \in [0, 1/2]$. Without loss of generality, assume that $f^{-1}(x, 1) < f^{-1}(x, 2)$. Let $a \in A_W$ satisfy $f^{-1}(x, 1) < a < f^{-1}(x, 2)$ and let $(y, m) = f(a)$. We must either have $y > x$ or $y < x$. We look at the case $y > x$. If f were stable, we would have both $V_W(a) + V_M(x, 2) \geq ax$ and $V_W(f^{-1}(x, 2)) + V_M(y, m) \geq f^{-1}(x, 2)y$. But then, using the supermodularity of the surplus function,

$$V_W(a) + V_M(x, 2) + V_W(f^{-1}(x, 2)) + V_M(y, m) \geq ax + f^{-1}(x, 2)y$$

$$> ay + f^{-1}(x, 2)x = V_W(a) + V_M(y, m) + V_W(f^{-1}(x, 2)) + V_M(x, 2),$$

which is impossible. A similar argument shows $y < x$ is not compatible with f being stable. Hence, f is not stable.

The next example shows how our notion of stability based on random sampling can be applied even when no “blocking pairs” have positive mass. It is similar to a distributional version of Example 1. Agents are not modeled now, but the distribution of their types is. As a by-product of the analysis, we obtain foundations for the assortativity assumption in Peters and Siow (2002) in terms of an explicit stability argument.

Example 3. The relevant characteristics of agents on both sides of the market can be summarized in a one-dimensional variable and the distributions of characteristics is uniform on $[0, 1]$.[6] Moreover, everyone prefers to be matched to someone with a higher number and being matched to not being matched at all. The distribution of characteristics of matched pairs is then given by a (Borel) probability measure μ on $[0, 1] \times [0, 1]$ with uniform marginals. Given the structure of the problem, we would expect positive assortative matching to be the only stable matching. If F is the two-dimensional cumulative distribution function of μ, this amounts to $F(x, y) = \min\{x, y\}$ for all $x, y \in [0, 1]$. Now, μ cannot have atoms, since the marginals have no atoms. So every matching would be stable if all we required is that there are no (w, m) and (w', m') with positive μ-mass such that $w > w'$ and $m' > m$ or such that $w' > w$ and $m > m'$. However, if μ is not positively assortative, we can find the hidden blocking matched pairs by random sampling with positive probability. Indeed, suppose $F(x, y) \neq \min\{x, y\}$. It follows from the marginal conditions that $F(x, y) \leq \min\{x, y\}$, so we must have $F(x, y) < \min\{x, y\}$. Given the symmetry of the problem, we can assume without loss of generality that $x \leq y$. So, with obvious notation, $\mu(w \leq x, m \leq y) < x$ and, a fortiori, $\mu(w \leq x, m \leq x) < x$. Together with the marginal condition

$$\mu(w \leq x, m \leq x) + \mu(w \leq x, m > x) = x,$$

we get $\mu(w \leq x, m > x) > 0$, which is equivalent to $\mu(w < x, m > x) > 0$. Similarly, from combining $\mu(w \leq x, m \leq x) < x$ with the marginal condition

$$\mu(w \leq x, m \leq x) + \mu(w > x, m \leq x) = x,$$

[6]The restriction to uniform distributions of characteristics is less restrictive than it might seem. As long as the distributions of characteristics on the real line have no atoms, the same argument works by interpreting x and y not as characteristics but as quantiles of characteristics. The relevant material on quantile functions and copulas may be found in Galichon (2016, Appendix C).

we get $\mu(w > x, m \leq x) > 0$, which is equivalent to $\mu(w > x, m < x) > 0$. Now $\mu(w < x, m > x) > 0$ and $\mu(w > x, m < x) > 0$ imply that we can detect instability of any matching that is not positively assortative with positive probability by independent random sampling.

The next example illustrates why our notion of stability cannot be based on a strong form of blocking.

Example 4. There is the same number of agents on both sides. One side has agents with types in the set $W = \{h, l\}$, equal in number, agents on the other side can have only one type; $M = \{m\}$. Every agent prefers being matched, no matter to whom, to being unmatched. An agent of type m strictly prefers to be matched with an agent of type h to being matched with an agent of type l. Since every agent prefers to be matched to being unmatched and every agent can be matched, a stable matching should not allow for unmatched agents in this case. But this means that half of all agents of type m are matched with an agent of type h, and half of all agents of type m are matched with an agent of type l. There is rationing of agents of type m and not all agents with the same type are equally well off.

Consider an agent of type h and an agent of type m matched to an agent of type l. Matching these agents would make the former agent not worse off and the latter agent strictly better off. If we want stable matchings to always exist, these two agents must not form a blocking pair. We will therefore consider a matching already to be stable if there are no two agents who strictly prefer to be matched to each other to their current position in the matching.

Example 4 depends crucially on the absence of transfers. Indeed we show in Section 7 in a setting with (imperfect) transfers under fairly weak conditions, that in any stable matching two agents with the same type will be equally well off and two agents with similar types will be similarly well off.

That not all preferences are strict might seem strange from the perspective of traditional matching theory, but is unavoidable if we want to admit general spaces of characteristics. If these characteristics are given by a Euclidean space with dimension larger than one, rational preferences cannot be both continuous and strict.[7]

3 The Environment

The model-relevant characteristics of agents on both sides of the market are given by nonempty sets of *types* W and M, respectively. We describe most of the theory in terms of a heterosexual marriage market in which W represents the types of women, M the types of men, and only marriages between women and men are under consideration. To allow agents to stay single, we also use extended spaces $W_ø = W \cup \{ø\}$ and $M_ø = M \cup \{ø\}$ with $ø \notin W \cup M$. An agent matched with an agent of type ø is really just single. We call a

[7]Pick a continuous utility representation. If preferences where strict, the representation would be injective and therefore a local homeomorphism between open sets of different dimensions. This is impossible.

pair in $W_\emptyset \times M_\emptyset$ a *couple type*. Matched agents can engage in activities and the activities available to them may depend on their couple type. Formally, there is a set A of *activities* and an *activity correspondence* $C : W_\emptyset \times M_\emptyset \to 2^A$ specifying the set of activities actually available to the matched agents. Depending on the context, activities might stand for intra-household allocations, transfers, wages, or contracts. A triple (w, m, a) with (w, m) a couple type and $a \in A$ is a *couple-activity type*. Types also specify (strict) preferences. For each $w \in W$, there is a relation $\succ_w$ on $M_\emptyset \times A$, and for each $m \in M$, there is a relation $\succ_m$ on $W_\emptyset \times A$. That $(m, a) \succ_w (m', a')$ means that a woman of type w prefers to engage in activity a with a man of type m to engaging in activity a' with a man of type m'. Similarly, $(w, a) \succ_m (w', a')$ means that a man of type m prefers to engage in activity a with a woman of type w to engaging in activity a' with a woman of type w'. Types play three roles: They specify the activities available to a couple, the preferences of an agent, and the characteristics of an agent that agents on the other side of the market might care about. We endow the sets W, M, and A with Polish (separable and completely metrizable) topologies, providing us with a notion of closeness for types and activities. We also endow $W_\emptyset$ and $M_\emptyset$ with the topologies that make $\emptyset$ an isolated point and such that the topologies on W and M are just the respective subspace topologies. For notational convenience, we let $\succ_\emptyset$ be the empty relation under which no elements are comparable, so that $\succ_w$ and $\succ_m$ are defined even when $w = \emptyset$ or $m = \emptyset$, respectively. Throughout the paper, we make the following three assumptions.

Acyclicity of Preferences: The relation $\succ_w$ on $M_\emptyset \times A$ is acyclic for each $w \in W$ and the relation $\succ_m$ on $W_\emptyset \times A$ is acyclic for each $m \in M$.[8]

Continuity of Preferences: The set

$$\{(m, a, m', a', w) \mid (m, a) \succ_w (m', a')\}$$

is open in

$$M_\emptyset \times A \times M_\emptyset \times A \times W$$

and the set

$$\{(w, a, w', a', m) \mid (w, a) \succ_m (w', a')\}$$

is open in

$$W_\emptyset \times A \times W_\emptyset \times A \times M.$$

Regularity of the Activity Correspondence: The correspondence $C : W_\emptyset \times M_\emptyset \to 2^A$ is continuous with nonempty and compact values.

The assumption that preferences are acyclic is extremely weak but will suffice for most of our results. The continuity assumption on preferences ties the notion of closeness specified by the topologies on W, M, and A to how the agents themselves view the types.[9]

[8]Recall that the relation $\succ$ on the set S is *acyclic* if there is no finite sequence $\langle s_1, s_2, \ldots, s_n \rangle$ with values in S such that $s_1 \succ s_2 \succ \cdots \succ s_n \succ s_1$.

[9]Our continuity condition is equivalent to the functions mapping w to $\succ_w$ and m to $\succ_m$, respectively, being continuous in the Kannai topology introduced in Kannai (1970), which can be equivalently defined to be the weakest topology that makes these preference functions continuous. However, the Kannai topology

In practice, it is usually more convenient to work with utility functions than preferences. Indeed, a sufficient condition for both assumptions on preferences to be satisfied is the existence of continuous functions $u_W : W \times M_\emptyset \times A \to \mathbb{R}$ and $u_M : W_\emptyset \times M \times A \to \mathbb{R}$ such that

$$(m, a) \succ_w (m', a') \text{ if and only if } u_W(w, m, a) > u_W(w, m', a')$$

and

$$(w, a) \succ_m (w', a') \text{ if and only if } u_M(w, m, a) > u_M(w', m, a').$$

If all preferences are asymmetric and negatively transitive (and therefore the asymmetric part of complete and transitive preference relations), and W, M, and A are Euclidean or, more generally, locally compact spaces, then this sufficient condition is also necessary by a Theorem of Mas-Colell (1977). Preferences can then always be represented by such parametrized jointly continuous utility functions. We will use such utility representations frequently in examples and assume them to hold when discussing transferable utility and imperfectly transferable utility.

Finally, to close the model, we specify nonzero, finite, Borel *population measures* ν_W and ν_M on W and M, respectively. In general, we denote the space of finite Borel measures on a Polish space X by $\mathcal{M}(X)$, so $\nu_W \in \mathcal{M}(W)$ and $\nu_M \in \mathcal{M}(M)$. We do allow for unbalanced markets in which $\nu_W(W) \neq \nu_M(M)$ and the population measures need not be probability measures. Our model is invariant to normalizing both measures jointly, so the absolute numbers $\nu_W(W)$ and $\nu(M)$ are economically meaningless. However, their relative size $\nu_W(W)/\nu_M(M)$ is meaningful and represents the number of women per men in the analyzed population. Finite agent models correspond to the case in which ν_W and ν_M both have finite support and all values are rational numbers. Indeed, by multiplying each of these rational numbers by the least common multiple of all denominators one obtains a model in which types (that are actually present) occur in positive integer quantities.

4 Specific Environments

We show in this section that widely used models with non-transferable utility, (perfectly) transferable utility, and imperfectly transferable utility can be represented in our model as special cases.

In the classic marriage model of Gale and Shapley (1962), no transfers or activities are allowed. In our framework, this corresponds to the degenerate case of a single activity $A = \{a\}$ and the activity correspondence C having the constant value $A = \{a\}$. Notationally, we will suppress A and C when working with the classic marriage model.

Assume now that preferences are given by jointly continuous functions $u_W : W \times M_\emptyset \times A \to \mathbb{R}$ and $u_M : W_\emptyset \times M \times A \to \mathbb{R}$. Suppose we require every couple type to make (weakly) efficient activity choices, so that not both agents involved could be made better off by choosing a different activity. If we only care about the utilities obtained, we

may fail to be even Hausdorff without some form of local non-satiation, so we will not work directly with the Kannai topology might let the couple choose directly from the utility possibility frontier, a space naturally homeomorphic to the unit interval under reasonable conditions.[10] We can then assume without loss of generality that $A = [0, 1]$ and that C is constant with value A. The rest of this section will focus on this case. The basic idea behind the natural homeomorphism is given in the picture below.

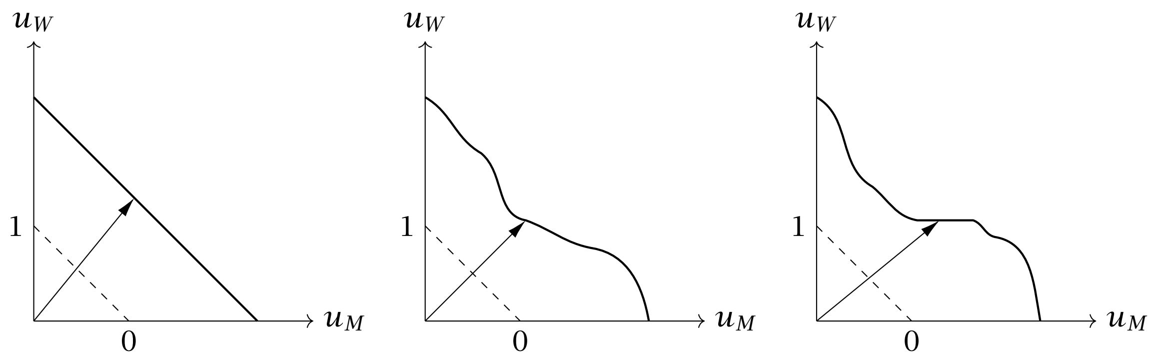

The utility possibility frontier is homeomorphic to $[0, 1]$.

If the utility possibility frontier is a line with constant slope -1 for all $w \in W$ and $m \in M$, as in the left diagram, we have a model of *transferable utility*. Matching with transferable utility has a long tradition, starting with the work of Koopmans and Beckmann (1957) and Shapley and Shubik (1971). With transferable utility, the value $u_W(w, m, a) + u_M(w, m, a)$ does not depend on a and we can define a surplus function $S : W_\emptyset \times M_\emptyset \to \mathbb{R}$ by

$$S(w, m) = u_W(w, m, 1) + u_M(w, m, 1) - u_W(w, \emptyset, 1) - u_M(\emptyset, m, 0).$$

The surplus function carries all relevant information for the analysis of stability. As formulated here, transferable utility depends on specific choices of the functions u_W and u_M and is not invariant under transformations that preserve ordinal rankings. The exact ordinal implications of transferable utility have recently been characterized by Chiappori and Gugl (2015).[11] Matching markets with transferable utility have a very special structure. Maximizing aggregate surplus is a linear program and the solutions to the dual program can be interpreted as stable matchings or market equilibria. This duality holds even if types are not discrete, so transferable utility is the one case in which existence is known for two-sided continuum models with general types, Gretsky, Ostroy, and Zame (1992, 1999), and Chiappori, McCann, and Nesheim (2010).

For the purposes of this paper equally well-behaved is the more general case of *imperfectly transferable utility*, as depicted in the middle figure. It is characterized by $u_W(w, m, \cdot)$ being an increasing[12] function and $u_M(w, m, \cdot)$ being a decreasing function for all $w \in W$ and $m \in M$. With imperfectly transferable utility, weak efficiency coincides with the more stringent usual definition of efficiency: It is impossible to make any agent in a couple better off without making the other worse off. Matching markets with imperfectly

[10]Essentially, the space of possible utility pairs should be closed, bounded above, comprehensive from below, and the outside options should lie in the relative interior. See Mas-Colell (1989, Proposition 4.6.1) for a proof.

[11]For a gentle exposition, see Section 3.1 in Chiappori (2017).

[12]We take increasing functions to be strictly increasing and decreasing functions to be strictly decreasing.

transferable utility have been studied in detail by Demange and Gale (1985). Existence proofs for such markets with finitely many agents may be found in Crawford and Knoer (1981) and Quinzii (1984). Parallel to our work, existence of stable matchings with imperfectly transferable utility in a distributional model with continuum (but compact) type spaces was shown by Nöldeke and Samuelson (2017). Nöldeke and Samuelson use a different stability notion than we do, but their stability notion coincides with ours under the assumptions they make, see Section 7.

It is possible that weak efficiency and strict efficiency diverge, as in the figure on the right. Importantly, the set of strictly efficient points might not be closed then. To guarantee the existence of stable matchings in the full generality of our model, we need to be content with weak efficiency.

An alternative to the approach taken in this section would be to let $A = \mathbb{R}^2$ and u_W and u_M be given by $u_W(w, m, (r, r')) = r$ and $u_M(w, m, (r, r')) = r'$. In that case, C would just provide the feasible utility allocations. This is the approach used traditionally in cooperative game theory with non-transferable utility, starting with the work of Aumann and Peleg (1960). It is also the approach used in Galichon, Kominers, and Weber (2016) to provide a general framework for the empirical analysis of matching problems with finitely many types. For most purposes, the two approaches are equivalent. The assumption commonly made in the literature on cooperative games that the set of feasible utility allocations is comprehensive from below ensures that one can recover this set from the utility possibility frontier.

5 Matchings and Stability

We are now ready to solve our model. Given our distributional point of view, a matching simply specifies how likely it is to observe certain couple types engaging in certain activities and is, therefore, a distribution of couple-activity types. Every woman is either single or married to a man and both couples and singles are engaged in activities. So the number of women whose type lies in the set B is exactly the number of couples engaging in some activity in which the type of the woman lies in the set B plus the number of single women engaged in some activity whose type lies within the set B. Every man is either single or married to a woman and both couples and singles are engaged in activities. So the number of men whose type lies in the set B is exactly the number of couples engaging in some activity in which the type of the man lies in the set B plus the number of single men engaged in some activity whose type lies within the set B. Also, couples can only engage in activities actually available, so the probability of observing a couple engaging in an unavailable activity is going to be zero. These considerations lead us naturally to our formal definition of a matching. Let $G_C \subseteq W_\emptyset \times M_\emptyset \times A$ be the graph of the activity correspondence C. A *matching* is a Borel measure $\mu \in \mathcal{M}(W_\emptyset \times M_\emptyset \times A)$ such that

(i) $\nu_W(B) = \mu(B \times M_\emptyset \times A)$ for every Borel set $B \subseteq W$,

(ii) $\nu_M(B) = \mu(W_\emptyset \times B \times A)$ for every Borel set $B \subseteq M$,

(iii) $\mu(W_\emptyset \times M_\emptyset \times A \setminus G_C) = 0$, and

(iv) $\mu\Big(\{(w, m, a) \mid w = m = \emptyset\}\Big) = 0.$

Condition (iv) is simply a convenient normalization. We want to emphasize that nothing in the definition of a matching is "random." We think of a matching (after normalizing the measure) as the empirical distribution of couple-activity types induced by a deterministic matching of agents. Section 9 supplies a formal foundation for this point of view.

Our next and crucial task is to define stability. Our definition cannot simply be based on the presence of blocking pairs of agents, since a single blocking pair is going to have measure zero in a continuum population and will therefore not be reflected in a distributional matching. We already saw this in Example 3. Our notion of stability should agree with the usual notion of stability in finite matching models and make only use of distributional information. We do this in two steps. We first define a set of pairs of couple-activity types that witness to the instability of a matching. That is, they signify the presence of an inefficient activity choice between a couple, that one agent would rather be single, or that two agents who are single or in a different relationship would rather be with each other than stay in their current arrangement.

In a finite matching model, a matching will be unstable if and only if such pairs of couple-activity types can be found with positive mass. Equivalently, the matching will be unstable if and only if such pairs of couple-activity types can be found with positive probability by selecting two couple-activity types at random. The latter approach generalizes and we define a stable matching in the general model so that the probability of finding two couple-activity types witnessing to instability by random sampling is zero. Informally, the institution of marriage is stable if it survives the temptations brought upon by random couples dinners. Formally, define the *instability set* I by

$$
\begin{aligned}
I = \Big\{ &((w, m, a), (w', m', a')) \in W_\emptyset \times M_\emptyset \times A \ \times \ W_\emptyset \times M_\emptyset \times A \mid \\
&(m, a'') \succ_w (m, a) \text{ and } (w, a'') \succ_m (w, a) \text{ for some } a'' \in C(w, m), \text{ or} \\
&(m', a'') \succ_{w'} (m', a') \text{ and } (w', a'') \succ_{m'} (w', a') \text{ for some } a'' \in C(w', m'), \text{ or} \\
&(\emptyset, a'') \succ_w (m, a) \text{ for some } a'' \in C(w, \emptyset), \text{ or} \\
&(\emptyset, a'') \succ_{w'} (m', a') \text{ for some } a'' \in C(w', \emptyset), \text{ or} \\
&(\emptyset, a'') \succ_m (w, a) \text{ for some } a'' \in C(m, \emptyset), \text{ or} \\
&(\emptyset, a'') \succ_{m'} (w', a') \text{ for some } a'' \in C(m', \emptyset), \text{ or} \\
&(m', a'') \succ_w (m, a) \text{ and } (w, a'') \succ_{m'} (w', a') \text{ for some } a'' \in C(w, m'), \text{ or} \\
&(m, a'') \succ_{w'} (m', a') \text{ and } (w', a'') \succ_m (w, a) \text{ for some } a'' \in C(w', m). \Big\}.
\end{aligned}
$$

The matching μ is *stable* if $\mu \otimes \mu(I) = 0$, where $\mu \otimes \mu$ is the product measure. If μ is a probability measure, $\mu \otimes \mu$ is simply the distribution of two independent draws from μ. Because of the essential symmetry of $\mu \otimes \mu$, one could drop four appropriately chosen conditions out of the eight conditions specifying the set I and still get an equivalent stability notion. Note also that the "individual rationality constraints" serve double duty; they also guarantee that single agents choose optimal actions. The instability set is open by the continuity of preferences and the continuity of C (only lower hemicontinuity of C matters here).

An equivalent way to define stability is to say that the matching μ is stable if there are no couple activity types (w, m, a) and (w', m', a') such that $\mu(V) > 0$, $\mu(V') > 0$, and $V \times V' \subseteq I$ for some neighborhoods V and V' of (w, m, a) and (w', m', a'), respectively. The equivalence is a straightforward consequence of the topology of $(W_\varnothing \times M_\varnothing \times A) \times (W_\varnothing \times M_\varnothing \times A)$ having a countable basis of open rectangles and I being open.

In our distributional formulation, it is impossible to guarantee that there are no blocking pairs of agents in the underlying economy even when the matching is stable in the way just defined. However, if agents only become aware of each other through random meetings, the probability of two agents forming a blocking pair finding each other is going to be zero.

Stable matchings always exist. To prove this, we approximate a given matching problem by finite matching problems for which Lemma 3 will guarantee the existence of stable matchings. A compactness argument allows us then to extract a stable matching for the limit matching problem. To make this argument work, we make use of the topology of weak convergence of measures.[13] Recall that the sequence $\langle \mu_n \rangle$ of measures in $\mathcal{M}(X)$ with X Polish converges to the measure $\mu \in \mathcal{M}(X)$ under the topology of weak convergence of measures if

$$\lim_{n \to \infty} \int g \, \mathrm{d}\mu_n = \int g \, \mathrm{d}\mu$$

for every bounded continuous function $g : X \to \mathbb{R}$. Whenever we make topological arguments for spaces of measures, it will be understood that we are using the topology of weak convergence of measures.

Our compactness argument can be split into two distinct parts. We first show that our sequence of finite matching problems must have a subsequence converging to some limit measure in Lemma 1 and then show using Lemma 2 that the limit measure will indeed be a matching.

Lemma 1. *Let $\langle \nu_W^n, \nu_M^n, \mu_n \rangle$ be a sequence in $\mathcal{M}(W) \times \mathcal{M}(M) \times \mathcal{M}(W_\varnothing \times M_\varnothing \times A)$ such that $\langle \nu_W^n \rangle$ converges to $\nu_W \in \mathcal{M}(W)$, $\langle \nu_M^n \rangle$ converges to $\nu_M \in \mathcal{M}(M)$, and μ_n is a matching for population distributions ν_W^n and ν_M^n for each natural number n. Then a subsequence of $\langle \mu_n \rangle$ converges.*

The proof of Lemma 1 is a straightforward application of Prohorov's characterization of relative compactness in the topology of weak convergence of measures. Intuitively, we need to make sure that no mass "escapes to infinity." No mass can escape to infinity unless mass of the sequence of population measures escapes to infinity, which is not possible when population measures converge. The argument does not require W, M, or A to be compact.

Lemma 2. *The following set is closed:*

$$\Big\{ (\nu_W, \nu_M, \mu) \in \mathcal{M}(W) \times \mathcal{M}(M) \times \mathcal{M}(W_\varnothing \times M_\varnothing \times A) \mid \mu \text{ is a matching for the population measures } \nu_W \text{ and } \nu_M \Big\}.$$

The larger part of the proof of Lemma 2 consists in showing that conditions (i) and (ii) in the definition of a matching are preserved under taking limits. The topology of weak convergence of measures does not guarantee convergence of the measure of every Borel

[13] The most important facts related to weak convergence can be found in the Appendix 12.

set, so we show that it suffices that convergence holds for an appropriately chosen class of well-behaved Borel sets for which we actually get convergence.

To get the compactness argument off the ground, we prove that stable matchings exist for discrete matching problems with finitely many agents by approximating the matching problem with one in which there are only finitely many activities.

Lemma 3. *A stable matching exists whenever ν_W and ν_M have finite supports and take on only rational values.*

We are now ready for the proof of our main existence theorem.

Theorem 1. *There is at least one stable matching.*

Proof. Let $\langle \nu_W^n, \nu_W^n \rangle$ be a sequence of pairs of measures on W and M, respectively, such that $\langle \nu_W^n \rangle$ converges to ν_W, $\langle \nu_M^n \rangle$ converges to ν_M and ν_W^n and ν_M^n have finite support and only rational values for all n. This is possible since measures with finite supports are dense in the space of all measures and, clearly, every measure with finite support is the limit of a sequence of measures with the same finite support and rational values.

For each n, we can choose a stable matching μ_n for the finite matching problem given by population distributions ν_W^n and ν_M^n by Lemma 3. By passing to a subsequence and using Lemma 1, we can assume without loss of generality that $\langle \mu_n \rangle$ converges to some measure μ, which is again a matching for the population measures ν_W and ν_M by Lemma 2. The continuity assumption on preferences and the lower hemicontinuity of C guarantee that I is open. Therefore,

$$\mu \otimes \mu(I) \leq \liminf_n \mu_n \otimes \mu_n(I) = 0$$

by the Portmanteau theorem and the fact that taking products preserves weak convergence. □

In the usual setting with finitely many agents, a stable matching continues to be stable if we remove matched couples from the population. Indeed, this can only reduce the blocking possibilities of other agents. The same holds true in our model and we note the following lemma for later reference.

Lemma 4. *Let $\mu, \lambda \in \mathcal{M}(W_\emptyset \times M_\emptyset \times A)$, with μ being a stable matching and $\lambda(B) \leq \mu(B)$ for every Borel set $B \subseteq W_\emptyset \times M_\emptyset \times A$. Then $\lambda \otimes \lambda(I) = 0$.*

6 Relation to Large Finite Matching Markets

We show in this section that stable matchings as defined by us are exactly the limits of stable matchings of finite agent matching problems that approximate our limit matching model.

If a sequence of stable matchings for matching problems with finitely many agents converges, it converges to a stable matching for the limiting population measures. If it does not converge, a subsequence will. Indeed, we showed this much when proving Theorem 1. Hence, *at least one* stable matching for the limiting population measures captures the limiting behavior of a sequence of stable matchings for large, finite populations. Next, we show that *all* stable matchings do so; they are all the limits of sequences of stable matchings for large, finite populations.

We first have to define what a matching problem is. We hold the type spaces, the set of activities, the activity correspondence, and the preferences fixed. So we only vary the population measures and define a *matching problem* to be a pair (ν_W, ν_M) of population measures. The matching problem (ν_W, ν_M) is *finite* if both ν_W and ν_M take on only rational values and have finite support.[14] As already mentioned above, by multiplying both ν_W and ν_M with a multiple of all denominators occurring in nonzero values, one obtains an equivalent matching problem in which finitely many types occur in positive integer quantities. Each of these quantities can be interpreted as the number of agents of this type. Similarly, a matching μ is *finite* if it has finite support and only rational values. Note that a stable matching for a finite matching problem need not be finite.

Theorem 2. *Let μ be a matching for the matching problem (ν_W, ν_M). Then μ is stable if and only if there are sequences $\langle \nu_W^n \rangle$, $\langle \nu_M^n \rangle$, and $\langle \mu_n \rangle$ such that*

(i) the matching problem (ν_W^n, ν_M^n) is finite for each n and μ_n is a finite stable matching for it,

(ii) the sequence $\langle \nu_W^n \rangle$ converges to ν_W, the sequence $\langle \nu_M^n \rangle$ converges to ν_M, and $\langle \mu_n \rangle$ converges to μ.

Proof. That (i) and (ii) imply that μ is stable was implicitly already shown in Lemma 2 and the proof of Theorem 1. For the converse, let μ be a stable matching. Normalize it to a probability measure $\bar{\mu} \in \mathcal{M}(W_\varnothing \times M_\varnothing \times A)$ by letting $\bar{\mu}(B) = \mu(B)/\mu(W_\varnothing \times M_\varnothing \times A)$ for every Borel set $B \subseteq W_\varnothing \times M_\varnothing \times A$. For each sequence $\omega = \langle \omega_n \rangle \in (W_\varnothing \times M_\varnothing \times A)^\infty$, we can form the sequence $\langle \bar{\mu}_n^\omega \rangle$ of sample distributions given by

$$\bar{\mu}_n^\omega(B) = n^{-1} \, \#\{m \leq n \mid \omega_m \in B\}$$

for every Borel set $B \subseteq W_\varnothing \times M_\varnothing \times A$ and each natural number n. For each ω and each natural number n, we have

$$\bar{\mu}_n^\omega \otimes \bar{\mu}_n^\omega(I) = n^{-2}\#\{(l, m) \mid (\omega_l, \omega_m) \in I \text{ and } l, m \leq n\}.$$

Since μ is stable, we have $\bar{\mu} \otimes \bar{\mu}(I) = 0$. It follows that for $\otimes_n \bar{\mu}$-almost all ω and each natural number n, $\bar{\mu}_n^\omega \otimes \bar{\mu}_n^\omega(I) = 0$. Moreover, by Varadarajan's version of the Glivenko–Cantelli theorem, Varadarajan (1958), the sequence $\langle \bar{\mu}_n^\omega \rangle$ converges to $\bar{\mu}$ for $\otimes_n \bar{\mu}$-almost all ω. So we can choose some sequence $\omega \in (W_\varnothing \times M_\varnothing \times A)^\infty$ such that $\bar{\mu}_n^\omega \otimes \bar{\mu}_n^\omega(I) = 0$ for each natural number n and such that ω converges to $\bar{\mu}$. Indeed, $\otimes_n \bar{\mu}$-almost every ω will do. Let$\langle q_n \rangle$ be a sequence of rational numbers converging to $\mu(W_\varnothing \times M_\varnothing \times A)$. For each natural number n, let $\mu_n = q_n \cdot \bar{\mu}_n^\omega$ and define ν_W^n and ν_M^n by

$$\nu_W^n(B) = \mu_n(B \times M_\varnothing \times A)$$

[14]It can be shown that a finite measure that takes on only rational values can take on only finitely many values. In the present context, such a measure must have finite support. Our definition of a finite matching problem is therefore redundant.

for every Borel set $B \subseteq W$ and

$$\nu_M^n(B) = \mu_n(W_\emptyset \times B \times A)$$

for every Borel set $B \subseteq M$, respectively. Clearly, $\langle \mu_n \rangle$ converges to μ and μ_n is a finite stable matching for the matching problem (ν_W^n, ν_M^n) for each natural number n. It remains to show that $\langle \nu_W^n \rangle$ converges to ν_W and $\langle \nu_M^n \rangle$ converges to ν_M. Let $B \subseteq W$ be a ν_W-continuity set. Then $B \times M_\emptyset \times A$ is a μ-continuity set and

$$\lim_n \nu_W^n(B) = \lim_n \mu_n(B \times M_\emptyset \times A) = \mu(B \times M_\emptyset \times A) = \nu_W(B).$$

It follows that $\langle \nu_W^n \rangle$ converges to ν_W. Similarly, $\langle \nu_M^n \rangle$ converges to ν_M. □

Let us take a look at what Theorem 2 does not say. The sequences of population measures $\langle \nu_W^n \rangle$ and $\langle \nu_M^n \rangle$ shown to exist do not just depend on the limiting population measures ν_W and ν_M, they depend on the matching μ itself. What is not true is that for any population measures ν_W and ν_M, we can find sequences of population measures $\langle \nu_W^n \rangle$ and $\langle \nu_M^n \rangle$ converging to ν_W and ν_M, respectively, such that for every stable matching μ for the population measures ν_W and ν_M, there exists a sequence $\langle \mu_n \rangle$ converging to μ such that μ_n is a stable matching for population measures ν_W^n and ν_M^n for each natural number n. Formally, the correspondence that maps each matching problem to its set of stable matchings may not be lower hemicontinuous.[15]

This problem is not just an artifact of our distributional model; the phenomenon is known to occur in finite matching theory. Indeed, Pittel (1989) has shown in a model with randomly drawn preferences, no transfers, and the same number of women and men, that the number of stable matchings grows incredibly fast with the number of agents, but Pittel's result is not robust to small changes of populations. Indeed, Ashlagi, Kanoria, and Leshno (2017) have shown that when the number of women differs from the number of men by even one, the set of stable matchings collapses essentially to a unique stable matching as the number of agents grows. A difference of only a single person in population sizes must vanish in the limit, so a reasonable limit model cannot preserve the distinction.

Our model is different, but we can use ideas inspired by Ashlagi, Kanoria, and Leshno (2017) to construct an example in which a sequence of population measures converges, but a stable matching for the limiting population measures is not the limit of any sequence of stable matchings for the given sequence of population measures.

Example 5. We are in a setting without transfers, so we suppress the space A. The types of women and men are simply points on a circle. Formally, we let $W = M = [0, 1)$, endowed with the metric d given by

$$d(x, y) = \min\{|x - y|, 1 - |x - y|\}.$$

Basically, we wrap the half-open unit interval $[0, 1)$ around a circle of circumference 1, or, equivalently, take the closed interval $[0, 1]$ and glue the end points together.

[15] The proofs of the results of the last section establish as a by-product that this correspondence is upper

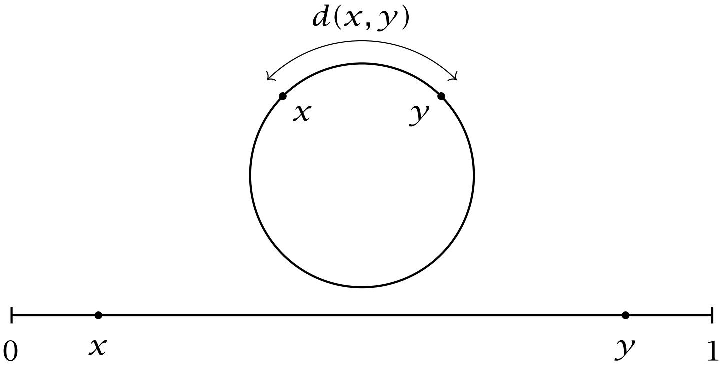

Now we let $u_W : W \times M_ø \to \mathbb{R}$ be given by $u_W(w, m) = d(w, m)$ and $u_W(w, ø) = -1$ for all $w \in W$ and $m \in M$, and we let $u_M : W_ø \times M \to \mathbb{R}$ be given by $u_M(w, m) = -d(w, m)$ and $u_M(ø, m) = -1$ for all $w \in W$ and $m \in M$. So every woman wants a man whose type is as far away as possible from her type on the circle, every man wants a woman whose type is as close as possible to his type on the circle, and everyone is desperate to avoid loneliness.

For any real number r, we let $[r]$ be the largest integer not larger than r. We also let $(r) = r - [r]$. Note that $(r) \in [0, 1)$ for ever real number r. Fix some irrational number θ. We define population measures ν_W^n and ν_M^n for each natural number n by

$$\nu_W^n = \nu_M^n = \frac{1}{n} \sum_{l=1}^{n} \delta_{(\theta l)},$$

where δ_x denotes the probability measure with support $\{x\}$. This means, we are allocating n points on a circle of circumference 1, with clockwise distance (θ) between consecutive points. The irrationality of θ ensures that all these points will be different. Let μ_n be the matching for ν_W^n and ν_M^n that pairs a man of type (θl) with a woman of type (θl). Since every man gets his top choice, μ_n is stable. We let $\nu_W = \nu_M$ be the uniform distribution on $[0, 1)$. It can be shown that $\langle \nu_W^n \rangle$ converges to ν_W and $\langle \nu_M^n \rangle$ converges to ν_M, see Kuipers and Niederreiter (1974, Theorem 1.1 and Example 2.1). Let μ be the uniform distribution on the diagonal $D = \{(x, y) \mid x = y, x, y \in [0, 1)\}$. Then $\langle \mu_n \rangle$ converges to μ, so

$$0 = \lim_{n \to \infty} \int d \, \mathrm{d}\mu_n = \int d \, \mathrm{d}\mu.$$

We now define a new sequence $\langle \nu_M^{n\prime} \rangle$ of population measures for men, adding a single man of type 0 to each ν_M^n. That is,

$$\nu_M^{n\prime} = \frac{1}{n}\delta_0 + \frac{1}{n} \sum_{l=1}^{n} \delta_{(\theta l)}.$$

Clearly, $\langle \nu_M^{n\prime} \rangle$ converges still to ν_M; the presence of one more man is not observable in the limit. But there will be no sequence $\langle \mu_n' \rangle$ such that μ_n' is a stable matching for the population measures μ_W^n and $\mu_M^{n\prime}$ for each natural number n, and such that $\langle \mu_n' \rangle$ converges to μ.

hemicontinuous and compact-valued.

To see this, take any sequence $\langle \mu_n' \rangle$ such that μ_n' is a stable matching for population distributions μ_W^n and $\mu_M^{n\prime}$ for each natural number n. Since ν_W is uniformly distributed on $[0,1)$, we get $\int d(\cdot, m)\, \mathrm{d}\nu_W = 1/4$, the average distance to any point on the circle under the given metric, for each $m \in [0,1)$. Now in each matching μ_n' exactly one man will stay single, and for large n, the average d-distance between the types of the women in the population and the type of the unmatched man will be close to $1/4$. Since μ_n' is stable, no woman prefers the unmatched man to her current partner and almost every woman must be matched with a partner whose type is at least as good. So the average distance between the type of a woman and her partner is close to at least $1/4$. Therefore,

$$1/4 \leq \liminf_n \int d \,\mathrm{d}\mu_n' \neq \int d \,\mathrm{d}\mu = 0.$$

The failure of lower hemicontinuity exhibited in Example 5 shows that one cannot simply transport the structure theory of stable matchings, as first reported by Knuth (1976), from finite matching theory to the distributional model by limit arguments; the limit of a sequence of men-optimal matchings for the sequence of "unbalanced" populations in Example 5 is far from the true men-optimal matching. The situation is better when we allow for transfers. Stable matchings have, under reasonable assumptions, the structure of a lattice in terms of payoff assignments. For transferable utility, this was shown by Gretsky, Ostroy, and Zame (1999). For imperfectly transferable utility, this was shown only recently by Nöldeke and Samuelson (2017). The case of imperfectly transferable utility is considerably harder, under transferable utility one can separate the matching between agents and the assignment of payoffs. In our general model, there are not even functions that assign payoffs to types as we saw in Example 4. Transfers help here as we will see in the next section.

7 Equal Treatment

This section provides sufficient conditions for equal treatment of types in a stable matching and relates our stability notion to the one used in the literature on matching with transfers.

We have seen in Example 4 that two agents with the same type may not be equally well off in a stable matching. In Example 4, no transfers are possible. For each individual agent, a matching can be seen as a decision problem in which they can choose to end up with anyone they want, provided they are willing to provide them with a higher level of satisfaction than their potential partner would obtain without them. With transfers, this level of satisfaction can be interpreted as a price, and transfers allow for competition in prices. In a stable matching, the law of one price should then guarantee that agents with the same type will end up equally well off; there will be equal treatment of equal types. Since we have a topological structure on types, we can look for a stronger form of equal treatment in which agents with similar types will end up similarly well off so that the function from types to utilities is continuous. This section is devoted to a precise formulation of these ideas.

Such equal treatment properties allow us to give alternative characterizations of stable matchings, the formulations of stability used by Nöldeke and Samuelson (2017) and the literature on optimal transport.[16] This shows that our model and stability concept nests these approaches.

We first look at the strong form of equal treatment for agents on one side only, in this case, women. We make the following assumptions:

Imperfectly Transferable Utility: We let $A = [0, 1]$, and for all $w \in W_\emptyset$ and $m \in M_\emptyset$, we let $C(w, m) = [0, 1]$. Preferences are represented by jointly continuous functions $u_W : W \times M_\emptyset \times A \to [0, 1] \to \mathbb{R}$ and $u_M : W_\emptyset \times M \times A \to [0, 1] \to \mathbb{R}$ with u_W increasing in the last coordinate and u_M decreasing in the last coordinate.

Limited Diversity: W is locally compact and M is compact.

We make the outside payoff an agent can obtain explicit by functions $u_W^\emptyset : W \to \mathbb{R}$ and $u_M^\emptyset : M \to \mathbb{R}$ defined by $u_W^\emptyset(w) = u_W(w, \emptyset, 1)$ and $u_M^\emptyset(m) = u_M(\emptyset, m, 0)$, respectively.

Bounds on Transfers do not matter: There are numbers $0 < b < t < 1$ such that

$$u_W(w, m, a) < u_W^\emptyset(w) \text{ whenever } a < b,$$
$$u_M(w, m, a) < u_M^\emptyset(m) \text{ whenever } a > t,$$

and such that

$$u_W(w, m, 1) \geq u_W(w, m', t) \geq u_W(w, m, 0) \text{ for all } w \in W \text{ and } m, m' \in M,$$
$$u_W(w, m, 1) \geq u_W^\emptyset(w) \text{ for all } w \in W \text{ and } m \in M,$$
$$u_M(w, m, 0) \geq u_M(w', m, b) \geq u_M(w, m, 1) \text{ for all } m \in M \text{ and } w, w' \in W,$$
$$u_M(w, m, 0) \geq u_M^\emptyset(m) \text{ for all } m \in M \text{ and } w \in W.$$

The assumption of imperfectly transferable utility was already discussed in Section 4. Clearly, transfers are indispensable in order to obtain a law of one price. The assumption of limited diversity has two aspects. The compactness assumption on M guarantees that there is a bound on the conceivable utility differences obtained between women of similar types; there will always be a type of man that maximizes the difference. We will see in Example 6 that the strong law of one price might fail if M is not compact. The assumption that W is locally compact is fairly weak and will be satisfied if W is a closed or open subspace of a Euclidean space. Still, there are known examples where natural type spaces are not locally compact. In the model of Chiappori and Reny (2016), each agent is endowed with a random variable representing a stochastic income and natural spaces of random variables are not locally compact. Other than that, local compactness imposes a very mild restriction on the diversity of preferences allowed.[17]

[16]Indeed, this section grew out of our attempts to understand how the framework of Nöldeke and Samuelson (2017) relates to ours.

[17]Suppose that we are in a quasilinear environment so that u_W can be interpreted in terms of some numeraire commodity. The function $(w, x) \mapsto u_W(w, \cdot, x)$ from W to the space $C(M_\emptyset)$ of continuous functions on the compact space $M_\emptyset$ endowed with the uniform topology is continuous. It follows that the range of this function is a countable union of compact sets since the local compactness of W implies that W is the countable union of compact sets. However, $C(M_\emptyset)$ is not the countable union of compact sets unless M is finite, so W can only include a proper subset of all conceivable preferences.

However, we need the assumption of limited diversity only to prove the strong form of equal treatment; if we are content with a possibly discontinuous function from types to utilities, the assumption is not needed.

The assumption that bounds on transfers do not matter ensures that every agent could provide every agent on the other side of the market with any utility level that could be achieved without violating anyone's individual rationality constraints. Of course, providing such a level of utility might violate the individual rationality constraint of the agent providing it. In Nöldeke and Samuelson (2017), as in Demange and Gale (1985), every agent can provide every agent on the other side with every conceivable utility level, so corner solutions are ruled out by the absence of corners. The assumption that bounds on transfers do not matter guarantees that the same holds true in our framework in which C being compact-valued forces the existence of corners.

Our proof of the strong equal treatment property establishes the existence of a function reminiscent of a "modulus of continuity" from a stability argument in Lemma 5 and shows that such a modulus of continuity can be used in a probabilistic way to show that a certain measure is supported on the graph of a unique continuous function in Lemma 6.

Lemma 5. *There exists a continuous function $\omega : W \times W \to \mathbb{R}_+$ such that $\omega(w,w) = 0$ for all $w \in W$ and such that*

$$|u_W(w,m,a) - u_W(w',m',a')| \leq \omega(w,w')$$

for all

$$((w,m,a),(w'm',a')) \in W \times M_\emptyset \times [0,1] \ \times \ W \times M_\emptyset \times [0,1]$$

that are not in the instability set I.

The idea behind the proof of Lemma 5 is fairly simple: If w is similar to w', then both can obtain similar payoffs for every conceivable partner: By the compactness of M, there is a conceivable partner who maximizes this difference and we take $\omega(w,w')$ essentially to be the largest such difference. This is the only place where we use compactness of M.

Lemma 6. *Let L be a locally compact Polish space and $\omega : L \times L \to \mathbb{R}$ a continuous function such that $\omega(x,x) = 0$ for all $x \in L$. If μ is a Borel measure on $L \times \mathbb{R}$ such that $|r - r'| \leq \omega(x,x')$ for $\mu \otimes \mu$-almost all pairs $((x,r),(x',r'))$, then μ is supported on a unique continuous function from the support of the L-marginal of μ to $\mathbb{R}$.*

Lemma 6 is somewhat technical, but here is the basic idea: We construct a sequence $\langle x_n, r_n \rangle$ of pairs in $L \times \mathbb{R}$ by random sampling and show, using the function ω, that this sequence of pairs forms, almost surely, the graph of a continuous function whose domain is dense in the support of the L-marginal of μ. We then extend this continuous function to the support of the L-marginal of μ by continuity. This way, we obtain the desired function, at least when L is compact. For noncompact L, one has to glue various such functions obtained for compact subspaces together.

With this machinery in place, we are ready for the main result of this section.

Theorem 3. *Let μ be a stable matching. Then there exists a unique continuous function $V_W : \operatorname{supp} \nu_W \to \mathbb{R}$ such that*

$$V_W(w) = u_W(w, m, a)$$

for μ-almost all $(w, m, a) \in W \times M_\emptyset \times [0, 1]$.

Without the assumption of limited diversity, we can still obtain a weaker form of equal treatment:

Theorem 4. *Let μ be a stable matching. Then there exists a measurable function $V_W : W \to \mathbb{R}$ such that*

$$V_W(w) = u_W(w, m, a)$$

for μ-almost all $(w, m, a) \in W \times M_\emptyset \times [0, 1]$. The assumption of limited diversity is not needed.

Theorem 4 is largely a corollary to Theorem 3. One can approximate the given stable matching by stable matchings for population measures supported on compact sets by a tightness argument. Theorem 3 ensures the existence of a measurable function from types to utilities for the approximate problem. One obtains V_W from these functions for the approximate problems by taking an appropriate limit. In general, one cannot obtain any continuity property for the resulting function. By Lusin's theorem, every Borel measurable function is continuous on a compact set whose complement can be taken to have arbitrarily small (positive) measure.

If both W and M are compact, the payoffs in a stable matching can be represented by two continuous value functions by Theorem 3. The following theorem characterizes stability given that payoffs in a stable matching can be represented by continuous value functions. It is essentially the stability notion used by Nöldeke and Samuelson (2017). Since Nöldeke and Samuelson (2017) assume W and M to be compact, this shows that our notion of stability is equivalent to theirs when W and M are compact.

Theorem 5. *Let μ be a matching. If ν_W and ν_M have full support and $V_W : W \to \mathbb{R}$ and $V_M : M \to \mathbb{R}$ are continuous functions such that*

$$V_W(w) = u_W(w, m, a)$$

for μ-almost all $(w, m, a) \in W \times M_\emptyset \times [0, 1]$ and

$$V_M(m) = u_M(w, m, a)$$

for μ-almost all $(w, m, a) \in W_\emptyset \times M \times [0, 1]$, then μ is a stable matching if and only if the following conditions are satisfied:

(i) $V_W(w) \geq u_W^\emptyset(w)$ for all $w \in W$.

(ii) $V_M(m) \geq u_M^\emptyset(m)$ for all $m \in M$.

(iii) $u_W(w, m, a) \leq V_W(w)$ whenever $u_M(w, m, a) \geq V_M(m)$.

Under transferable utility, the relevant notion of stability is usually given by the dual of the corresponding optimal transport problem. The dual solutions are usually not required to be continuous. The following theorem characterizes stability in terms of value functions that need only be measurable. Such value functions exist always by Theorem 4. In particular, our theory covers the optimal transport approach to stable matchings under transferable utility when the surplus function is continuous.

Theorem 6. *Let μ be a matching. If $V_W : W \to \mathbb{R}$ and $V_M : M \to \mathbb{R}$ are measurable functions such that*

$$V_W(w) = u_W(w, m, a)$$

for μ-almost all $(w, m, a) \in W \times M_\emptyset \times [0, 1]$ and

$$V_M(m) = u_M(w, m, a)$$

for μ-almost all $(w, m, a) \in W_\emptyset \times M \times [0, 1]$, then μ is a stable matching if and only if the following conditions are satisfied for $\nu_W \otimes \nu_M$-almost all $(w, m) \in W \times M$:

(i) $V_W(w) \geq u_W^\emptyset(w)$.

(ii) $V_M(m) \geq u_M^\emptyset(m)$.

(iii) $u_W(w, m, a) \leq V_W(w)$ *if* $u_M(w, m, a) \geq V_M(m)$.

It should be noted that Gretsky, Ostroy, and Zame (1992) prove that when type spaces are compact and the surplus function continuous, then the solution of the dual optimal transport problem can be taken to have continuous values. Their method of proof is different from ours. They construct continuous value functions from measurable value functions by what they call a “shrink-wrap”-argument. Our approach delivers continuous value functions directly. We then obtain the existence of measurable value functions under weaker conditions from the continuous value functions. Nöldeke and Samuelson (2017) obtain continuity directly from their duality theory. Their arguments actually show that any functions V_W and W_M satisfying (i)-(iii) in Theorem 5 must already be continuous when W and M are compact.

That compactness of M is needed to obtain continuity of V_W is shown in the following example.

Example 6. Let $W = [0, 1]$, ν_W be the uniform distribution, $M = \mathbb{R}_+$, and let ν_M be any continuous full support distribution with infinite expectation, such as a truncated Cauchy distribution. We consider a transferable utility setting with a surplus function given by $S(w, m) = w \cdot m$. The value of staying single is zero for everyone; there is no need for anyone to stay alone. A stable matching continues to be stable when restricted to some subpopulation by Lemma 4, so we can use Theorem 5 and Galichon (2016, Theorem 4.7), a result for compact type spaces, to conclude the matching is positively assortative and supported on the graph of the function $T : W \to M$, where T is the quantile function of ν_M. Similarly, it follows using Galichon (2016, Theorem 4.8) that almost every woman w obtains a payoff of at least $\int_0^w T \, \mathrm{d}t$. Since ν_W is uniform and T the quantile function of ν_M, we have $\nu_M = \nu_W \circ T^{-1}$. Since the first moment does not exist for ν_M, there is, therefore, no upper bound to the payoffs that can be obtained by a nonnegligible number of women. If there were a continuous function $V_W : W \to \mathbb{R}$ such that $V_W(w)$ would be the value obtained by w for ν_W-almost all w, there would be a maximum value, since W is compact, and V_W would be bounded.

8 Externalities

Our model can easily be extended to allow for widespread externalities in which both preferences of agents and the activities available to agents might depend on the matching itself as we show in this section. Besides classical externalities and peer effects, this allows for modeling market forms and institutions outside the matching market under consideration. With finitely many agents, indivisibilities may preclude existence of stable matchings as in the following simple example.

Example 7. Let $W = [0,1]$. We now interpret the members of this side of the market as high school graduates who have to decide whether to go to college or join the labor force directly. There is little use in making the other side explicit. The net value for a high school graduate of going to college is αw for some $\alpha > 0$, so $R(w) = \alpha w$ is their reservation wage. There is a uniform wage paid on the labor market that depends on the fraction λ of high school graduates joining the labor force; the decreasing inverse labor demand function is given by P. We consider a finite population of n students placed on the points $1/(2n) + (2k)/(2n)$ for $k = 0, \ldots, n-1$. Assume that $R(1/2) = \alpha/2 = P(1/2)$, so that the market would clear exactly when the mass of students on the labor market is exactly $1/2$. But for n odd, this is impossible. There will be a high school graduate placed exactly at $1/2$, but since they have mass $1/n$ in the population, the mass on the labor market would be strictly below or above the market clearing mass, depending on whether they join the labor market or go to college. The figure below illustrates. The black dots on the inverse labor demand function correspond to wages that are feasible depending on who joins the labor market. Clearly, none of them exactly clears the market.

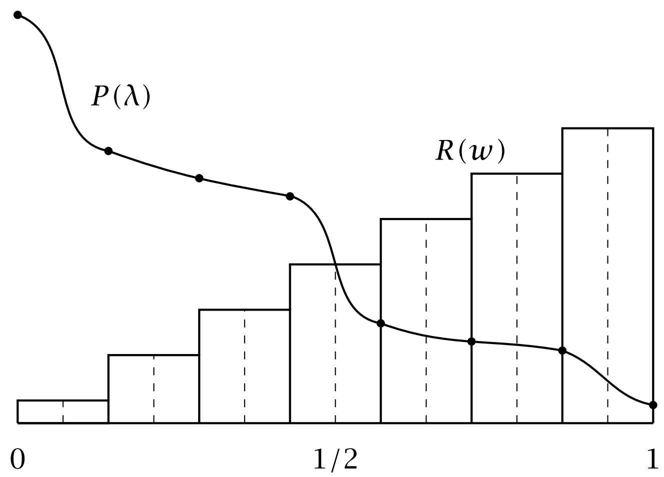

There is no stable matching if agents are indivisible.

But the population distribution converges to the uniform distribution on $[0,1]$ as n increases to infinity,[18] and market clearing is possible in the limit.

[18]This follows from Kuipers and Niederreiter (1974, Theorem 1.1).

Apart from the problems with indivisibilities in the example above, externalities pose some conceptual problems for matching markets with finitely many agents. In that case, each agent has to have some notion of how their behavior impacts others and what kind of response it might cause. Starting with Sasaki and Toda (1996), a number of authors have analyzed such matching markets with finitely many agents and externalities using fairly sophisticated farsightedness ideas. Our approach sidesteps the main problems occurring with finitely many agents and allows for a much simpler treatment. The idea that large aggregate externalities might be compatible with stability and finitely many agents was explored by Fisher and Hafalir (2016), but they still had to make special assumptions to deal with indivisibilities. Closer to our approach is the treatment of widespread externalities by Hammond, Kaneko, and Wooders (1989). The main difference is in how they topologize allocations or matchings. The topology used by them is not compact and their solution concept can be interpreted as an approximate solution concept. Most closely related to our approach is the treatment by Noguchi and Zame (2006) who study the existence of Walrasian equilibria under widespread externalities. This problem is technically more demanding; commodity spaces are never compact and these authors have to employ clever truncations. Our approach to existence is very simple and combines our previous existence result with a simple compactness argument.

We first have to adapt our environment and our assumptions to a setting with externalities. Preferences now include matchings and even more general measures over couple-activity types and C can depend on these measures too.

Acyclicity of Preferences with Externalities: The relation $\succ_w$ on $M_\emptyset \times A \times \mathcal{M}(W_\emptyset \times M_\emptyset \times A)$ is acyclic for each $w \in W$ and the relation $\succ_m$ on $W_\emptyset \times A \times \mathcal{M}(W_\emptyset \times M_\emptyset \times A)$ is acyclic for each $m \in M$.

Continuity of Preferences with Externalities: The set

$$\Big\{(m, a, \mu, m', a', \mu', w) \mid (m, a, \mu) \succ_w (m', a', \mu')\Big\}$$

is open in

$$M_\emptyset \times A \times \mathcal{M}(W_\emptyset \times M_\emptyset \times A) \times M_\emptyset \times A \times \mathcal{M}(W_\emptyset \times M_\emptyset \times A) \times W$$

and the set

$$\Big\{(w, a, \mu, w', a', \mu', m) \mid (w, a, \mu) \succ_m (w', a', \mu')\Big\}$$

is open in

$$W_\emptyset \times A \times \mathcal{M}(W_\emptyset \times W_\emptyset \times A) \times W_\emptyset \times A \times \mathcal{M}(W_\emptyset \times W_\emptyset \times A) \times M.$$

Regularity of the Activity Correspondence with Externalities: The correspondence $C : W_\emptyset \times M_\emptyset \times \mathcal{M}(W_\emptyset \times M_\emptyset) \to 2^A$ is continuous with nonempty and compact values.

The interpretation of these assumptions is similar as in the model without externalities, but the continuity of preferences in externalities is a serious restriction. Agents now have preferences over matchings and even distributions over couple-activity types that are not matchings. Agents cannot individually influence matchings, they have preferences defined on matchings only so we can compare their welfare in different matchings. That they have preferences even over non-matchings is a technical convenience for our existence result, the space of matchings is closed and even compact. For finite approximations introduced below, we need continuity on the larger space though. Fixing the parameter μ in the preferences, we have a standard matching problem that comes with an *instability set* $I(\mu)$. We now say that the matching μ is *stable* if $\mu \otimes \mu(I(\mu)) = 0$.

Theorem 7. *There is at least one stable matching.*

Proof. Let μ_1 be an arbitrary, not necessarily stable, matching. Now recursively construct, using Theorem 1, a sequence $\langle \mu_n \rangle$ of matchings such that $\mu_{n+1} \otimes \mu_{n+1}(I(\mu_n)) = 0$. Using Lemma 1 and Lemma 2, we can assume by passing to a subsequence that $\langle \mu_n \rangle$ converges to a matching μ.

The matching μ is stable. Indeed, for each pair $p = ((w,m,a),(w',m',a'))$ of couple-activity types in $I(\mu)$ we can choose by our strengthened continuity assumption and the lower hemicontinuity of C an open neighborhood O_p of p and an open neighborhood U_p of μ such that $O_p \subseteq I(\mu')$ for $\mu' \in U_p$. Since $\mu_n \in U_p$ for n large enough, we have $O_p \subseteq I(\mu_n)$ for n large enough. Hence, $\mu \otimes \mu(O_p) \leq \liminf_n \mu_n \otimes \mu_n(O_p) = 0$ by the Portmanteau theorem. Now suppose for the sake of contradiction that $\mu \otimes \mu(I(\mu)) > 0$. Since Borel measures are regular, there exists then a compact set $K \subseteq I(\mu)$ such that $\mu \otimes \mu(K) > 0$. Now the family $(O_p)_{p \in K}$ is an open cover of K and K is therefore covered by finitely many open sets of $\mu \otimes \mu$-measure zero, in contradiction to $\mu \otimes \mu(K) > 0$. □

Without externalities, stable matchings correspond exactly to the limits of stable matchings for finite matching problems. Clearly, this is not the case here. But under a mild assumption, each stable matching corresponds to the limit of "nearly stable" finite matchings for finite matching problems. We assume the following.

Equicontinuous utility: Preferences are represented by fixed continuous functions $u_W : W \times M_\emptyset \times A \times \mathcal{M}(W_\emptyset \times M_\emptyset \times A) \to \mathbb{R}$ and $u_M : W_\emptyset \times M \times A \times \mathcal{M}(W_\emptyset \times M_\emptyset \times A) \to \mathbb{R}$ vary equicontinuously in $\mathcal{M}(W_\emptyset \times M_\emptyset \times A)$. That is, for each $\mu \in \mathcal{M}(W_\emptyset \times M_\emptyset \times A)$ and each $\epsilon > 0$, there is a neighborhood U of μ such that

$$\sup_{w,m,a} \left| u_W(w,m,a,\mu) - u_W(w,m,a,\mu') \right| < \epsilon$$

and

$$\sup_{w,m,a} \left| u_M(w,m,a,\mu) - u_M(w,m,a,\mu') \right| < \epsilon$$

for all $\mu' \in U$.

Importantly, cardinal properties of the specific functions u_W and u_M are taken to be meaningful. If W, M, and A are compact and u_W and u_M are continuous, equi-continuity follows from the Arzelà-Ascoli theorem.

Now define for each $\epsilon > 0$ and matching μ the ϵ-*instability set* $I_\epsilon(\mu)$ so that $I_\epsilon(\mu)$ replaces the strict preferences in $I(\mu)$ by strict inequalities in terms of u_W and u_M that have to hold with a gap of at least ϵ. The notion of a matching problem translates directly to the model with externalities. Adapting the proof of Theorem 2 using arguments from the proof of Theorem 7, one can easily prove the following theorem.

Theorem 8. *Assume equicontinuous utility. Let μ be a matching for the matching problem (ν_W, ν_M). Then μ is stable if and only if there are sequences $\langle \nu_W^n \rangle$, $\langle \nu_M^n \rangle$, and $\langle \mu_n \rangle$ such that*

(i) the matching problem (ν_W^n, ν_M^n) is finite for each n and μ_n is a finite matching for it,

(ii) the sequence $\langle \nu_W^n \rangle$ converges to ν_W, the sequence $\langle \nu_M^n \rangle$ converges to ν_M, and $\langle \mu_n \rangle$ converges to μ,

(iii) and for all $\epsilon > 0$, there exists a natural number N such that $\mu_n \otimes \mu_n(I_\epsilon(\mu_n)) = 0$ for each $n \geq N$.

Proof. We first show that (i)-(iii) implies that μ is stable. Now $I(\mu) = \bigcup_{\epsilon \downarrow 0} I_\epsilon(\mu)$, so $\mu \otimes \mu(I(\mu)) > 0$ would imply $\mu \otimes \mu(I_\epsilon(\mu)) > 0$ for some $\epsilon > 0$. But exactly as as in the proof of Theorem 7, one can show $\mu \otimes \mu(I_\epsilon(\mu)) = 0$.

For the other direction, assume that $\mu(I(\mu)) = 0$. As in the proof of Theorem 2, we can show that there are sequences $\langle \nu_W^n \rangle$, $\langle \nu_M^n \rangle$, and $\langle \mu_n \rangle$ such that (i) and (ii) holds and such that $\mu_n \otimes \mu_n(I(\mu)) = 0$ for all n. We show that (iii) holds too. So let $\epsilon > 0$. By equicontinuous utility, there exists a neighborhood U of μ such that

$$\sup_{w,m,a} |u_W(w,m,a,\mu) - u_W(w,m,a,\mu')| < \epsilon/2$$

and

$$\sup_{w,m,a} |u_M(w,m,a,\mu) - u_M(w,m,a,\mu')| < \epsilon/2$$

for all $\mu' \in U$. Note that $\mu_n \otimes \mu_n(I(\mu)) = 0$ is equivalent to $I(\mu) \cap \operatorname{supp} \mu_n \times \operatorname{supp} \mu_n = \emptyset$ for each n since μ_n is finite. Similarly, $\mu_n \otimes \mu_n(I_\epsilon(\mu_n)) = 0$ is equivalent to $I_\epsilon(\mu_n) \cap \operatorname{supp} \mu_n \times \operatorname{supp} \mu_n = \emptyset$. Let N be such that $\mu_n \in U$ for $n \geq N$. It is straightforward but slightly tedious to verify that $I_\epsilon(\mu_n) \cap \operatorname{supp} \mu_n \times \operatorname{supp} \mu_n = \emptyset$ for $n \geq N$. □

9 Individualistic Representation

We talked about agents, but our model has none. In this section, we show that one can enrich the model to take account of every single agent. There are measure spaces of women and men, and each matching matches a unique woman to a unique man or lets her be by herself. This exercise has two purposes: First, it shows there is nothing random about a matching in our distributional model; the underlying matching of agents is deterministic. Second, it clarifies our notion of stability by taking it to the level of agents. Nevertheless, our distributional model is much easier to handle for practical purposes.

Examples 1 and 2 show that not every space of agents will work. The problem in these examples is that types cannot be arbitrarily split up. To deal with this problem, the spaces of agents we construct have to satisfy a very strong form of nonatomicity that no Borel probability measure on a Polish space satisfies. On the level of agents, a matching is a function and we have to find such a function that induces the matching, in measure form, on the level of couple-activity types. Formally, this is a so-called "purification"-problem for measure-valued maps. Our purification theorem is related to, but does not follow from existing results on the purification of measure-valued maps such as Podczeck (2009), Loeb and Sun (2009), Wang and Zhang (2012), and Greinecker and Podczeck (2015). The additional complication we face comes from requiring the matching to be represented by a measurable, measure-preserving isomorphism. This poses additional demand on the spaces of agents, we must have many agents of every type.[19] A related, somewhat weaker, such representation is given in Gretsky, Ostroy, and Zame (1992, Section 1.5.1), but the individualistic matchings obtained there need not be invertible and measurable in both directions.

In the individualistic representation of a matching, we require all couples and individuals to choose feasible activities. Not even a measure zero set of exceptions is allowed. In the individualistic representation of a stable matching, we further require all couples and individuals to choose efficient feasible activities. Again, not even a measure zero set of exceptions is allowed.

We say that $a \in C(w,m)$ is *efficient* if there is no $a' \in C(w,m)$ such that both $(m,a') \succ_w (m,a)$ and $(w,a') \succ_m (w,a)$, or if $w \in W$ and $m = \emptyset$ and there is no $a' \in C(w,m)$ such that $(\emptyset,a') \succ_w (\emptyset,a)$, or if $w = \emptyset$ and $m \in M$ and there is no $a' \in C(w,m)$ such that $(\emptyset,a') \succ_m (\emptyset,a)$, or if $w = \emptyset = m$. We are now able to state our representation theorem on the level of agents.

Theorem 9. *There exist measure spaces $(A_W, \mathcal{A}_W, \tau_W)$ and $(A_M, \mathcal{A}_M, \tau_M)$, and measurable type functions $t_W : A_W \to W_\emptyset$ and $t_M : A_M \to M_\emptyset$ such that $\nu_W(B) = \tau_W \circ t_W^{-1}(B)$ for every Borel set $B \subseteq W$, such that $\nu_M(B) = \tau_M \circ t_M^{-1}(B)$ for every Borel set $B \subseteq M$, and such that for every matching μ, there is a pair of measurable functions $\phi : A_W \to A_M$ and $\alpha : W_\emptyset \to A$ such that*

(i) the measurable function ϕ is a bijection with a measurable inverse that preserves measure; $\tau_M(S) = \tau_W \circ \phi^{-1}$ for every measurable set $S \subseteq A_M$,

(ii) for every Borel set $B \subseteq W \times M \times A$,

$$\mu(B) = \tau_W\Big(\big\{a_W \in A_W \mid (t_W(a_W), t_M(\phi(a_W)), \alpha(a_W)) \in B\big\}\Big),$$

(iii) and for every $a_W \in A_W$,

$$\alpha(a_W) \in C\big(t_W(a_W), t_M(\phi(a_W))\big).$$

Moreover, if μ is stable, then ϕ and α can be chosen to satisfy the following condition:

(iv) For every $a_W \in A_W$, the activity choice $\alpha(a_W)$ is efficient.

[19] However, saturation or, equivalently, super-nonatomicity as in the purification results mentioned above does not suffice. Every nonatomic Borel probability measure on a Polish space extends to a saturated (superatomless) probability measure on a larger σ-algebra by the main result of the appendix of Podczeck (2009). In particular, Examples 1 and 2 are compatible with the spaces of agents being saturated probability spaces.

In our representation for stable matchings, there might still be some "blocking pairs." What the representation does ensure is that two agents that could form a blocking pair have a hard time finding each other. It is tedious but straightforward to verify that the set of $a_W \in A_W$ and $a_M \in A_M$ that could form a blocking pair has $\tau_W \otimes \tau_M$-measure zero. Part of the tedium is that one has to define blocking pairs for pairs of agents and take account of the outside options ø. We sketch the argument for the simplest case, the case of the marriage model without activities, non-binding individual rationality constraints, and both ν_W and ν_M being probability measures. Let μ be a stable matching. Then $\phi : A_W \to A_M$ is a measure preserving measurable bijection with a measurable inverse such that the function $a_W \mapsto (t_W(a_W), t_M(\phi(a_W)))$ has τ_W-distribution μ. That ϕ is measure preserving implies that the function $a_M \mapsto (t_W(\phi^{-1}(a_M)), t_M(a_M))$ has τ_M-distribution μ too. It follows that the function $(a_W, a_M) \mapsto (t_W(a_w), t_M(\phi(a_W)), t_W(\phi^{-1}(a_M)), t_M(a_M))$ has $\tau_W \otimes \tau_M$-distribution $\mu \otimes \mu$. So for $\tau_W \otimes \tau_M$-almost all a_W and a_M,

$$(t_W(a_w), t_M(\phi(a_W)), t_W(\phi^{-1}(a_M)), t_M(a_M)) \notin I$$

since $\mu \otimes \mu(I) = 0$. And for such a_W and a_M, it is not the case that a_W prefers a_M to $\phi(a_W)$ and a_M prefers a_W to $\phi^{-1}(a_M)$.

10 Concluding Remarks

We have provided foundations for large two-sided matching markets from the distributional point of view. Our model represents exactly the distributional properties of large finite matching markets that are preserved under weak convergence. Even though individual agents are negligible, stability has a simple and natural interpretation in the limit model, and we nest existing models with transfers. We also provided an individualist interpretation of our distributional model and used it to clarify the economic meaning of our stability notion.

There is a problem in treating our model as a limit model of econometric matching models. In econometric models of matching, the payoff usually includes an idiosyncratic additive component that is stochastically independent between pairs of agents that might be matched. This implies in our distributional framework that there is a jointly measurable function $\epsilon : W \times M \to \mathbb{R}$, representing the purely idiosyncratic part, such that for ν_W-almost all w and ν_W almost all $w' \in W$, the random variables $\epsilon(w, \cdot)$ and $\epsilon(w', \cdot)$ are stochastically independent. By Sun (2006, Proposition 2.1), the random variable $\epsilon(w, \cdot)$ must be constant for ν_W-almost all w and therefore deterministic. One approach to deal with the problem is to characterize the projection of stable matchings onto observable characteristics, not to worry whether unobservable characteristics converge or not, and only look at the limit of the observable part. In a special econometric version of the marriage model, Menzel (2015) did exactly that. The details of how the idiosyncratic part of preferences is modeled will generally matter. A technical hurdle in adapting our approach to this more general problem is that we make much use of the instability set being open. When only some characteristics are observable, we have to work with the projection of an open set and such a projection need not be open. These problems are of course not particular to our approach, they haunt all the existing models nested by ours.

We consider the definition of stability to be one of the main contributions of this paper. This definition allows us to work in a distributional model with a rich set of types even though all effective coalitions are finite and seemingly invisible at the level of analysis. There is no reason to think this approach is restricted to matching theory alone. In club theory as developed in Cole and Prescott (1997) and Ellickson, Grodal, Scotchmer, and Zame (1999), agents in individualistic continuum economies form clubs with finitely many members to organize their consumption and production decisions. A major restriction in this literature is that the external characteristics of agents (those relevant to other agents) and characteristics of clubs must belong to a finite set, a restriction that could be overcome by the distributional approach when club internal decisions are magnified by random sampling as in our stability definition.

11 Omitted Proofs

11.1 Proofs omitted from Section 5

Proof of Lemma 1. It suffices to show that the sequence $\langle \mu^n \rangle$ is tight. Take some $\epsilon > 0$. Since the families $\{\nu_W^n\}$ and $\{\nu_M^n\}$ of population distributions come from converging sequences, they have compact closure and are therefore tight. So there are compact sets $K_W \subseteq W$ and $K_M \subseteq M$ such that both $\nu_W^n(W \setminus K_W) < \epsilon/2$ and $\nu_M^n(M \setminus K_M) < \epsilon/2$ for every natural number n. Let $K_W^ø = K_W \cup \{ø\}$ and $K_M^ø = K_M \cup \{ø\}$. Now

$$C(K_W^ø \times K_M^ø) = \bigcup_{(w,m)\in K_W^ø \times K_M^ø} C(w,m)$$

is compact as the forward image of a compact set under a compact-valued upper hemicontinuous correspondence, Aliprantis and Border (2006, 17.8). We show that

$$\mu_n\Big(W_ø \times M_ø \times A \ \setminus \ K_W^ø \times K_M^ø \times C(K_W^ø \times K_M^ø)\Big) < \epsilon$$

for every natural number n. Indeed, in order for a couple-activity type to be in $W_ø \times M_ø \times A$ but not in $K_W^ø \times K_M^ø \times C(K_W^ø \times K_M^ø)$ it has to either have a first term not in $K_W^ø$ and therefore be in $W_ø \backslash K_W^ø \times M_ø \times A$, or it has to have a second term not in $K_M^ø$ and therefore be in $W_ø \times (M_ø \backslash K_M^ø) \times A$, or it has to have both first two terms in $K_W^ø \times K_M^ø$ but the third term not in $C(K_W^ø \times K_B^ø)$ and therefore be in $K_W^ø \times K_M^ø \times A \backslash C(K_W^ø \times K_B^ø)$. Now $W_ø \backslash K_W^ø \times M_ø \times A$ has μ_n-measure $\nu_W^n(W \backslash K_W) < \epsilon/2$ by condition (i) in the definition of a matching, $W_ø \times (M_ø \backslash K_M^ø) \times A$ has μ_n-measure $\nu_M^n(M \backslash K_M) < \epsilon/2$ by condition (ii) in the definition of a matching, and, finally, $K_W^ø \times K_M^ø \times A \backslash C(K_W^ø \times K_B^ø)$ has μ_n measure zero by condition (iii) in the definition of a matching. Now $\epsilon/2 + \epsilon/2 + 0 = \epsilon$, so we obtain the desired inequality and therefore the tightness of the sequence $\langle \mu_n \rangle$. □

Proof of Lemma 2. Since the topology of weak convergence of measures is metrizable, it suffices to prove that the limit (ν_W, ν_M, μ) of a convergent sequence $\langle \nu_W^n, \nu_M^n, \mu_n \rangle$ with values in the set must again lie in the set.

We first show that for every Borel set $B \subseteq W$, we have $\mu\,(B \times M_ø \times A) = \nu_W\,(B)$, which is equivalent to showing $\mu \circ \pi_W^{-1}(B) = \nu_W(B)$, with $\pi_W : W \times M_\emptyset \times A \to W$ being the canonical projection. Since Borel measures on Polish spaces are tight, for each Borel set $B \subseteq W$,

$$\mu \circ \pi_W^{-1}(B) = \sup\left\{\mu \circ \pi_W^{-1}(K) \mid K \text{ compact and } K \subseteq B\right\}.$$

and

$$\nu_W(B) = \sup\{\nu_W(K) \mid K \text{ compact and } K \subseteq B\}.$$

It therefore suffices to prove the result for B compact. So let B be compact. For each $\delta > 0$, let $B_\delta = \{x \in G \mid d(x,B) < \delta\}$. Note that $B_\delta \downarrow B$ as $\delta \downarrow 0$. For each $x \in B$ and $\epsilon > 0$, the boundary $\partial B_\epsilon(x)$ of the ball $B_\epsilon(x)$ is a subset of the sphere $S_\epsilon(x) = \{y \in W \mid d(x,y) = \epsilon\}$. This boundary can therefore only have positive $\mu \circ \pi_W^{-1}$-measure or ν_W-measure for countably many ϵ, since no finite measure space can allow for an uncountable disjoint family of measurable set with positive measure.[20] For each $x \in B$, there is therefore some ϵ_x^1 such that $0 < \epsilon_x^1 < 1$ and

$$\mu \circ \pi_W^{-1}\left(\partial B_{\epsilon_x^1}(x)\right) = \nu_W\left(\partial B_{\epsilon_x^1}(x)\right) = 0.$$

The family $\{B_{\epsilon_x^1} \mid x \in B\}$ is an open cover of the compact set B and has, therefore, a finite subcover. Let B^1 be the union of this subcover. Then $B \subseteq B^1 \subseteq B_1$, B^1 is open and since the boundary of a finite union of sets is a subset of their boundaries, we have

$$\mu \circ \pi_W^{-1}\left(\partial B^1\right) = \nu_W\left(\partial B^1\right) = 0.$$

Given that B^n is defined, we can repeat the procedure to obtain B^{n+1} so that $B^{n+1} \subseteq B^n$ and $B^{n+1} \subseteq B_{1/n}$ and

$$\mu \circ \pi_W^{-1}(\partial B^n) = \nu_W(\partial B^n) = 0.$$

By the Portmanteau theorem,

$$\mu \circ \pi_W^{-1}(B^n) = \lim_m \mu_m \circ \pi_W^{-1}(B^n) = \lim_m \nu_W^m(B^n) = \nu_W(B^n)$$

for all n. Now, since measures are downward-continuous,

$$\mu \circ \pi_W^{-1}(B) = \lim_n \mu \circ \pi_W^{-1}(B^n) = \lim_n \nu_W(B^n) = \nu_W(B).$$

Similarly, one can show that $\mu(W_\emptyset \times B \times A) = \nu_M(B)$ for every Borel set $B \subseteq M$.

Finally, since G_C is closed as the graph of a compact-valued and upper hemicontinuous correspondence, $W_\emptyset \times M_\emptyset \times A \setminus G_C$ is open and therefore, by another use of the Portmanteau theorem,

$$\mu(W_\emptyset \times M_\emptyset \times A \backslash G_C) \leq \liminf_n \mu_n(W_\emptyset \times M_\emptyset \times A \backslash G_C) = 0.$$

□

Proof of Lemma 3. We can assume without loss of generality that $W = \operatorname{supp} \nu_W$ and $M = \operatorname{supp} \nu_M$. For each $w \in W_\emptyset$ and $m \in M_\emptyset$ let $\langle a_{wm}^n \rangle$ be a sequence in $C(w,m)$ such

[20]Indeed, if $\mathcal{F}$ is a disjoint family of measurable sets, the family of all $F \in \mathcal{F}$ whose measure exceeds $1/n$ must be finite for every natural number n. Since the countable union of finite sets is countable, the conclusion follows.

that $\{a^n_{wm} \mid n \in \mathbb{N}\}$ is dense in $C(w,m)$. Define $C_n : W_\emptyset \times M_\emptyset \to 2^A$ by $C_n(w,m) = \{a^k_{wm} : k \leq n\}$. A stable matching exists when we replace C by C_n. To see this, represent the matching problem by an individualistic matching problem with actual agents. This is possible since ν_W and ν_M have finite support and take on only rational values. Extend all preferences to strict linear orders. This is possible since preferences are acyclic. With strict linear orders, weakly efficient and strictly efficient activity choices coincide. So one can apply the extended deferred acceptance algorithm with wages of Crawford and Knoer (1981) and Kelso and Crawford (1982) with the set of efficient activities to obtain a stable matching for the extended preferences. Since extending preferences cannot reduce blocking possibilities, the matching continues to be stable under the original preferences. The induced distribution of couple-activity types gives us a distributional stable matching.

Note that every matching, stable or not, for the restricted correspondence C_n is also a, not necessarily stable, matching for the unrestricted correspondence C. So we can find a sequence $\langle \mu_n \rangle$ of matchings such that μ_n is a stable matching for the restricted correspondence C_n. By passing to a subsequence and using Lemma 1, we can assume without loss of generality that $\langle \mu_n \rangle$ converges to some measure μ, which is again a matching by Lemma 2.

It remains to prove that μ is stable. For each $n \in \mathbb{N}$, let I_n be the instability set for the matching problem with the restricted correspondence C_n. The continuity assumption on preferences guarantees that I_n is open for every natural number n, and together with $\bigcup_n C_n(w,m) = \{a^n_{wm} \mid n \in \mathbb{N}\}$ being dense in $C(w,m)$ also that $I = \bigcup_n I_n$. If $k \leq n$, then $I_k \subseteq I_n$, so $\mu_n \otimes \mu_n(I_k) = 0$ for $k \leq n$. Therefore

$$\mu \otimes \mu(I_k) \leq \liminf_n \mu_n \otimes \mu_n(I_k) = 0$$

by the Portmanteau theorem. Finally,

$$\mu \otimes \mu(I) = \mu \otimes \mu\Big(\bigcup_k I_k\Big) \leq \sum_k \mu \otimes \mu(I_k) = 0.$$

□

Proof of Lemma 4. Clearly, λ is absolutely continuous with respect to μ and has a Radon-Nikodym derivative g with values in $[0,1]$. Using Fubini's theorem,

$$\begin{aligned}
\lambda \otimes \lambda(I) &= \int 1_I \,\mathrm{d}\lambda \otimes \lambda \\
&= \int \int 1_I(x,y)\,\mathrm{d}\lambda(x)\mathrm{d}\lambda(y) \\
&= \int g(y) \int g(x) 1_I(x,y)\,\mathrm{d}\mu(x)\mathrm{d}\mu(y) \\
&= \int g(x)g(y)1_I(x,y)\,\mathrm{d}\mu \otimes \mu(x,y) \\
&\leq \int 1_I \,\mathrm{d}\mu \otimes \mu \\
&= \mu \otimes \mu(I) \\
&= 0.
\end{aligned}$$

□

11.2 Proofs omitted from Section 7

We first need a preliminary lemma:

Lemma 7. *There exists a unique function* $c : W \times (M \times W \times [b,t]) \to [0,1]$ *such that*

$$u_M(w,m,a) = u_M\big(w', m, c(w',(w,m,a))\big).$$

The function c *is continuous.*

We call the unique function $c : W \times (M \times W \times [b,t]) \to A$ shown to exist in Lemma 7 the *compensation function*. For notational ease, we write $c_{w'}(w,m,a)$ for $c(w',(w,m,a))$.

Proof of Lemma 7. For each $w, w' \in W$, $m \in M$ and $a \in [b,t]$,

$$u_M(w',m,0) \geq u_M(w,m,a) \geq u_M(w',m,1)$$

by the assumptions that bounds on transfers do not matter. By the intermediate value theorem, there exists $a' \in [0,1]$ such that $u_M(w,m,a) = u_M(w',m,a')$. Since u_M is decreasing in its third argument, there can be at most one $a' \in [0,1]$ such that $u_M(w,m,a) = u_M(w',m,a')$. Therefore, c is well defined and unique. To see that c is continuous, note that

$$c(w',(w,m,a)) = \arg\min_{a' \in [0,1]} |u_M(w,m,a) - u_M(w',m,a')|$$

and apply the maximum theorem. □

Proof of Lemma 5. Let $((w,m,a),(w',m',a')) \notin I$. We derive a number of inequalities by simple stability and continuity arguments. If both $m \in M$ and $m' \in M$, then

$$u_W(w,m,a) \geq u_W(w,m',c_w(w',m',a')),$$
$$u_W(w',m',a') \geq u_W(w',m,c_{w'}(w,m,a)),$$

which implies

$$|u_W(w,m,a) - u_W(w',m',a')| \leq |u_W(w,m,a) - u_W(w',m,c_{w'}(w,m,a))| + |u_W(w',m',a') - u_W(w,m',c_w(w',m',a'))|.$$

If $m = \emptyset = m'$, then

$$|u_W(w,m,a) - u_W(w',m',a')| = |u_W^{\emptyset}(w) - u_W^{\emptyset}(w')|.$$

If $m \in M$ and $m' = \emptyset$, then

$$u_W(w,m,a) - u_W(w',m',a') \leq u_W(w,m,a) - u_W(w',m,c_{w'}(w,m,a)),$$
$$u_W(w',m',a') - u_W(w,m,a) = u_W^{\emptyset}(w') - u_W(w,m,a) \leq u_W^{\emptyset}(w') - u_W^{\emptyset}(w).$$

Similarly, if $m = \varnothing$ and $m' \in M$, then

$$u_W(w', m', a') - u_W(w, m, a) \le u_W(w', m', a') - u_W(w, m', c_{w'}(w, m, a)),$$
$$u_W(w, m, a) - u_W(w', m', a') = u_W^{\varnothing}(w) - u_W(w', m', a') \le u_W^{\varnothing}(w) - u_W^{\varnothing}(w').$$

Collecting inequalities, we obtain

$$\begin{aligned}|u_W(w, m, a) - u_W(w', m', a')| \le & |u_W(w, m, a) - u_W(w', m, c_{w'}(w, m, a))| \\ & + |u_W(w', m', a') - u_W(w, m', c_w(w', m', a'))| \\ & + |u_W^{\varnothing}(w) - u_W^{\varnothing}(w')|,\end{aligned}$$

with the first two terms only being effective if $m \in M$ or $m' \in M$, respectively. A fortiori, $|u_W(w, m, a) - u_W(w', m', a')|$ can be no larger than

$$\begin{aligned}&\max_{m \in M, a \in [b,t]} |u_W(w, m, a) - u_W(w', m, c_{w'}(w, m, a))| \\ &+ \max_{m' \in M, a' \in [b,t]} |u_W(w', m', a') - u_W(w, m', c_w(w', m', a'))| \\ &+ |u_W^{\varnothing}(w) - u_W^{\varnothing}(w')|.\end{aligned}$$

This last expression depends only on w and w' and we take it to be the value of $\omega(w, w')$. Clearly, $\omega(w, w') = 0$ if $w = w'$. The continuity of ω follows from the maximum theorem. □

Proof of Lemma 6. Without loss of generality, we can assume that μ is a probability measure. Consider the space $(L \times \mathbb{R})^\infty$ endowed with the product measure $\mu^\infty = \otimes_n \mu$ and let $\langle x_n, r_n \rangle \in (L \times \mathbb{R})^\infty$ be a random sequence.

The space L has a countable basis; pick an open set O in such a basis. If $\mu(O \times \mathbb{R}) = 0$, then μ^∞-almost surely $x_n \notin O$ for each natural number n. If $\mu(O \times \mathbb{R}) > 0$, then μ^∞-almost surely $x_n \in O$ for some natural number n. So μ^∞-almost surely, the closure of the random set $\{x_n \mid n \in \mathbb{N}\}$ equals the support of the L-marginal of μ.

Assume first that L is compact. Now μ^∞-almost surely, $|r_m - r_n| \le \omega(x_m, x_n)$ for $m, n \in \mathbb{N}$. Indeed, this holds, by assumption, for fixed m and n, and there are only countably many such pairs of natural numbers. In particular, $r_m = r_n$ whenever $x_m = x_n$ holds μ^∞-almost surely, so the random set $\{(x_n, r_n) \mid n \in \mathbb{N}\}$ is μ^∞-almost surely the graph of a function g^∞. Let d be any metric that metrizes L. We show that g^∞ is uniformly continuous with respect to d. Let $\epsilon > 0$. The set $\omega^{-1}([0, \epsilon))$ is an open neighborhood of the diagonal $D_L = \{(x, y) \in L \times L \mid x = y\}$. Define the metric d_1 on $L \times L$ by $d_1((x, y), (x', y')) = d(x, x') + d(y, y')$ and observe that $d_1((x, y), D_L)) = d(x, y)$. Since D_L is compact and the function $(x, x) \mapsto d_1((x, x), \omega^{-1}([\epsilon, \infty)))$ continuous, the function must take on a minimal value $\delta > 0$. Then for $d(x_m, x_n) < \delta$, we get $\omega(x_m, x_n) < \epsilon$ and, since $|r_m - r_n| \le \omega(x_m, x_n)$, also $|r_m - r_n| < \epsilon$. So g^∞ is uniformly continuous and extends, by Aliprantis and Border (2006, 3.11), to a unique continuous function g defined on the closure of $\{x_n \mid n \in \mathbb{N}\}$, which equals the support of the L-marginal of μ. Now for μ-almost all (x, r), we must have $|r - r_n| \le \omega(x, x_n)$ for each natural number n. But this implies that (x, r) lies on the graph of g, since g is continuous.

Next, to see that g is unique, assume that g' is another continuous function from the support of the K-marginal of μ to $\mathbb{R}$ whose graph supports μ. Take another random sequence $\langle x'_n, r'_n \rangle \in (K \times \mathbb{R})^\infty$. Now μ^∞-almost surely, the closure of the set $\{x'_n \mid n \in \mathbb{N}\}$ equals the support of the K-marginal of μ as above. Since g and g' coincide μ^∞-almost surely, we have μ^∞-almost surely that $g(x'_n) = r'_n = g'(x'_n)$ for each natural number n. But two continuous functions that agree on a dense set must coincide, so $g' = g$.

Finally, we dispose of the assumption that L is compact and assume only that L is a locally compact Polish space. Without loss of generality, we can take μ to have support L. Indeed, every closed subspace of a locally compact Hausdorff space is easily shown to be locally compact. Now by Aliprantis and Border (2006, 2.76 and 2.77), there exist an increasing sequence $\langle K_n \rangle$ of compact sets such that $\bigcup_n K_n = L$ and such that K_n is a subset of the interior of K_{n+1} for each natural number n. Let the Borel measure μ_n be defined by $\mu_n(B) = \mu(B \cap K_n)$ for each natural number n and each Borel set $B \subseteq L$. By what we have shown above, there exists a continuous function $g_n : K_n \to \mathbb{R}$ such that μ_n is supported on the graph g_n and any two continuous functions with this property must agree on the support of μ_n. Now for every point $x \in L$, there is some natural number n such that $x \in K_n$ and we let $n(x)$ be the smallest natural number with this property. We define $g : L \to \mathbb{R}$ by $g(x) = g_{n(x)+1}(x)$.

Next, we show that $g(x) = g_l(x)$ for each $l \geq n(x) + 1$. Indeed, x is in the interior of K_l for $l \geq n(x) + 1$. By the full support assumption, the interior of K_l is a subset of the support of μ_l. The support of $\mu_{n(x)+1}$ is a subset of the support of μ_l for each $l \geq n(x)+1$. Now g_l restricted to the support of $\mu_{n(x)+1}$ is a continuous function such that $\mu_{n(x)+1}$ is supported on its graph. But then this restriction must coincide with $g_{n(x)+1}$ on the support of $\mu_{n(x)+1}$. It follows that $g(x) = g_l(x)$ for each $l \geq n(x) + 1$.

To see that g is continuous, take any $x \in L$. By assumption, $x \in K_{n(x)}$ and $K_{n(x)}$ is a subset of the interior of $K_{n(x)+1}$. So there is an open neighborhood U of x that is wholly included in the interior of $K_{n(x)+1}$. This implies $n(y) \leq n(x) + 1$ and therefore $g(y) = g_{n(x)+2}(y)$ for all $y \in U$. So g is continuous at x because $g_{n(x)+2}$ is.

We are almost done; two details are left. First, note that $\mu(B) = \lim_n \mu_n(B)$ for every Borel set $B \subseteq L$, so the measure μ is supported on the graph of g. Second, note that g is the only continuous functions whose graph supports μ since the uniqueness argument given above does not rely on L being compact. □

Proof of Theorem 3. Let $\omega : W \times W \to \mathbb{R}_+$ be a function as guaranteed to exist by Lemma 5. Let μ^W be the trace of μ on $W \times M_\emptyset \times [0,1]$. That is, $\mu^W(B) = \mu(B \cap W \times M_\emptyset \times [0,1])$ for every Borel set $B \subseteq W_\emptyset \times M_\emptyset \times [0,1]$. Define $h : W \times M_\emptyset \times [0,1] \to W_\emptyset \times M_\emptyset \times [0,1] \times \mathbb{R}$ by

$$h(w,m,a) = (w,m,a,u_W(w,m,a)).$$

We show that the $W \times \mathbb{R}$-marginal of $\mu^W \circ h^{-1}$ satisfies the conditions of Lemma 6. To see this, let $\pi : W \times M_\emptyset \times [0,1] \times \mathbb{R} \to W \times \mathbb{R}$ be the canonical projection. The $W \times \mathbb{R}$-marginal of $\mu^W \circ h^{-1}$ is then simply $\mu^W \circ h^{-1} \circ \pi^{-1}$. Now

$$\begin{aligned}
&\mu^W \circ h^{-1} \circ \pi^{-1} \otimes \mu^W \circ h^{-1} \circ \pi^{-1}\Big(\{(w,r),(w',r') \mid |r-r'| > \omega(w,w')\}\Big) \\
&= \mu^W \circ h^{-1} \otimes \mu^W \circ h^{-1}\Big(\{(w,m,a,r),(w',m',a',r') \mid |r-r'| > \omega(w,w')\}\Big) \\
&\leq \mu^W \circ h^{-1} \otimes \mu^W \circ h^{-1}\Big(\{(w,m,a,r),(w',m',a',r') \mid (w,m,a),(w',m',a') \in I\}\Big) \\
&= \mu^W \otimes \mu^W(I) \leq \mu \otimes \mu(I) = 0.
\end{aligned}$$

Let $V_W : S_W \to \mathbb{R}$ be the unique function shown to exist by Lemma 6. We have

$$\begin{aligned}
0 &= \mu^W \circ h^{-1} \circ \pi^{-1}\Big(\{(w,r) \mid V_W(w) \neq r\}\Big) \\
&= \mu^W\Big(\{(w,m,a) \mid V_W(w) \neq u_W(w,m,a)\}\Big),
\end{aligned}$$

so V_W has the desired properties. Moreover, since any other function V'_W with the desired properties must satisfy the last two equations in place of V_W, uniqueness follows from the uniqueness part of Lemma 6. □

Proof of Theorem 4. Since ν_W and ν_M are tight, there exist increasing sequences $\langle K_W^n \rangle$ and $\langle K_M^n \rangle$ of compact subsets of W and M, respectively, such that $\nu_W(W) = \lim_n \nu_W(K_W^n)$ and $\nu_M(M) = \lim_n \nu_M(K_M^n)$. Let $K_{W\emptyset}^n = K_W^n \cup \{\emptyset\}$ and $K_{M\emptyset}^n = K_M^n \cup \{\emptyset\}$. For each natural number n, define μ_n by

$$\mu_n(B) = \mu(B \cap K_{W\emptyset}^n \times K_{M\emptyset}^n \times [0,1])$$

for every Borel set $B \subseteq W_\emptyset \times M_\emptyset \times [0,1] \times W_\emptyset \times M_\emptyset \times [0,1]$. Then μ_n is a stable matching for appropriately chosen population measures supported on compact sets by Lemma 4. By Theorem 3, there exists for each natural number n a measurable function $V_n : W \to \mathbb{R}$ such that $V_n(w) = u_W(w,m,a)$ for μ_n-almost all $(w,m,a) \in K_W^n \times K_{M\emptyset}^n \times [0,1]$. Let $V : W \to \mathbb{R} \cup \{\infty\}$ be given by $V(w) = \limsup_n V_n(w)$. Construct V_W from V by changing the value ∞ to some real number. We claim that V_W has the desired property. Consider the set

$$N = \{(w,m,a) \in W \times M_\emptyset \times [0,1] \mid V_W(w) \neq u_W(w,m,a)\}.$$

It suffices to show that $\mu(N) = 0$. Suppose not. Since $\mu(N) = \lim_n \mu_n(N)$, there exists some natural number k such that $\mu_k(N) > 0$. Let $n \geq k$. We claim that $V_k(w) = V_n(w)$ for μ_k-almost all $(w,m,a) \in K_W^k \times K_{M\emptyset}^k \times [0,1]$. Indeed, every set of μ_n-measure zero has μ_k-measure zero, so $V_n(w) = u_W(w,m,a) = V_k(w)$ for μ_k-almost all $(w,m,a) \in K_W^k \times K_{M\emptyset}^k \times [0,1]$. It follows that $V_W(w) = \lim_n V_n(w) = V_k(w)$ for μ_k-almost all $(w,m,a) \in K_W^k \times K_{M\emptyset}^k \times [0,1]$. Therefore $\mu_k(N) > 0$ is equivalent to

$$\mu_k\Big(\{(w,m,a) \in W \times M_\emptyset \times [0,1] \mid V_k(w) \neq u_W(w,m,a)\}\Big) > 0,$$

which is impossible. □

Proof of Theorem 5. We first show that (i)-(iii) are satisfied if μ is stable. Let

$$N = \{(w,m,a) \in W \times M_\emptyset \times [0,1] \mid u_W(w,m,a) < u_W^\emptyset\}.$$

Now, $N \times N \subseteq I$. Since $\mu \otimes \mu(I) = 0$, we get $\mu \otimes \mu(N \times N) = \mu(N)\mu(N) = 0$ and therefore $\mu(N) = 0$. Together with $V_W(w) = u_W(w, m, a)$ for μ-almost all $(w, m, a) \in W \times M_\varnothing \times [0, 1]$, this implies

$$\mu\Big(\{(w, m, a) \in W \times M_\varnothing \times [0, 1] \mid V_W(w) < u_W^\varnothing\}\Big) = 0.$$

Since μ is a matching, the open set $\{w \in W \mid V_W(w) < u_W^\varnothing\}$ has therefore ν_W-measure zero. But since ν_W has full support, every open set with ν_W-measure zero must be empty. This proves (i) and an analogous argument applies to (ii).

Next, we deal with (iii). Suppose that $u_M(w, m, a) \geq V_M(m)$, but $u_W(w, m, a) > V_W(w)$. We know from (i) and (ii) and the assumption that bounds on transfers don't matter that we can assume $a \neq 0$. So there is some a^* slightly smaller than a such that $V_M(m) < u_M(w, m, a^*)$ and $V_W(w) < u_W(w, m, a^*)$ by continuity. Also by continuity, there exists open neighborhoods O_w of w and O_m of m, such that $V_M(m') < u_M(w', m', a^*)$ and $V_W(w') < u_W(w', m', a^*)$ for all $w' \in O_w$ and $m' \in O_m$. Now $O_w \times M_\varnothing \times [0, 1] \ \times \ W_\varnothing \times O_m \times [0, 1]$ is a subset of

$$\Big\{((w, m, a), (w', m', a')) \in W \times M_\varnothing \times [0, 1] \ \times \ W_\varnothing \times M \times [0, 1] \mid \\ u_W(w, m', a^*) > V_W(w) \text{ and } u_M(w, m', a^*) > V_M(m')\Big\}$$

and the latter set coincides $\mu \otimes \mu$-almost surely with

$$\Big\{((w, m, a), (w', m', a')) \in W \times M_\varnothing \times [0, 1] \ \times \ W_\varnothing \times M \times [0, 1] \mid \\ u_W(w, m', a^*) > u_W(w, m, a) \text{ and } u_M(w, m', a^*) > u_M(w', m', a')\Big\},$$

a subset of the instability set I. It follows that

$$\mu \otimes \mu(O_w \times M_\varnothing \times [0, 1] \ \times \ W_\varnothing \times O_m \times [0, 1]) = 0.$$

Since this is the measure of a measurable rectangle and μ is a matching, this shows that

$$0 = \mu(O_w \times M_\varnothing \times [0, 1])\mu(W_\varnothing \times O_m \times [0, 1]) = \nu_W(O_w)\nu_M(O_m),$$

so $\nu_W(O) = 0$ or $\nu_M(O_m) = 0$. If $\nu_W(O_w) = 0$, then O_w is empty since ν_W has full support. If $\nu_M(O_m) = 0$, then O_m is empty since ν_M has full support. In either case, we obtain a contradiction.

For the other direction, assume that (i)-(iii) hold. Proving that $\mu \otimes \mu(I) = 0$ is somewhat tedious since I is defined by no less than eight conditions. Each of these conditions defines an open subset of $W_\varnothing \times M_\varnothing \times A \ \times \ W_\varnothing \times M_\varnothing \times A$ and I is the union of these eight open sets. It suffices, therefore, to show separately that each of these eight open sets has $\mu \otimes \mu$-measure zero. We do one case here and leave the others to the industrious reader.[21]

[21] The other condition involving blocking pairs is completely analogous to the one we verify here, showing that the four individual rationality conditions hold is straighforward, and the two conditions concerning efficient activity choices for both couples hold vacuously since there can be no inefficient activity choices under the assumption of imperfectly transferable utility.

So let

$$I' = \Big\{((w,m,a),(w',m',a')) \in W_\emptyset \times M_\emptyset \times A \ \times \ W_\emptyset \times M_\emptyset \times A \mid \\ (m',a'') \succ_w (m,a) \text{ and } (w,a'') \succ_{m'} (w',a') \text{ for some } a''\}.$$

We show that $\mu \otimes \mu(I') = 0$. We can rewrite I' as

$$\Big\{((w,m,a),(w',m',a')) \in W_\emptyset \times M_\emptyset \times A \ \times \ W_\emptyset \times M_\emptyset \times A \mid \\ u_W(w,m',a'') > u_W(w,m,a) \text{ and } u_M(w,m',a'') > u_M(w',m',a') \text{ for some } a''\Big\},$$

which $\mu \otimes \mu$-almost surely coincides with

$$\Big\{((w,m,a),(w',m',a')) \in W_\emptyset \times M_\emptyset \times A \ \times \ W_\emptyset \times M_\emptyset \times A \mid \\ u_W(w,m',a'') > V_W(w) \text{ and } u_M(w,m',a'') > V_M(m') \text{ for some } a''\Big\}.$$

This last set must be empty by (iii) and therefore have $\mu \otimes \mu$-measure zero. □

Proof of Theorem 6. Showing that (i)-(iii) hold almost surely if μ is stable, follows almost exactly as in the proof of Theorem 5. But whenever we showed that some set violating the condition is an open set of measure zero and therefore empty under the full support assumption, it now suffices that the set is measurable with measure zero.

Showing that μ is stable if conditions (i)-(iii) holds, works exactly as in the proof of Theorem 5, with the tiny modification that the set discussed at the end may not be empty, but is already assumed to have measure zero. Neither the continuity of the value function nor the support being full played any other role in proving that direction. □

11.3 Proof of Theorem 9

We need some definitions for the proof of Theorem 9. A probability space (Ω, Σ, τ) is *saturated* if for every two Polish spaces X and Y, every Borel probability measure μ on $X \times Y$ and every measurable function $f : \Omega \to X$ with distribution equal to the X-marginal of μ, there exists a measurable function $g : \Omega \to Y$ such that the function $(f,g) : \Omega \to X \times Y$ given by $(f,g)(\omega) = \big(f(\omega), g(\omega)\big)$ has distribution μ, that is, $\mu = \tau \circ (f,g)^{-1}$.

An *automorphism* of the probability space (Ω, Σ, τ) is a measurable bijection $h : \Omega \to \Omega$ with measurable inverse such that $\tau(A) = \tau\big(h(A)\big)$ for all $A \in \Sigma$. A probability space (Ω, Σ, τ) is *homogeneous* if for every two measurable functions $f : \Omega \to X$ and $g : \Omega \to X$ with X Polish such that $\tau \circ f^{-1} = \tau \circ g^{-1}$, there exists an automorphism h such that $f(\omega) = g(h(\omega))$ for almost all ω.

An extensive discussion of these concepts can be found in Fajardo and Keisler (2002), where it is also shown that probability spaces that are both saturated and homogeneous exist.[22]

[22]The notion of homogeneity used in Fajardo and Keisler (2002) is more permissive in that they require only automorphisms of sets of measure 1 that may be smaller than the whole probability space. But in their proof of their Theorem 3B.12, which shows that homogeneous and saturated probability spaces exist, they obtain the automorphisms as the realization of automorphisms of the underlying measure algebra using a result from Maharam (1958), which actually delivers automorphisms in our stronger sense.

Proof of Theorem 9. We first ignore (iii) and (iv) and then patch up our solution so that even these conditions hold. Extend ν_W to all of $W_\emptyset$ by assigning mass $\nu_M(M)$ to the point $\emptyset \in W_\emptyset$, and extend ν_M to all of $M_\emptyset$ by assigning mass $\nu_W(W)$ to the point $\emptyset \in M_\emptyset$. The measures ν_W and ν_M thus extended satisfy $\nu_W(W_\emptyset) = \nu_W(W) + \nu_M(M) = \nu_M(M_\emptyset)$ and we take them without loss of generality to be probability measures. We take $(A_W, \mathcal{A}_W, \tau_W)$ and $(A_M, \mathcal{A}_M, \tau_M)$ to be the same saturated and homogeneous, but otherwise arbitrary, probability space (Ω, Σ, τ).

Let X be any Polish space and $g : \Omega \to X$ be any measurable function. By saturation, there exists $h : \Omega \to W_\emptyset$ such that $\tau \circ (g, h)^{-1} \circ \tau = \tau \circ g^{-1} \otimes \nu_W$. In particular, $\tau \circ h^{-1} = \nu_W$ and we can take t_W to be h. Similarly, we can find a function $t_M : \Omega \to M_\emptyset$ such that $\tau \circ t_M^{-1} = \nu_M$.

Now let μ be a matching and let $\mu_\emptyset$ be a measure on $W_\emptyset \times M_\emptyset \times A$ obtained from μ by letting $\mu_\emptyset(B) = \mu(B)$ for every Borel set $B \subseteq W_\emptyset \times M_\emptyset \times A \setminus \{\emptyset, \emptyset\} \times A$, but

$$\mu_\emptyset(W_\emptyset \times M_\emptyset \times A) = 1$$

and such that $\mu_\emptyset(G_C) = 1$. So $\mu_\emptyset(W_\emptyset \times M_\emptyset \times A) = 1$, the $W_\emptyset$-marginal of $\mu_\emptyset$ is ν_W, and the $M_\emptyset$-marginal of $\mu_\emptyset$ is ν_M. By saturation, there exists measurable functions $f_W : \Omega \to M_\emptyset$, $\alpha_W : \Omega \to A$, $f_M : \Omega \to W_\emptyset$, and $\alpha_M : \Omega \to A$ such that

$$\tau \circ (t_W, f_W, \alpha_W)^{-1} = \mu_\emptyset = \tau \circ (f_M, t_M, \alpha_M)^{-1}.$$

By homogeneity, there exists an automorphism $\phi : \Omega \to \Omega$ such that

$$\Big(t_W(\omega), f_W(\omega), \alpha_W(\omega)\Big) = \Big(f_M(\phi(\omega)), t_M(\phi(\omega)), \alpha_M(\phi(\omega))\Big)$$

for τ-almost all $\omega \in \Omega$. In particular,

$$\tau \circ \Big(t_W, t_M(\phi(\omega)), \alpha_W\Big)^{-1} = \mu_\emptyset.$$

There might still be some $\omega \in \Omega$ such that $\alpha_W(\omega) \notin C\big(t_W(\omega), t_M(\phi(\omega))\big)$. The correspondence $C : W_\emptyset \times M_\emptyset \to 2^A$ is upper hemicontinuous with nonempty and compact values and therefore also measurable with nonempty and closed values. By the Kuratowski–Ryll-Nardzewski measurable selection theorem, Aliprantis and Border (2006, 18.13), there exists a measurable function $s : W_\emptyset \times M_\emptyset \to A$ such that $s(w, m) \in C(w, m)$ for all $w \in W_\emptyset$ and $w \in W_\emptyset$. Let

$$\begin{aligned} N &= \Big\{\omega \in \Omega \;\Big|\; \alpha_W(\omega) \notin C\big(t_W(\omega), t_M(\phi(\omega))\big)\Big\} \\ &= \Big\{\omega \in \Omega \;\Big|\; \big(t_W(\omega), t_M(\phi(\omega)), \alpha_W(\omega)\big) \notin G_C\Big\}. \end{aligned}$$

Since $\mu_\emptyset(G_C) = 1$, we have $\tau(N) = 0$. Define $\alpha : \Omega \to A$ by

$$\alpha(\omega) = \begin{cases} s\Big(t_W(\omega), t_M(\phi(\omega))\Big) & \text{if } \omega \in N, \\ \alpha_W(\omega) & \text{otherwise.} \end{cases}$$

The functions ϕ and α have the desired properties apart from, possibly, (iv).

Now assume that μ is in addition stable. Let $E : W_\emptyset \times M_\emptyset \to 2^A$ be the correspondence such that $E(w, m)$ consists of all efficient $a \in C(w, m)$. Note that efficient activity choices are maximal elements under the weak Pareto ordering for couples and this ordering is acyclic with an open graph. It follows from a version of the maximum theorem that E is upper hemicontinuous with nonempty and compact values, Hildenbrand (1974, Theorem 3 on page 29).[23] In particular, E is a measurable correspondence with nonempty closed values. By the Kuratowski–Ryll-Nardzewski measurable selection theorem, there exists a measurable function $s' : W_\emptyset \times M_\emptyset \to A$ such that $s'(w, m) \in E(w, m)$ for all $w \in W_\emptyset$ and $m \in M_\emptyset$. Now let

$$N' = \Big\{\omega \in \Omega \;\Big|\; \alpha_W(\omega) \notin E\big(t_W(\omega), t_M(\phi(\omega))\big)\Big\}.$$

We show that $\tau(N') = 0$. Since μ is stable, we have $\mu \otimes \mu(I) = 0$ and therefore also $\mu_\emptyset \otimes \mu_\emptyset(I) = 0$. Define $I' \subseteq I$ by

$$\begin{aligned} I' = \Big\{&\big((w, m, a), (w', m', a')\big) \in W_\emptyset \times M_\emptyset \times A \;\times\; W_\emptyset \times M_\emptyset \times A \;\big|\; \\ &a \notin E(w, m) \text{ or } a' \notin E(w', m')\Big\}. \end{aligned}$$

Note that

$$\mu_\emptyset\Big(\{(w, m, a) : a \notin E(w, m)\}\Big) \leq \mu_\emptyset \otimes \mu_\emptyset(I') \leq \mu_\emptyset \otimes \mu_\emptyset(I) = 0.$$

Since $\mu_\emptyset = \tau \circ \Big(t_W, t_M(\phi(\omega)), \alpha_W\Big)^{-1}$, we must have $\tau(N') = 0$. In the present case, define $\alpha : \Omega \to A$ by

$$\alpha(\omega) = \begin{cases} s'\Big(t_W(\omega), t_M(\phi(\omega))\Big) & \text{if } \omega \in N', \\ \alpha_W(\omega) & \text{otherwise.} \end{cases}$$

The functions ϕ and α have the desired properties. □

12 Mathematical Appendix

Here we collect some mathematical background information used throughout the paper without much ado. The reader is assumed to be familiar with basic notions of general topology and a bit of measure and integration theory. The material on weak convergence of measures can be found in Parthasarathy (1967) and Billingsley (1999), with the caveat that these books only deal with probability measures. Non-probability measures are dealt with in Bogachev (2007, Chapter II.8), but that book is considerably less accessible. However, there is a mechanical way to identify a family of uniformly bounded measures with a family of probability measures that allows one to transfer results on probability measures to the more general case. Let $\mathcal{F}$ be a family of measures on a measurable space $(X, \mathcal{X})$ such that for some $b > 0$, $\mu(X) < b$ for all $\mu \in \mathcal{F}$. Define a new measurable space $(X^*, \mathcal{X}^*)$ such that for some $* \notin X$, $X^* = X \cup \{*\}$ and $\mathcal{X}^* = \mathcal{X} \cup \{A \cup \{*\} \mid A \in \mathcal{X}\}$. For each $\mu \in \mathcal{F}$, let μ^* be the probability measure on $(X^*, \mathcal{X}^*)$ such that $\mu^*(A) = \mu(A)/b$ for

[23] The cited result assumes that preferences are transitive and irreflexive, not merely acyclic. But the proof works without modification for acyclic preferences using the fact that maximal elements exist for acyclic relations on nonempty finite sets.

$A \in \mathcal{X}$ and $\mu^*(\{*\}) = 1 - \mu(X)/b$. The function $\mu \mapsto \mu^*$ identifies measures in $\mathcal{F}$ with probability measures. If X has a Polish topology, to be defined below, there is a unique Polish topology on X^* such that X is a subspace and $*$ an isolated point. A continuous real-valued function on X can then be identified with a continuous real-valued function on X^* that vanishes on $*$. With these tools at hand, the reader should be able to obtain the general results from the special case of probability measures.

A topological space is *metrizable* if there exists a metric that induces the topology. A topological space is *completely metrizable* if there exists a complete metric that induces the topology. A subset of a topological space is *dense* if it intersects every nonempty open set or, equivalently, its closure is the whole space. A topological space is *separable* if there is some countable dense subset. A metrizable topological space is separable if and only if it has a countable basis, that is if there is a countable family of open sets such that every open set is a union of open sets in this family. A topological space is *Polish* if it is separable and completely metrizable. The distinction between Polish spaces and separable complete metric spaces is not just nitpicking. A metric subspace S of a separable complete metric space is a separable complete metric space if and only if S is closed. But a topological subspace S of a Polish space is Polish if and only if S is the countable intersection of open sets (which includes closed sets). The countable topological product of Polish spaces is again Polish. We usually view products of topological spaces as being endowed with the product topology without further comment. A topological space is *locally compact* if every point is in the interior of a compact set. Examples of Polish spaces that fail to be locally compact are infinite dimensional separable Banach spaces.

We endow each Polish space X with the *Borel σ-algebra*, the smallest σ-algebra that includes all open sets. Measurable sets in this σ-algebra are *Borel sets*. We only consider measures with real values (∞ is not allowed as the value of a measure). A measure defined on the Borel σ-algebra is a *Borel measure*. It is a *Borel probability measure* if X has measure 1. A Borel measure μ on a Polish space is always *regular*, that is, for each Borel set $B \subseteq X$,

$$\mu(B) = \sup\{\mu(K) \mid K \text{ is compact and } K \subseteq B\} = \inf\{\mu(O) \mid O \text{ is open and } O \supseteq B\}.$$

If X is a Polish space, we let $\mathcal{M}(X)$ be the corresponding set of Borel measures and $\mathcal{P}(X)$ be the corresponding space of Borel probability measures. We endow $\mathcal{M}(X)$ with the *topology of weak convergence*. This is the weakest topology such that for every bounded continuous function $g : X \to \mathbb{R}$, the function $\mu \mapsto \int g \, d\mu$ is continuous. Endowed with the topology of weak convergence, $\mathcal{M}(X)$ is again a Polish space and $\mathcal{P}(X)$ a closed subspace. Convergence of sequences of measures will always be understood to be with respect to this topology. Write ∂B for the *boundary* of B, that is the set of closure points of B that are not interior points. If μ is a Borel measure, the set B is a *μ-continuity set* if $\mu(\partial B) = 0$. Note that X itself has an empty boundary and is therefore always a μ-continuity set. The so-called *Portmanteau theorem* states that the following are equivalent for a sequence $\langle \mu_n \rangle$ in $\mathcal{M}(X)$ and a measure $\mu \in \mathcal{M}(X)$.

(i) The sequence $\langle \mu_n \rangle$ converges to μ,

(ii) $\limsup_n \mu_n(F) \geq \mu(F)$ for every closed set $F \subseteq X$ and $\lim_n \mu_n(X) = \mu(X)$,

(iii) $\liminf_n \mu(O) \leq \mu(O)$ for every open set $O \subseteq X$ and $\lim_n \mu_n(X) = \mu(X)$,

(iv) $\lim_n \mu_n(B) = \mu(B)$ for every μ-continuity set $B \subseteq X$.

We say that a family $\mathcal{F} \subseteq \mathcal{M}(X)$ of Borel measures is *tight* if for each $\epsilon > 0$ there is a compact set $K_\epsilon \subseteq X$ such that $\mu(X \setminus K_\epsilon) < \epsilon$ for all $\mu \in \mathcal{F}$. Similarly, we say that a sequence $\langle \mu_n \rangle$ of elements of $\mathcal{M}(X)$ is *tight* if the family $\{\mu_n \mid n \in \mathbb{N}\}$ is. *Prohorov's theorem* states that $\mathcal{F} \subseteq \mathcal{M}(X)$ is relatively compact (has compact closure) if and only if $\mathcal{F}$ is tight and $\sup_{\mu \in \mathcal{F}} \mu(X) < \infty$. The *support* $\operatorname{supp} \mu$ of a Borel measure $\mu \in \mathcal{M}(X)$ is the largest closed set whose complement has μ-measure zero. In particular, μ has *full support* if $\operatorname{supp} \mu = X$; this is equivalent to every open set of μ-measure zero being empty. The family of all Borel measures with finite support is dense in $\mathcal{M}(X)$.

If $(X_i, \mathcal{X}_i)_{i \in I}$ is a family of measurable spaces, the *product σ-algebra* $\otimes_i \mathcal{X}_i$ on $\prod_i X_i$ is the smallest σ-algebra that makes the coordinate projections measurable. Alternatively, it is the smallest σ-algebra on $\prod_i X_i$ that includes every *measurable rectangle*, where a measurable rectangle is a set of the form $\prod_i A_i$ with $A_i \in \mathcal{X}_i$ for all i and $A_i = X_i$ for all but finitely many i. For the countable topological product of Polish spaces, the Borel σ-algebra of the topological product coincides with the product σ-algebra of the individual Borel σ-algebras. If we look at only two measurable spaces $(X, \mathcal{X})$ and $(Y, \mathcal{Y})$, we write $\mathcal{X} \otimes \mathcal{Y}$ for the product σ-algebra. If $(X, \mathcal{X}, \nu)$ and $(Y, \mathcal{Y}, \mu)$ are measure spaces, there is a unique measure $\nu \otimes \mu$ defined on the product σ-algebra, the *product measure*, such that $\nu \otimes \mu(A \times B) = \nu(A)\mu(B)$ for each measurable rectangle $A \times B$. We heavily rely on the fact that for two Polish spaces X and Y and sequences $\langle \mu_n \rangle$ in $\mathcal{M}(X)$ and $\langle \nu_n \rangle$ in $\mathcal{M}(Y)$, $\langle \mu_n \rangle$ converges to $\mu \in \mathcal{M}(X)$ and $\langle \nu_n \rangle$ converges to $\nu \in \mathcal{M}(Y)$ if and only if the sequence $\langle \mu_n \otimes \nu_n \rangle$ converges to $\mu \otimes \nu \in \mathcal{M}(X \times Y)$ (see Billingsley (1999, Theorem 2.8)).

For probability measures, product measures can be defined even with infinitely many factors. If $(X_i, \mathcal{X}_i, \mu_i)_{i \in I}$ is a family of probability spaces, there is a unique probability measure $\otimes_i \mu_i$, the *independent product* or, again, *product measure*, defined on $\otimes_i \mathcal{X}_i$ such that $\mu(\prod_i A_i) = \prod_{i: A_i \neq X_i} \mu_i(A_i)$ for every measurable rectangle $\prod_i A_i$. If X is a Polish space and $\omega = \langle \omega_n \rangle$ a sequence in X and n a natural number, we let $\mu_n^\omega \in \mathcal{P}(X)$ be the *n-th sample distribution* given by

$$\mu_n^\omega = n^{-1} \# \{m \leq n \mid \omega_n \in B\}$$

for each Borel set $B \subseteq X$ ($\#A$ is the cardinality of A). The Varadarajan (1958) version of the Glivenko–Cantelli theorem says that for each $\mu \in \mathcal{P}(X)$, the random sequence $\langle \mu_n^\omega \rangle$ converges to μ for $\otimes_n \mu$-almost all $\omega \in \prod_n X$.

References

Charalambos D. Aliprantis and Kim C. Border. *Infinite dimensional analysis*. Springer, Berlin, third edition, 2006. ISBN 978-3-540-32696-0; 3-540-32696-0.

Itai Ashlagi, Yash Kanoria, and Jacob D. Leshno. Unbalanced random matching markets: The stark effect of competition. *Journal of Political Economy*, 125:69–98, 2017.

Robert J. Aumann. Markets with a continuum of traders. *Econometrica*, 32:39–50, 1964.

Robert J. Aumann and Bezalel Peleg. Von Neumann–Morgenstern solutions to cooperative games without side payments. *Bulletin of the American Mathematical Society*, 66:173–179, 1960.

Eduardo M. Azevedo and John William Hatfield. Existence of equilibrium in large matching markets with complementarities. 2015.

Eduardo M. Azevedo and Jacob D. Leshno. A supply and demand framework for two-sided matching markets. *Journal of Political Economy*, 124:1235–1268, 2016.

Mourad Baïou and Michel Balinski. Erratum: The stable allocation (or ordinal transportation) problem. *Mathematics of Operations Research*, 27:662–680, 2002.

Patrick Billingsley. *Convergence of probability measures*. Wiley Series in Probability and Statistics. John Wiley & Sons, Inc., New York, second edition, 1999. ISBN 0-471-19745-9.

Vladimir I. Bogachev. *Measure theory. Vol. I, II*. Springer-Verlag, Berlin, 2007. ISBN 978-3-540-34513-8; 3-540-34513-2.

Yeon-Koo Che, Jinwoo Kim, and Fuhito Kojima. Stable matching in large economies. 2015.

Pierre-André Chiappori. *Matching with Transfers: The Economics of Love and Marriage*. Princeton University Press, 2017. ISBN 978-0691171739.

Pierre-André Chiappori and Elisabeth Gugl. Transferable utility and demand functions: A complete characterization, 2015.

Pierre-André Chiappori and Philip J. Reny. Matching to share risk. *Theoretical Economics*, 11: 227–251, 2016.

Pierre-André Chiappori and Bernard Salanié. The econometrics of matching models. *Journal of Economic Literature*, 54:832–861, 2016.

Pierre-André Chiappori, Robert J. McCann, and Lars P. Nesheim. Hedonic price equilibria, stable matching, and optimal transport: equivalence, topology, and uniqueness. *Economic Theory*, 42: 317–354, 2010.

Harold L. Cole and Edward C. Prescott. Valuation equilibrium with clubs. *Journal of Economic Theory*, 74:19–39, 1997.

Harold L. Cole, George J. Mailath, and Andrew Postlewaite. Efficient non-contractible investments in large economies. *J. Econom. Theory*, 101:333–373, 2001.

Vincent P. Crawford and Elsie Marie Knoer. Job matching with heterogeneous firms and workers. *Econometrica*, 49:437–450, 1981.

Gabrielle Demange and David Gale. The strategy structure of two-sided matching markets. *Econometrica*, 53:873–888, 1985.

Federico Echenique, Sangmok Lee, and Matthew Shum. Aggregate matchings. *mimeo*, 2010.

Federico Echenique, Sangmok Lee, Matthew Shum, and M. Bumin Yenmez. The revealed preference theory of stable and extremal stable matchings. *Econometrica*, 81:153–171, 2013.

Bryan Ellickson, Birgit Grodal, Suzanne Scotchmer, and William R. Zame. Clubs and the market. *Econometrica*, 67:1185–1217, 1999.

Sergio Fajardo and H. Jerome Keisler. *Model theory of stochastic processes*, volume 14 of *Lecture Notes in Logic*. Association for Symbolic Logic, Urbana, IL; A K Peters, Ltd., Natick, MA, 2002. ISBN 1-56881-167-5; 1-56881-172-1.

James C. D. Fisher and Isa E. Hafalir. Matching with aggregate externalities. *Mathematical Social Sciences*, 81:1–7, 2016.

Tamás Fleiner. A fixed-point approach to stable matchings and some applications. *Mathematics of Operations Research*, 28:103–126, 2003.

David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.

Alfred Galichon. *Optimal transport methods in economics*. Princeton University Press, Princeton, N.J., 2016. ISBN 978-0-691-17276-7.

Alfred Galichon, Scott Duke Kominers, and Simon Weber. Costly concessions: An empirical framework for matching with imperfectly transferable utility. 2016.

Michael Greinecker and Konrad Podczeck. Purification and roulette wheels. *Economic Theory*, 58: 255–272, 2015.

Neil E. Gretsky, Joseph M. Ostroy, and William R. Zame. The nonatomic assignment model. *Economic Theory*, 2:103–127, 1992.

Neil E. Gretsky, Joseph M. Ostroy, and William R. Zame. Perfect competition in the continuous assignment model. *Journal of Economic Theory*, 88:60–118, 1999.

Peter J. Hammond, Mamoru Kaneko, and Myrna Holtz Wooders. Continuum economies with finite coalitions: Core, equilibria, and widespread externalities. *Journal of Economic Theory*, 49: 135–168, 1989.

Sergiu Hart, Werner Hildenbrand, and Elon Kohlberg. On equilibrium allocations as distributions on the commodity space. *Journal of Mathematical Economics*, 1:159–166, 1974.

Werner Hildenbrand. *Core and equilibria of a large economy*. Princeton University Press, Princeton, N.J., 1974.

Werner Hildenbrand. Distributions of agents' characteristics. *Journal of Mathematical Economics*, 2:129–138, 1975.

Murat Iyigun and Randall P. Walsh. Building the family nest: Premarital investments, marriage markets, and spousal allocations. *The Review of Economic Studies*, 74:507–535, 2007.

Mamoru Kaneko and Myrna Holtz Wooders. The core of a game with a continuum of players and finite coalitions: the model and some results. *Mathematical Social Sciences*, 12:105–137, 1986.

Mamoru Kaneko and Myrna Holtz Wooders. The nonemptiness of the f-core of a game without side payments. *International Journal of Game Theory*, 25:245–258, 1996.

Yakar Kannai. Continuity properties of the core of a market. *Econometrica*, 38:791–815, 1970.

Alexander S. Kelso, Jr. and Vincent P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50:1483–1504, 1982.

Donald E. Knuth. *Mariages stables et leurs relations avec d'autres problèmes combinatoires*. Les Presses de l'Université de Montréal, Montreal, Que., 1976. ISBN 0-8405-0342-3.

Fuhito Kojima. *Recent Developments in Matching Theory and Their Practical Applications*, volume 1 of *Econometric Society Monographs*, pages 138–175. Cambridge University Press, 2017.

Tjalling C. Koopmans and Martin Beckmann. Assignment problems and the location of economic activities. *Econometrica*, 25:53–76, 1957.

Lauwerens Kuipers and Harald Niederreiter. *Uniform distribution of sequences.* John Wiley & Sons, New York-London-Sydney, 1974.

Peter Loeb and Yeneng Sun. Purification and saturation. *Proceedings of the American Mathematical Society*, 137:2719–2724, 2009.

Dorothy Maharam. Automorphisms of products of measure spaces. *Proceedings of the American Mathematical Society*, 9:702–707, 1958.

Andreu Mas-Colell. On the continuous representation of preorders. *International Economic Review*, 18:509–513, 1977.

Andreu Mas-Colell. On a theorem of Schmeidler. *Journal of Mathematical Economics*, 13:201–206, 1984.

Andreu Mas-Colell. *The theory of general economic equilibrium*, volume 9 of *Econometric Society Monographs.* Cambridge University Press, Cambridge, 1989. A differentiable approach.

Konrad Menzel. Large matching markets as two-sided demand systems. *Econometrica*, 83: 897–941, 2015.

Mitsunori Noguchi and William R. Zame. Competitive markets with externalities. *Theoretical Economics*, 1:143–166, 2006.

Georg Nöldeke and Larry Samuelson. Investment and competitive matching. *Econometrica*, 83: 835–896, 2015.

Georg Nöldeke and Larry Samuelson. The implementation duality. 2017. Cowles Foundation Discussion Paper No. 1993R.

Kalyanapuram Rangachari Parthasarathy. *Probability measures on metric spaces.* Probability and Mathematical Statistics, No. 3. Academic Press, Inc., New York–London, 1967.

Michael Peters and Aloysius Siow. Competing premarital investments. *Journal of Political Economy*, 110:592–608, 2002.

Boris Pittel. The average number of stable matchings. *SIAM Journal on Discrete Mathematics*, 2: 530–549, 1989.

Konrad Podczeck. On purification of measure-valued maps. *Economic Theory*, 38:399–418, 2009.

Martine Quinzii. Core and competitive equilibria with indivisibilities. *International Journal of Game Theory*, 13:41–60, 1984.

Hiroo Sasaki and Manabu Toda. Two-sided matching problems with externalities. *Journal of Economic Theory*, 70:93–108, 1996.

David Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7: 295–300, 1973.

Lloyd S. Shapley and Martin Shubik. The assignment game. I: The core. *International Journal of Game Theory*, 1:111–130, 1971.

Yeneng Sun. The exact law of large numbers via Fubini extension and characterization of insurable risks. *Journal of Economic Theory*, 126:31–69, 2006.

Veeravalli S. Varadarajan. On the convergence of sample probability distributions. *Sankhyā*, 19: 23–26, 1958.

Cédric Villani. *Topics in optimal transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, R.I., 2003. ISBN 0-8218-3312-X.

Cédric Villani. *Optimal transport: Old and New*, volume 338 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, 2009. ISBN 978-3-540-71049-3.

Jianwei Wang and Yongchao Zhang. Purification, saturation and the exact law of large numbers. *Economic Theory*, 50:527–545, 2012.