

The Emergence of Cooperative Outcomes: A Stochastic Stability Analysis*

Petros Milionis[†] Sohei Kaihatsu

University of Groningen Bank of Japan

November 2011

Abstract

The concept of the core capturing allocations that exhaust all possible gains from trade is central in economic theory. Are all allocations included in the core, though, equally likely to emerge in a decentralized exchange economy? We investigate this question by studying the evolution of exchange in such an economy via a dynamic trading process in the spirit of Edgeworth's recontracting. Assuming that along this process individual agents might make mistakes with a small probability, we characterize the stochastically stable distribution of the process and use this distribution to obtain a measure of the relative frequency with which each core allocation will emerge in the long run. Based on this measure, we demonstrate that there exist particular allocations inside the core which are more likely to emerge relative to others. These are allocations which are welfare improving for a larger number of agents, or alternatively, allocations which can be reached relatively easily with only a small number of agents involved in trading.

Keywords: Core, Edgeworth's Recontracting, Stochastic Stability, Housing Economy, Incentives to Trade, Decentralization.

JEL Classification: C71, C73, D51, D61.

*We are indebted to Roberto Serrano for his guidance and support throughout this project. We have also benefited greatly from the valuable input of Amy Greenwald, John Wicks, as well as that of seminar participants at Brown University Economics and Computer Science Departments. The views expressed in this paper are those of the authors and not necessarily those of the Bank of Japan.

[†]Corresponding Author: Petros Milionis, Department of Economics, Econometrics & Finance, University of Groningen, P.O. Box 800, 9700 AV Groningen, The Netherlands, e-mail: P.Milionis@rug.nl.

1 Introduction

One of the central and most widely used solution concepts in modern economic theory is definitely the *core*. This notion of the core has a long history in economics, dating back to the work of Francis Ysidro Edgeworth (Edgeworth 1881), who first conceived the idea of allocations that are immune to coalitional deviations. However, his idea of the core did not receive much attention in subsequent work and gradually fell into oblivion up until the early 1950s. At that point in time, the notion of the core was reinvented and formalized, as part of the emerging field of game theory, by Donald B. Gillies (Gillies 1953) and Lloyd Shapley (Shapley 1952). Finally, a few years later, also the connection between the work of Edgeworth and the modern game-theoretic view of the core was established by Martin Shubik (Shubik 1959) .

Since then, the core has remained one of the fundamental concepts in the theory of cooperative games as well as that of general equilibrium. Its wide usage and applicability to various economic environments stems from the simple and intuitive way in which the core characterizes the allocations that exhaust all gains from trade or cooperation, more generally. These are exactly the allocations upon which no group of agents can improve. Thus, given the interest of rational agents not to leave any gains unexploited, we should expect that if an allocation *"...does not belong to the core, one should not expect to see it as prediction of the theory if there is full cooperation,"* as Serrano (2009) puts it.

Yet, in many economic environments the number of allocations that exhaust all gains from trade can be quite large¹. In such environments, the prediction given by the core as a solution concept regarding the outcome of cooperation will typically not be that sharp. Hence, in the case of a non-singleton core, a natural question to ask is whether all allocations

¹There are also economic environments in which the core may be empty and thus there is again no clear-cut prediction regarding what outcome one should expect. Although this is an important issue as well, we will not explicitly deal with it in this research. Moreover, as shown in Scarf (1967), balancedness provides a key sufficient condition for the non-emptiness of the core. For more information on this issue, the reader is referred to the survey by Serrano (2009) as well as the discussion in Chapter 13 of Osborne and Rubinstein (1994).

included in the core constitute equally good predictions of the expected outcome. Specifically, should one treat all these allocations as having the same likelihood of emerging from trade or cooperation or are there particular allocations which are more likely to emerge compared to the others.

In order to answer this question, we need to shift our focus from the static description of the core and to a more careful investigation of how the allocations inside the core can generally be reached. Interestingly, this problem was also first addressed by Edgeworth, who proposed a dynamic trading process, the well-known *recontracting process* (Edgeworth 1881), in an attempt to rationalize his notion of the core. This process is based on a series of proposed allocations made by various groups of agents, each of which can only be countered by an improved proposal from a subsequent group. The process continues until a proposal is made that can not be countered by any group of agents and which constitutes, by definition, a core allocation exhausting the gains from trade. Thus, in Edgeworth's mind, the economic agents would contract and recontract until a final settlement is reached.

The most important feature of the Edgeworthian recontracting process is the fact that it is "*...based on the same behavioral postulate, blocking by coalitions, that is used to define the solution concept, the core,*" as pointed out by Jerry Green (Green 1974). Yet, Edgeworth himself was notoriously unclear in his descriptions about several aspects of recontracting and often presented conflicting examples². As a result, the alleged convergence of the process to the core was initially treated simply as a conjecture. It was only about hundred years later that this convergence result was formally established by Feldman (1974) and Green (1974), who formalized close variants of the recontracting process and demonstrated its convergence to the set of core allocations in a large class of economic environments, as Edgeworth had hypothesized.

This line of research was recently revisited by Serrano and Volij (2008), who utilized a similar variant of the Edgeworthian recontracting process in order to capture dynamic

²For more details regarding this point, there is an excellent summary of the difficulties in interpreting Edgeworth's original views on recontracting in Section 1.4 of Serrano and Volij (2008).

trading in the context of the housing economy of Shapley and Scarf (1974). One important novelty of their approach was that Serrano and Volij also allowed for the possibility of agents making small *mistakes* in the process. This feature added an element of randomness into the recontracting process and enabled them to invoke the criterion of *stochastic stability* in the same spirit as Kandori, Mailath, and Rob (1993) and Young (1993) did in a non-cooperative game-theoretic setup. Hence, by considering the limiting case with the likelihood of mistakes going to zero, they were able to identify the allocations that would emerge in the long run with positive probability. These allocations are referred to as *stochastically stable* and can be thought of as the most natural ones in an environment where mistakes are possible, although highly unlikely events.

Yet, as Serrano and Volij demonstrated, in the context of the housing economy this set of stochastically stable allocations does not always coincide with the core, nor with any particular subset of it, such as the strong core or the competitive allocations³. Moreover, it is shown that the set may even include allocations which are individually rational although not in the core. These findings naturally pose an important puzzle for the dynamic analysis of exchange and cooperation, as dynamic outcomes appear to be different than those predicted by static solution concepts. With that in mind Serrano and Volij hypothesized that: *"It may happen that some regions of the core are hard to access, while some non-core allocations may have strong dynamic attractor properties."* Hence, the identification of the subset of all allocations that constitute strong dynamic attractors provides an additional reason for the study of dynamic exchange, beyond the potential multiplicity problem.

The approach followed by Serrano and Volij (2008), though, is limited by the fact that it only allows for the distinction between stochastically stable allocations, namely allocations emerging with positive probability in the long run, and non-stochastically stable ones,

³This is the case whenever the preferences of the agents over the set of available houses are non-strict. On the other hand, if preferences do not allow for indifferences, Serrano and Volij (2008) show that there is a unique stochastically stable allocation corresponding to the single competitive equilibrium of the economy. However, given the importance of indifferences in economies with indivisible goods, it may be misleading to rule out indifferences.

namely those whose occurrence is just a zero probability event. Nevertheless, any statement regarding which allocations are more likely to emerge from dynamic exchange -whether part of the core or not- require also some discrimination among allocations which are stochastically stable based on their relative frequency of occurrence. Therefore, it is necessary that we go beyond the simple characterization of the set of stochastically stable allocations in order to address such questions.

This is where the main contribution of this research lies. Using the same setup as Serrano and Volij (2008), Shapley and Scarf's housing economy, we will proceed to characterize the long-run probability distribution of the recontracting process, often referred to as *stochastically stable distribution*. Particularly, we will provide both explicit numerical results based on the recent algorithm of Greenwald and Wicks (2005), as well as analytical results using a novel graph-theoretic approach. Treating the probabilities attributed by the distribution to each core allocation as well as other allocations outside the core as a measure of the relative likelihood with which they will emerge in the long run, we will demonstrate the existence of particular allocations, which are more likely to emerge compared to others from coalitional interaction.

Given this finding, we will then attempt to account for the observed differences in the relative frequencies of emergence for various allocations by investigating the factors affecting these frequencies. These differences, as we will demonstrate, are associated with particular characteristics of each allocation such as the number of agents for which the allocation is welfare improving, on the one hand, and the number of agents whose cooperation is necessary in order for the allocation to be achieved. Particularly, our main result is that allocations which improve upon a larger number of agents, as well as allocations that can be reached with a relatively small number of agents involved in trading are more natural to emerge in an environment of unrestrained coalitional interaction. This suggests that, although the economy gravitates towards allocations that are beneficial to more and more agents, considerations regarding the degree of decentralization behind each exchange, as reflected in

the number of agents involved in it, will also matter for the cooperative outcomes observed in simple exchange economies.

The remainder of this paper is organized as follows. In Section 2, we begin with a brief overview of the rationale behind stochastic stability analysis and contrast our approach with those of the existing literature. The setting of the housing economy, which we are going to use in the rest of the paper, together with the basics of the recontracting process is described in Section 3. Then, in Section 4, the focus shifts on our main object of interest, the stochastically stable distribution, where we demonstrate how the distribution can help us distinguish among the different core and non-core allocations. Our main theoretical results on the factors influencing coalitional recontracting and the observed differences in the relative frequencies of emergence of the various allocations are presented in Section 5. Finally, the paper ends with a few concluding remarks.

2 The Basics of Stochastic Stability Analysis

The notion of stochastic stability originates from the analysis of stochastic dynamical systems. It was developed specifically to capture the stability of a dynamical system that is constantly subject to small random shocks. In economics it was first introduced by Foster and Young (1989), who proposed it in the context of evolutionary games as an alternative to the notion of evolutionary stability. This was followed by the seminal contributions of Kandori, Mailath, and Rob (1993) and Young (1993) who employed stochastic stability analysis in order to refine the predictions given by Nash equilibrium in the context of static non-cooperative games. Their approach paved the way for a long list of contributions in the field of non-cooperative game theory that provided equilibrium refinement results based on the application of stochastic stability.

The typical application of stochastic stability in a game-theoretic setting can be roughly described as follows. First, a particular static game is specified. Then a repetition of this

game is considered together with a dynamic process prescribing the strategy choices for each individual player in each round of the game. These strategy choices, though, are not necessarily based on rational best-response calculations but on the history of previous interactions, thus leading to a large degree of inertia in the strategy choices of individual players. Yet, apart from inertia, the process also incorporates some element of randomness giving rise to a stochastic dynamical system. The resulting system is then analyzed based on the notion of stochastic stability.

Assuming the dynamic process satisfies the Markov property and given an appropriate definition of the process' states, the analysis of the resulting system can be greatly simplified and its stability easily assessed. One has to just look in this case at the invariant distribution of the process, which given the appropriate assumptions on the nature of the shocks is easily shown to be unique. Specifically, most of the focus of the literature has been on the limiting case where randomness vanishes. This captures the outcomes of the process in an environment where small random shocks are present, although quite uncommon events.

Here is where the methodology of Freidlin and Wentzell (1984), which enables the analysis of the dynamical system's flow via graph-theoretic techniques, has been proven quite useful. The states of the process are treated as vertices of a complete weighted directed graph whose weights on each edge correspond to the respective transition probabilities. Studying the flows on this graph, one can assess the likelihood of process reaching each state. This will be proportional to the product of the edges of a spanning tree rooted at the corresponding vertex. Hence, by comparing the minimum spanning trees for each vertex, one can assess which states of the process are stochastically stable, namely have a positive probability of being obtained in the long run.

This is exactly the approach followed by most of the literature, which concentrates on distinguishing states of process that are stochastically stable from states that are not, and hence their relative likelihood of occurrence is zero. However, even among states that are stochastically stable, there can be large differences in these relative likelihoods. In the

presence of such differences, it is natural to think of states with higher long-run probabilities as better predictors regarding the outcome of the process. Yet, any distinction between such states would only be possible if the process' limit invariant distribution was known.

Fortunately, Freidlin and Wentzell (1984) also provides us with the means for computing that distribution. The relative likelihood of each state of the process is just proportional to the total probability of all minimum spanning trees corresponding to that particular state. This approach, however, has the drawback of requiring the construction all possible minimum spanning trees for each vertex/state, which makes the computation of the process invariant distribution, even in the limiting case where randomness vanishes, quite complicated. More importantly, the actual dimensionality of the problem rises exponentially with the number of states. Thus, most researchers have refrained from using the distribution in their analysis.

One exception is the recent work of Klaus, Bochet, and Walzl (forthcoming) who discuss the extent to which the stochastically stable distribution of an appropriately defined dynamic trading process can provide a refined prediction regarding the outcome of exchange in economies with multiple indivisible goods. Yet, their analysis is only suggestive of this possibility. This is because the above authors just analyze simple examples of particular 3-agent economies, where the stochastically stable distribution can be easily calculated⁴. Hence, apart from providing a useful illustration, the evidence presented by Klaus, Bochet, and Walzl (forthcoming) is inconclusive regarding the potential of the stochastically stable distribution to discriminate among stochastically stable states

Recently, though, Greenwald and Wicks (2005) have proposed an exact algorithm that can deal with the computational difficulties of working with the stochastically stable distribution. The idea of the algorithm is based on a series of elementary operations performed on the Markov matrix of the process. These operations end up reducing its dimensionality to a degree that is manageable. Then, keeping track of these operations, the invariant

⁴Actually, the authors' calculations, which are included in the online appendix, are indicative of the complications involved in the computation of the stochastically stable distribution even for simple 3-agent economies.

distribution embedded in the original matrix can be recovered. This procedure enables the computation of the invariant distribution directly from the Markov matrix, instead of going through all possible minimum spanning trees. This simplifies a lot the problem of working with the process' stochastically stable distribution.

Moreover, in addition to the use of the Greenwald-Wicks algorithm, in the sections below we will be presenting some novel graph-theoretic results that will enable us to identify which elements of the stochastically stable set have a higher likelihood of occurrence in the long run. Based on those results, we will try to trace back differences in long-run relative frequencies among states of dynamic process to features of the underlying economic environment. This way, we believe, that our contribution will provide an important extension of the standard stochastic stability analysis. Now, in order to facilitate a comparison with the existing literature, we will restrict our attention to the housing economy of Shapley and Scarf (1974). We will begin with the necessary preliminaries regarding this particular environment.

3 The Housing Economy Framework

3.1 Basic Definitions

The housing economy of Shapley and Scarf (1974) is a simple exchange economy with indivisible goods. Formally, it comprises a finite set of agents N and a finite set of houses H where $|H| = |N|$. Each agent i is endowed with one house e_i , and has a complete and transitive preference relation \succsim_i over the whole set of houses. Following the original article of Shapley and Scarf, an abundant literature on housing economies has emerged with different variants reflecting different assumptions regarding the ownership of houses or the degree of indivisibility. For example, Bogomolnaia and Moulin (2001) considered the case of social ownership of houses, while Abdulkadiroglu and Sonmez (1999) distinguished between existing and new tenants. Recently, Athanassoglou and Sethuraman (forthcoming) introduced

the idea of agents owning fractional endowments⁵. However, for our purposes, we will stick to the original formulation of Shapley and Scarf.

In this environment, we define a *coalition*, S , to be any group of agents with $S \subseteq N$ and a *house allocation* x any redistribution of the existing houses where $\forall i \in N, \exists! j \in N : x_i = e_j$, so that each agent continues to occupy one and only one house. We call an allocation x to be feasible for coalition S if allocation x can just be achieved via a redistribution of the endowed houses among the members of coalition S . Moreover, we say that a coalition S can improve upon allocation x if there exists another allocation y which is feasible for coalition S and which is strictly preferred to x by all members of the coalition, namely $\forall i \in S, y_i \succ_i x_i$.

Note that if an allocation is feasible for coalition S , then it should also be feasible for any coalition $T \supseteq S$. Given that, we define A_S to be the set of allocations that are feasible for coalition S and S_x^{\min} the smallest possible coalition for which allocation x is feasible. Since any allocation is feasible for the grand coalition N , we can denote the set of all possible house allocations as simply A_N . Two important subsets of A_N are the set of individually rational allocations IR and the set of core allocations C , where $C \subseteq IR$. The former consists of all allocations $x \in A_N$, for which no agent resides in a house deemed worse than his or her endowed one, namely $\forall i \in N, x_i \succeq_i e_i$. The latter consists of all allocations $x \in A_N$ which cannot be improved upon by any coalition S , or more formally $\nexists S \subseteq N \wedge y \in A_S : \forall i \in S, y_i \succ_i x_i$.⁶

The set of core allocations is an important reference point for the housing economy. It consists of those house allocations that are not dominated by any other allocation and thus exhaust all possible gains from house-trading. Moreover, this set is always non-empty, though typically it includes more than one allocations. Yet, to what extent do these allocations constitute good predictions for the outcome of trade in our simple exchange economy? To

⁵For a more detailed description of the literature the reader can have a look in summary provided by Athanassoglou and Sethuraman (forthcoming).

⁶The set of core allocations should not be confused with the set of strong core allocations $SC \equiv \{x \in A_N \mid \nexists S \subseteq N \wedge y \in A_S : \forall i \in S, y_i \succeq_i x_i \wedge \exists j \in S \ni y_j \succ_j x_j\}$ which is a subset of C is composed only of allocations which cannot be strictly improved upon.

answer this question we need to move beyond the original static analysis of the housing economy and consider the nature of dynamic trading within the economy.

3.2 The Recontracting Process

Although one could conceive several different ways in which the dynamic exchange of houses could take place among the agents, given our focus on the solution concept of the core, there is a natural candidate for this dynamic analysis, the Edgeworthian recontracting process⁷. This is a dynamic trading process proposed by Edgeworth (1881) in an attempt to motivate the notion of the core. For this reason, it is also based on the principle of coalitional blocking used in the definition of the core. The process begins at the endowment allocation with various coalitions of agents contracting and recontracting iteratively up until a house allocation is reached upon which no coalition can improve. This is what Edgeworth called a "final settlement," which in our terminology corresponds to allocation inside the core.

Yet, it was Feldman (1974) and Green (1974) that managed to formally demonstrate that the process, as envisioned by Edgeworth, does actually converge to the core under reasonable assumptions. According to their setup, there is a given exchange economy, with fixed preferences and endowments, in which trade takes place repeatedly over time. In each time period, there is one particular coalition that forms and seeks a reallocation of the goods -in our case houses- owned by the members of the coalition, that would leave each individual member strictly better off. Once such a reallocation is found, it is implemented with positive probability. Otherwise, the existing allocation of houses from the previous period persists. Note that an essential feature of the process is that the ownership of houses does not change over time. Owners simply offer their houses to different tenants withholding the right to request them back at any point in time. Thus, the recontracting process can be described

⁷It is interesting to contrast this point with the non-cooperative game-theoretic applications of stochastic stability. In those cases there is no natural dynamic process to be associated with the underlying static game and hence different researchers have come up with alternative suggestions. For example one can look at the choices made by Kandori, Mailath, and Rob (1993) on the one hand and Young (1993) on the other when studying 2×2 coordination games.

briefly as follows:

- A.** In each period t , the economy begins at the allocation $x(t)$, with $x(0) = e$.
- B.** At the beginning of the period, a coalition S is randomly chosen with probability p_S , with $p_S > 0, \forall S \subseteq N$.
- C.** If that coalition can improve on $x(t)$ with $y_S \in A_S^8: y_i \succ_i x_i, \forall i \in S$, then with probability $\delta > 0$ the new allocation will be⁹:

$$x(t+1) = \left\{ \begin{array}{ll} (y_S, x_{-S}(t)) & \text{if } x_{-S}(t) \in A_{N-S} \\ (y_S, e_{-S}) & \text{if } x_{-S}(t) \notin A_{N-S} \end{array} \right\}.$$

Otherwise $x(t+1) = x(t)$.

Note that, by construction, the recontracting process will never leave a core allocation, once being reached. Hence, the question of convergence to the core is just a matter of whether the process will eventually reach such an allocation. This fact is true for any exchange economy fulfilling the conditions identified by Feldman (1974) and Green (1974). Yet, in the context of the simple housing economy, these conditions are not satisfied due to the inherent indivisibilities of the traded goods¹⁰. This fact raises an additional complication for our analysis, which is the possibility of the recontracting process constantly cycling among non-core allocations, as discussed in Serrano and Volij (2008). For this reason, in our subsequent treatment, we will also have to allow for the possibility of non-convergence.

⁸Note that there could be more than one allocation to satisfy this condition. In such cases, one of them is randomly selected.

⁹Actually, there are more than one possibilities regarding the "fate" of agents not included in the selected coalition after a coalitional move. Here we follow Serrano and Volij (2008) and assume that the allocation of houses outside the coalition is not affected, as long as this includes no house owned by a member of the coalition. Alternative assumptions, though, will not alter our main results.

¹⁰Hence, the failure of convergence does not depend on the actual details of the recontracting process.

3.3 Introducing Mistakes

Having presented the fashion in which the dynamic exchange of houses will take place in the economy, our next issue is how to measure the relative likelihood with which each allocation, either in the core or not- is going to emerge from recontracting. One possibility here would be to consider several identical housing economies, in terms of preferences and endowments, observe the outcome of the recontracting process in each of them, and then consider the relative frequency with which each allocation is reached. Instead of that, we will follow Serrano and Volij (2008) and looking into a variant of the recontracting process that also incorporates for some degree of randomness in the agents' decision-making. This randomness will allow for any possible transition to occur, and as a result the recontracting process will not converge to any particular allocation but constantly transit among various allocations. This way we will be able to assess the likelihood of core allocations as well as non-core cycles by observing the number of times the given allocation will be visited by the process.

To justify this additional randomness in the recontracting process, Serrano and Volij emphasize the possibility of agents making small mistakes in the coalition formation stage. Particularly, they assume that when a coalition forms, there is a small probability that one or more agents might make a mistake and agree on a reallocation that does not leave them better off. In addition, Serrano and Volij distinguish between two kinds of mistakes. The minor ones are committed if an agent agrees on a reallocation $z \rightarrow z'$, even though $z_i \sim_i z'_i$. The serious ones are committed whenever an agent agrees on a reallocation $z \rightarrow z'$, even though $z_i \succ_i z'_i$. Both kinds of mistakes are assumed only to happen with a very small probability, with the serious ones occurring less frequently. Thus, if the probability of each agent making a minor mistake is a small number $\epsilon \in (0, 1)$, then the probability of a serious mistake should be a even smaller one ϵ^λ , with $\lambda > 1$ ¹¹.

Since mistakes are taken to be small probability events, transitions for which mistakes

¹¹Thus, mistakes here mimic the role played by the various mutations or experimentations schemes often assumed in evolutionary game theory literature.

are necessary will also be less likely to occur. This number of mistakes involved can actually be measured in the following way. For each particular transition $z \longrightarrow z'$, where $z, z' \in A_N$, induced by a particular coalition S for which the transition $z \longrightarrow z'$ is feasible, namely $z' \in A_S$, we can compute the number of minor mistakes, $n_I(S, z, z') = |\{i \in S : z_i \sim_i z'_i\}|$, and the number of serious mistakes, $n_W(S, z, z') = |\{i \in S : z_i \succ_i z'_i\}|$, that are necessary in order for the transition to occur. These two numbers can then be combined into a weighted number of mistakes:

$$n(S, z, z') = \lambda n_W(S, z, z') + n_I(S, z, z').$$

Of course, for each transition $z \longrightarrow z'$, there could be more than one coalitions that can induce it. Therefore, the total probability of the recontracting process transiting from allocation z to allocation z' can be obtained by the formula:

$$\Pr(x_{t+1} = z' | x_t = z) = \sum_{\{S | n(S, z, z') > 0\}} p_S \cdot \epsilon^{n(S, z, z')} + \sum_{\{S | n(S, z, z') = 0\}} p_S \cdot \delta.$$

Note that the formula captures both transitions based on mistakes as well as mistake-free transitions. Observe, though, that the probability of a particular transition $z \longrightarrow z'$ is inversely related to the number of mistakes involved and at the same time is independent of the past history of the process before reaching allocation z . Thus, our recontracting process can be treated as a Markov process over the set of all possible allocations of the underlying housing economy.

4 The Stochastically Stable Distribution

4.1 The Recontracting Process as a Markov Process

Having presented the underlying framework of the housing economy and discussed the details of the recontracting process, we now turn to the actual predictions regarding the outcome

of dynamic exchange in the economy. Here is where the Markov nature of the recontracting process with mistakes will ease the analysis substantially. Particularly, the introduction of mistakes, even as events of very small probability, allows for any possible transition to take place, making the resulting process irreducible. This means that it will encompass a unique invariant distribution.

This distribution will capture the probabilities with which the process is going to visit each allocation in the long run. Of course, this long-run distribution, denoted by μ^ϵ , will depend on the likelihood of an agent committing a mistake ϵ . Yet, as we mentioned in the Section 2, our focus will be on the limiting case when $\epsilon \rightarrow 0$. Given the nature of the recontracting process with mistakes the limit distribution, $\mu^* = \lim_{\epsilon \rightarrow 0} \mu^\epsilon$, is well defined and will be our main object of interest in the subsequent analysis. Following the literature we will refer to μ^* as the stochastically stable distribution and to all allocations $x \in A_N$ for which $\mu^*(x) > 0$ as the stochastically stable states.

Our focus on the stochastically stable distribution can simply be justified by the fact that the distribution captures exactly the probabilistic outcome of dynamic exchange in the economy under the assumption that mistakes are possible though highly unlikely events. Thus, in such an environment, the knowledge of μ^* conveys all the necessary information regarding which states are stochastically stable, namely visited by the process for a positive proportion of time, as well as what their long-run probabilities are. This information is what we are going to compare with the static predictions given by the core, with the probabilities reflecting the relative frequencies with which each allocation is going to emerge in the long-run from dynamic exchange. Yet, before we getting there, we believe that a brief discussion regarding how the stochastically stable distribution can be obtained is **in order**.

4.2 Computing the Stochastically Stable Distribution

Contrary to the simple identification of the stochastically stable states, which can be done using the "mistake-counting" techniques of Kandori, Mailath, and Rob (1993) and Young

(1993), the computation of the stochastically stable distribution can actually be far more complicated. The reason, as we discussed in Section 2, has to do with the fact that the latter requires the construction of all possible minimum-resistance trees, while for the former the construction of one such tree suffices. Thus, the distribution can be computed analytically only for housing economies consisting of only a small number of agents, since the number of minimum-resistance trees that need to be considered rises exponentially as the size of the economy increases.

Fortunately though, even for those cases where the analytic computation of the stochastically stable distribution becomes intractable, there is the possibility of resorting to the numerically algorithm of Greenwald and Wicks (2005). This is -to our knowledge- the only existing algorithm for computing the exact, and not approximate, limit invariant distribution of any perturbed Markov matrix, as the perturbation rate goes to zero. The algorithm is based on a series of simple operations on Markov matrices and can thus be easily implemented in the context of the recontracting process with mistakes to obtain the stochastically stable distribution.

However, before proceeding with actual computations, some additional discussion regarding particular aspects of the recontracting process is necessary. First of all, we need to specify the probabilities p_S with which each coalition S is going to be selected at the coalition formation stage. Here we will make the natural assumption and treat all coalitions as equiprobable. This means that $p_S = \frac{1}{2^{|N|-1}}$. Of course, one could possibly think of alternative setups in which these probabilities varied depending on the size of S . For example, Green (1974), in his treatment of the recontracting process, assumed that smaller coalitions were more likely to form, possibly due to the existence of communication costs that increase with the size of the coalition. Yet, we prefer the assumption of all coalitions forming with equal probability, because any other assumption would simply add a slight bias to the obtained stochastically stable distribution.

A second important parameter of the recontracting process with mistakes is the weight λ

falling on serious mistakes, namely on the mistakes that make agents worse off compared to the ones that leave them simply indifferent. In theory, given that such mistakes are considered to be more serious, this weight λ has to be strictly greater than 1. Yet, in practice, in order for the numerical algorithm to deliver robust results, this number has to also be substantially higher from 1. This is because with values of λ very close to 1, the two types of mistakes become almost indistinguishable and this may influence the resulting distribution. As long as this is not the case though, the actual value of λ will have no effect on the stochastically stable distribution.

A last key variable in the recontracting process is δ . This parameter reflects a factor of inertia in the process. This is because it captures the probability of the coalition proceeding to implement any reallocation of houses that involves no mistakes. The closer δ is to 0, the more inertia there is in the process. Moreover, as Serrano and Volij (2008) point out δ should in principle lie in the interval $(0, \frac{1}{|A|})$. So when implementing the Wicks-Greenwald algorithm we will simply let $\delta = \frac{1}{|A|+1}$ to minimize the additional inertia. Nevertheless, it is not hard to show that this choice will have no effect on the relative likelihoods of the various core allocations, which is our main focus here¹².

Finally, in light of the critique of Bergin and Lipman (1996), it is important to emphasize that the resulting stochastically stable distribution is not the result of an ad hoc "mistake model" applied to the recontracting process, but one which is consistent with the basic principles of coalitional decision-making. Apart from that, we would like to point out that our results also extend to alternative mistake models maintaining the distinction between serious and less serious mistakes. These include natural alternatives such as the case where

¹²The parameter δ only affects the mistake-free transitions. Thus, if all the stochastically stable allocations belong to the set of the core, the value of δ will have no effect at all on the stochastically stable distribution. This is because, for the economy to transit from a core allocation, there has to be at least one agent who makes a mistake and hence there are no mistake-free way-outs. Needless to say, there could be other mistake-free transitions in the process, but those would only be transitions from transient states and hence less relevant.

The case where the value of δ will matter for the stochastically stable distribution is when some non-core cycles are also stochastically stable. In those cases, a larger value of δ increases the chances of the economy staying in the cycle, which leads to higher long-run probabilities for the cycle allocations. However, even in those cases, the effect of δ on the distribution is not substantive.

the weight on each serious mistake is higher than λ if an agent i accepts a house ranked lower than his or her endowment e_i , or the case where the seriousness of the mistake increases for agent i the lower the new house $x_i(t+1)$ is ranked compared to the previous one $x_i(t)$.

4.3 Core Allocations and the Stochastically Stable Distribution

Having defined the stochastically stable distribution and presented the details of how it can be computed, we move on to discuss how it can be used in two simple examples. For tractability reasons, we concentrate on housing economies consisting of just three agents, so that the number of possible allocations is only six and the number of coalitions to consider is seven. In both cases we compute the exact stochastically stable distribution of the recontracting process with mistakes and contrast it with the set of core allocations. This will allow us to demonstrate how coalitional interaction may favor a particular allocation inside the core. Moreover, using these examples as a starting point, we attempt to provide some intuition for the observed differences in the relative frequencies with which the various core allocations will emerge in the long run. This intuition will then be formalized in Section 5, in a way that will enable us to pin down the main factors that influencing the long-run outcome of the coalitional recontracting process.

4.3.1 Example I: Different Incentives to Trade

As our first example, let us consider the following 3-agent housing economy, where the three agents hold the following preferences over the three available houses:

$$e_2 \succ_1 e_3 \succ_1 e_1$$

$$e_1 \sim_2 e_2 \succ_2 e_3$$

$$e_1 \sim_3 e_3 \succ_3 e_2$$

Note that only three out of the six possible house allocations are individually rational and these allocations also constitute the economy's core:

$$C = IR = \left\{ \begin{array}{l} c_1 = (e_1, e_2, e_3) \\ c_2 = (e_2, e_1, e_3) \\ c_3 = (e_3, e_2, e_1) \end{array} \right\}.$$

These three core allocations will all emerge with positive probability in the long run. Yet, their relative frequencies, as given by the stochastically stable distribution, are actually quite different:

$$c_1 \longrightarrow 0.3333$$

$$c_2 \longrightarrow 0.5000$$

$$c_3 \longrightarrow 0.1666$$

To understand these differences in the relative frequencies, note that agents 2 and 3 already own a top ranked house. This means that they do not have any incentive to exchange the house they own with that of any other agent. The only agent who really wants a house exchange is agent 1; her aim should be to exchange it with either the house of agent 2 or the house of agent 3, with a slight preference for that of agent 3. Yet, in both cases, the other party is not interested in that exchange. Hence, this reallocation of houses will not occur without mistakes, in which case agent 1 has to remain at her endowed house.

Thus, in this case, the variation in the relative frequencies seems to be the result from the differences in the incentives that agents have to exchange their houses.

4.3.2 Example II: Degree of Decentralization

Let us now turn to another example of the following 3-agent housing economy with the preference profile:

$$e_2 \succ_1 e_1 \succ_1 e_3$$

$$e_1 \sim_2 e_3 \succ_2 e_2$$

$$e_1 \sim_3 e_2 \succ_3 e_3$$

Again a moment's thought should reveal that, in this economy, there are four individually rational allocations from which the following three constitute the economy's core:

$$C = \left\{ \begin{array}{l} c_1 = (e_1, e_3, e_2) \\ c_2 = (e_2, e_1, e_3) \\ c_3 = (e_2, e_3, e_1) \end{array} \right\}.$$

As in the previous example, the three core allocations will emerge in the long run with different relative frequencies:

$$\begin{aligned} c_1 &\longrightarrow 0.3596 \\ c_2 &\longrightarrow 0.3596 \\ c_3 &\longrightarrow 0.2808 \end{aligned}$$

What is striking in this example is that although there is a core allocation c_3 which gives everybody a top ranked house, this allocation will emerge less often compared to the remaining two. This is because the allocation c_3 has the drawback that it can be achieved only if all 3 agents get together and decide to exchange houses, while the allocations c_1 and c_2 can be obtained by pairwise trades between agents $\{1, 2\}$ and $\{1, 3\}$ respectively.

Thus, in this case, the variation in the relative frequencies may seem to be the result of the differences in the degree of decentralization behind each core allocation.

5 Accounting for the Different Frequencies of Emergence

5.1 Overview of Main Results

The examples presented in the previous section underscore two factors that seem to influence the relative frequencies with which different allocations are going to emerge from the recontracting process. These are the incentives that each agent has to trade in each case and

the degree of decentralization behind each allocation. Particularly, as the above examples indicate, coalitional recontracting process seems to favor allocations which leave more agents better off, as well as allocations that can be achieved via smaller coalitions of agents. In this section we will make an attempt to formalize this intuition via a series of theoretical results. However, let us begin with a overview of our main theoretical contribution.

First of all, we will try to capture the incentives factor by the proportion of the agents that are inclined towards a particular exchange to those that are indifferent to it, while for the decentralization factor we will consider the minimum size of the coalition that needs to form in order for the allocation to be obtained. Using these definitions we will compare different house allocations in terms of the underlying incentives and decentralization.

Definition 1 (Decentralization) *An allocation $z \in A_N$ is said to be more decentralized than another allocation z' , a fact that we will denote as $d(z) > d(z')$, if $|S_z^{\min}| < |S_{z'}^{\min}|$, namely the size of the minimum coalition S_z^{\min} is smaller than that of the minimum coalition $S_{z'}^{\min}$.*

Definition 2 (Incentives) *An allocation $z \in A_N$ is said to be beneficial to more agents than another allocation z' , a fact that we will denote as $b(z) > b(z')$, if $|\{i \in N : z_i \succ_i e_i\}| > |\{i \in N : z'_i \succ_i e_i\}|$, namely the number of agents residing in a house strictly preferred to their endowed one is greater in z compared to z' .¹³*

Our main proposition now will demonstrate how these differences in the incentives and decentralization among the distinct allocations of houses that are possible in this economy will be reflected in different long-run probabilities of emergence.

Proposition 1 (Relative Frequencies) *Let z and z' be two house allocations that share the same degree of decentralization, while z is beneficial to more agents compared to z' , then*

¹³This definition can be stated alternatively using the minimum coalitions S_z^{\min} and $S_{z'}^{\min}$ as $|\{i \in S_z^{\min} : z \succ_i e\}| > |\{i \in S_{z'}^{\min} : z' \succ_i e\}|$, simply because agents not included in the minimum coalitions will necessarily resides in their endowed houses.

we must have that $\mu^*(z) \geq \mu^*(z')$. Similarly, let z and z' be two house allocations where the number of agents benefiting is the same, but z is more decentralized than z' , then again we must have $\mu^*(z) \geq \mu^*(z')$.

Hence, as it is evident from the statement of the proposition, even within the core of the economy, allocations that are beneficial for a larger fraction of agents and allocations that can be obtained in a more decentralized fashion should be considered more natural to emerge in an environment of unfettered coalitional interaction. We should note that these two factors could work in the same or in the opposite directions, in which case the long-run outcome of the recontracting process will be determined by their interplay. However, even in such cases where the theory does not provide a clear prediction for the long-run probabilities, investigating the roles of these two factors would still enable us to rationalize the reasons why particular allocations are more likely to emerge as a result of cooperation.

In the following sections we will formally establish the above proposition by displaying the key results one by one starting from basic principles.

5.2 Transition Resistances

We will begin with a few definitions that are going to be central in our subsequent analysis of dynamic trading. First of all, we need to have a measure of how easy it is to get from any one allocation to another. The following definition provides us with such a measure.

Definition 3 (Resistance) *Let A_N be the set of feasible allocations. For any two allocations $z, z' \in A_N$, we can define the resistance of the direct transition $z \longrightarrow z'$ as,*

$$r(z, z') \equiv \min_{S \subseteq N} n(S, z, z') = \min_{S \subseteq N} [\lambda \cdot n_w(S, z, z') + n_I(S, z, z')], \quad \lambda > 1,$$

which is equivalent to the least weighted number of mistakes needed for the given transition to take place. Similarly, we can define the resistance $\tilde{r}(z, z')$ of an indirect transition from

z to z' as the least weighted number of mistakes needed along any path $z \longrightarrow x_1 \longrightarrow \dots \longrightarrow x_{L-1} \longrightarrow z'$ that begins at z and ends at z' , namely:

$$\tilde{r}(z, z') \equiv \min_{\{x_l\}_{l=0}^L: x_0=z \wedge x_L=z'} \sum_{l=0}^{L-1} r(x_l, x_{l+1}).$$

Based on this notion of resistance, we can construct measures of how easy it is to reach and how easy it is to leave any particular allocation.

Definition 4 (Minimum Resistance In and Out) For each allocation $z \in A_N$, we can define the minimum resistance to leave that allocation as $r_{out}^{\min}(z) \equiv \min_{z' \in A_N, z' \neq z} r(z, z')$, and the minimum resistance to reach that allocation as $r_{in}^{\min}(z) \equiv \min_{z' \in A_N, z' \neq z} r(z', z)$.¹⁴

Using the above definition, we can characterize the set of core allocations as follows:

Proposition 2 (The Core) A feasible allocation $z \in A_N$ is included in the core if and only if $r_{out}^{\min}(z) > 0$. Hence, $C = \{z \in A_N \mid r_{out}^{\min}(z) > 0\}$.

Proof. Suppose $z \in C$. Then this means that $\nexists S \subseteq N \wedge y \in A_S : y_i \succ_i z_i, \forall i \in S$. Consider now any $z' \in A_N$ with $z' \neq z$ and pick any coalition S for which the transition z to z' is feasible. Given that $z \in C, \exists i \in S : z'_i \sim_i z_i \vee z'_i \prec_i z_i$. Hence, we have that $n_I(S, z, z') \geq 1 \vee n_W(S, z, z') \geq 1 \implies n(S, z, z') > 0 \implies r(z, z') > 0$ and since z' was chosen arbitrarily, it must be that $r_{out}^{\min}(z) > 0$.

Suppose $z \notin C$. Then this means that $\exists S \subseteq N \wedge y \in A_S : y_i \succ_i z_i, \forall i \in S$. Hence, when coalition S forms, $\exists z' = (y, z_{-S}) \vee (y, e_{-S}) : z'_i \succ_i z_i, \forall i \in S$. This means that: $n(S, z, z') = 0 \implies r(z, z') = 0 \implies r_{out}^{\min}(z) = 0$. ■

Let us now extend our notion of resistance into any partition \mathbf{P} of the set of feasible allocations A_N .

¹⁴The concepts of r_{out}^{\min} and r_{in}^{\min} have some common features with Ellison (2000). However, our results are not directly comparable to those of Ellison.

Definition 5 (Cell Resistance) Let $\mathbf{P} = \{P^1, P^2, \dots, P^I\}$ be a partition of the set of feasible allocations A_N or any subset of it. For any two cells $P^i, P^j \in \mathbf{P}$, we can define the resistance of the transition $P^i \rightarrow P^j$ as $\tilde{r}(P^i, P^j) \equiv \min_{z^i \in P^i, z^j \in P^j} \tilde{r}(z^i, z^j)$, namely the least weighted number of mistakes needed to transit from an allocation in cell P^i to an allocation in cell P^j .

Based on this extended notion of resistance, we can similarly define:

Definition 6 (Minimum Resistance In and Out for Cells) Let $\mathbf{P} = \{P^1, P^2, \dots, P^I\}$ be a partition of the set of feasible allocations A_N or any subset of it. For any cell $P^i \in \mathbf{P}$, we can define the minimum resistance of leaving that cell to any other as $\tilde{r}_{out}^{\min}(P^i|\mathbf{P}) \equiv \min_{P^j \in \mathbf{P}, P^j \neq P^i} \tilde{r}(P^i, P^j)$ and the minimum resistance of reaching that cell from any other as $\tilde{r}_{in}^{\min}(P^i|\mathbf{P}) \equiv \min_{P^j \in \mathbf{P}, P^j \neq P^i} \tilde{r}(P^j, P^i)$.

Despite the generality of the above definition in what follows we will be focusing on a particular partition of the set of feasible allocations A_N . Given that the recontracting process constitutes a Markov process over the set of feasible allocations, a natural partition of that set would be based on the recurrent classes of the process.¹⁵ This partition will be denoted by $\mathbf{R} = \{R^1, R^2, \dots, R^\Omega, T\}$ where each R^ω represents a distinct recurrent class, while the cell T contains all the transient states of the process. Moreover, as pointed out by Serrano and Volij (2008), in the context of a housing economy, we should expect to see two particular types of recurrent classes.

Lemma 1 (Serrano and Volij) *The recurrent classes of the unperturbed recontracting process take the following two forms:*

1. singleton recurrent classes, each of which contains one core allocation.

¹⁵A state z of a Markov process is called recurrent if once visited by the process, there is positive probability that the process will return to it at some point, namely $P(x_{t+N} = z | x_t = z) > 0$. Otherwise that state is called transient. Now any ergodic set of recurrent allocations constitutes a recurrent class of the process.

2. non-singleton recurrent classes, in each of which the allocations are individually rational, but are not core allocations.

With this result of Serrano and Volij in mind, we can now state and prove the following claim:

Claim 1 (Recurrent Classes I) *For every recurrent class R^ω of the unperturbed recontracting process, it must be that $\tilde{r}_{out}^{\min}(R^\omega|\mathbf{R}) > 0$.*

Proof. If R^ω is a singleton recurrent class, $R^\omega = \{z\}$, then z constitutes a core allocation for which we know by Proposition 2 that $r_{out}^{\min}(z) > 0$. Yet, $\tilde{r}_{out}^{\min}(R^\omega|\mathbf{R}) = \min_{R^\psi \in \mathbf{R}, R^\psi \neq R^\omega} \tilde{r}(R^\omega, R^\psi) = \min_{R^\psi \in \mathbf{R}, R^\psi \neq R^\omega} \min_{z' \in R^\psi} r(z, z') = \min_{z' \in A_N, z' \neq z} r(z, z') = r_{out}^{\min}(z)$. Hence, the statement holds trivially.

Suppose now that R^ω is a non-singleton recurrent class. If $\tilde{r}_{out}^{\min}(R^\omega|\mathbf{R}) = 0$ this requires that $\exists w \in R^\omega \wedge y \notin R^\omega : r(w, y) = 0$ meaning $\exists S \subseteq N : w \xrightarrow{S} y \wedge y_i \succ_i w_i, \forall i \in S$. Hence, even for the unperturbed process we should have $P(x_{t+1} = y | x_t = w) > 0$. However, this contradicts the fact that R^ω is a recurrent class. As $\forall z \in R^\omega \wedge \forall z' \notin R^\omega$ it must be that $P(x_{t+1} = z' | x_t = z) = 0$. ■

Furthermore, for the recurrent classes that do not include the endowment allocation e we can state and prove the following claim.

Claim 2 (Recurrent Classes II) *The recontracting process can leave any recurrent class R^ω not containing the endowment allocation with either one minor or one serious mistake, namely $\tilde{r}_{out}^{\min}(R^\omega|\mathbf{R}) = \{1, \lambda\}$.*

Proof. Select a recurrent class, R^ω and pick an allocation z included in that class, $z \in R^\omega$. Since $z \neq e$ the process can exit R^ω via any singleton coalition $\{i\}$ where $i \in S_z^{\min}$ with agents i simply requesting his endowed house back. Since z should be an individually rational allocation by Lemma 1, it must be that $z_i \succsim_i e_i, \forall i \in N$. Thus, if $\exists j \in S_z^{\min} : z_j \sim_j e_j$ then $n(\{j\}, z, e) = 1$. If not, namely $z_i \succ_i e_i, \forall i \in S_z^{\min}$ then $n(\{i\}, z, e) = \lambda$. In the latter case, there could be another allocation z' to which the process can move with just one

minor mistake. In any case though, we should note that the process can always leave and return to the endowment allocation with at most one serious mistake. Combining that with Claim 1 we get that $r_{out}^{\min}(z) = \{1, \lambda\}$ and since z was picked randomly we conclude that: $\tilde{r}_{out}^{\min}(R^\omega | \mathbf{R}) = \{1, \lambda\}$. ■

5.3 Spanning Trees and Stochastic Stability

Having described our various resistance measures, we will now discuss how these measures can be used to assess the stochastic stability of particular allocations as well as sets of allocations. This will provide an alternative to the method of constructing minimum spanning trees that has become standard in the literature on stochastic stability.

Definition 7 (Spanning Trees) *Consider a complete weighted directed graph with vertex set the set of recurrent classes $\mathbf{R}^* = \{R^1, R^2, \dots, R^\Omega\}$ where for each edge $R^\psi \rightarrow R^\omega$ the corresponding weight equals the value $\tilde{r}(R^\psi, R^\omega)$, a weight to which we will also refer to as the resistance of the given edge. Starting from this graph we will then consider all spanning subgraphs connecting all vertices in a way that creates a unique directed path leading to a particular recurrent class R^ω . This type of graph is called a spanning tree routed at recurrent class R^ω or simply an R^ω -tree and will be denoted subsequently by T_{R^ω} . Moreover, for any such tree T_{R^ω} , we will define its resistance $r(T_{R^\omega})$ as the sum of the resistance of all its edges,*

$$r(T_{R^\omega}) = \sum_{(R^\chi, R^\psi) \in T_{R^\omega}} \tilde{r}(R^\chi, R^\psi).$$

From the above definition, it is evident that there is a lower bound to the resistance of any spanning tree.

Remark 1 (Tree Resistance Lower Bound) *Let T_{R^ω} be a spanning tree rooted at recurrent class R^ω . Its resistance should then be bounded below by $r(T_{R^\omega}) \geq \sum_{R^\psi \neq R^\omega} \tilde{r}_{out}^{\min}(R^\psi | \mathbf{R}^*)$.*

Proof. Using the definition of $r(T_{R^\omega})$ we have that

$$\begin{aligned}
r(T_{R^\omega}) &= \sum_{(R^\chi, R^\psi) \in T_{R^\omega}} \tilde{r}(R^\chi, R^\psi) \\
&= \sum_{\{R^\chi \neq R^\omega : (R^\chi, R^\psi) \in T_{R^\omega}\}} \tilde{r}(R^\chi, R^\psi) \\
&= \tilde{r}(R^\alpha, R^\beta) + \tilde{r}(R^\beta, R^\gamma) + \dots \\
&\geq \tilde{r}_{out}^{\min}(R^\alpha | \mathbf{R}^*) + \tilde{r}_{out}^{\min}(R^\beta | \mathbf{R}^*) + \dots \\
&= \sum_{\{R^\chi \neq R^\omega : (R^\chi, R^\psi) \in T_{R^\omega}\}} \tilde{r}_{out}^{\min}(R^\chi | \mathbf{R}^*) \\
&= \sum_{R^\psi \neq R^\omega} \tilde{r}_{out}^{\min}(R^\psi | \mathbf{R}^*).
\end{aligned}$$

■

Yet, this lower bound will not always be reached as not all trees will have the same resistance. Hence, there will be differences in the resistance of different spanning trees rooted at the same recurrent class as well as differences in the resistance of spanning trees rooted at different recurrent classes. From these differences we can determine which recurrent classes will be stochastically stable.

Definition 8 (Stochastic Stability) *For each recurrent class R^ω of the unperturbed reconstructing process, let us consider the set \mathcal{T}_{R^ω} of all spanning trees T_{R^ω} rooted at R^ω . Based on this set we can define the stochastic potential $sp(R^\omega)$ of a recurrent class R^ω as the minimum resistance attained by any of its spanning trees, namely $sp(R^\omega) \equiv \min_{T_{R^\omega} \in \mathcal{T}_{R^\omega}} r(T_{R^\omega})$. Using the notion of stochastic potential, we can state that a recurrent class R^ω is stochastically stable if and only if*

$$sp(R^\omega) = \min_{\{R^1, R^2, \dots, R^\Omega\}} sp(R^\psi) \equiv sp^{\min}.$$

Also a feasible allocation $z \in A_N$ is said to be stochastically stable if it belongs to one of these recurrent classes with minimum stochastic potential. The set of all these stochastically stable allocations is often referred to as the stochastically stable set and will be denoted by SSA .

As it is obvious from the above definition, if the process has only one recurrent class, that class will have minimum stochastic potential and hence will be stochastically stable. In what follows, we will provide a way of identifying which classes are stochastically stable in case the number of recurrent classes is greater than one utilizing our minimum resistance measure in and out of each class. However, before presenting this method it is important to understand how a minimum resistance tree for each recurrent class can be constructed. This can be done using the following procedure which is based on Dijkstra's Graph Search Algorithm:

1. Consider the set of recurrent classes $\mathbf{R}^* = \{R^1, R^2, \dots, R^\Omega\}$.
2. Pick a particular recurrent class, R^ω , and identify among all classes in $\mathbf{A}_0 = \mathbf{R}^*/\{R^\omega\}$ the one for which the resistance of the edge $R^\psi \longrightarrow R^\omega$ satisfies $\tilde{r}(R^\psi, R^\omega) = \min_{R^i \in \mathbf{A}_0} \tilde{r}(R^i, R^\omega)$.
3. Connect R^ψ to R^ω , and then look among the remaining classes in $\mathbf{A}_1 = \mathbf{R}^*/\{R^\omega, R^\psi\}$ to identify the one, R^χ for which the $\min\{\tilde{r}(R^\chi, R^\omega), \tilde{r}(R^\chi, R^\psi)\} = \min_{\substack{R^i \in \mathbf{A}_1 \\ R^j \in \{R^\omega, R^\psi\}}} \tilde{r}(R^i, R^j)$.
4. Continue in a similar fashion up until all recurrent classes are connected to R^ω .

Proposition 3 (Stochastically Stable Set) *Consider the set of recurrent classes $\mathbf{R}^* = \{R^1, R^2, \dots, R^\Omega\}$ and a function $G : \mathcal{P}(\mathbf{R}^*) \longrightarrow \mathcal{P}(\mathbf{R}^*)$ where $G(\mathbf{R}) \equiv \{R^\omega \in \mathbf{R} \mid \tilde{r}_{out}^{\min}(R^\omega | \mathbf{R}) - \tilde{r}_{in}^{\min}(R^\omega | \mathbf{R}) \geq 0\}, \forall \mathbf{R} \subseteq \mathbf{R}^*$. Then the set of stochastically stable allocations, denoted as \mathbf{R}^{SS} , is the largest subset \mathbf{R} of \mathbf{R}^* for which $G(\mathbf{R}) = \mathbf{R}$.*

Proof. The fact that there is always a subset of \mathbf{R}^* for which $G(\mathbf{R}) = \mathbf{R}$ is obvious, as this holds true when \mathbf{R}^* is the singleton $\{R^\omega\}$. This is because trivially $\tilde{r}_{out}^{\min}(R^\omega | \{R^\omega\}) - \tilde{r}_{in}^{\min}(R^\omega | \{R^\omega\}) = 0 - 0 = 0$.¹⁶

¹⁶Things are not that trivial when \mathbf{R}^* is not a singleton. For example in case $\mathbf{R}^* = \{R^\psi, R^\omega\}$, then there are three possibilities:

$$\begin{aligned} \tilde{r}(R^\psi, R^\omega) &> \tilde{r}(R^\omega, R^\psi) \rightsquigarrow G(\mathbf{R}^*) = \{R^\psi\}, \\ \tilde{r}(R^\psi, R^\omega) &< \tilde{r}(R^\omega, R^\psi) \rightsquigarrow G(\mathbf{R}^*) = \{R^\omega\}, \\ \tilde{r}(R^\psi, R^\omega) &= \tilde{r}(R^\omega, R^\psi) \rightsquigarrow G(\mathbf{R}^*) = \{R^\psi, R^\omega\}. \end{aligned}$$

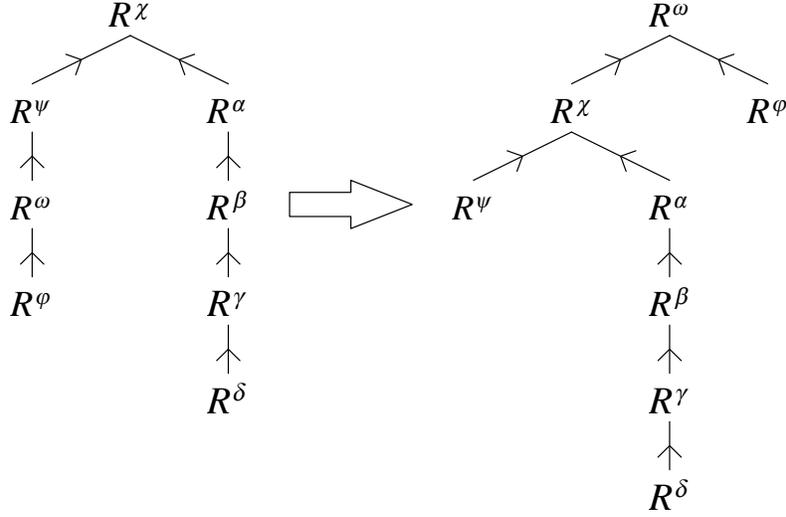


Figure 1: Tree Surgery 1

Having made this point, we just have to show that if $\bar{\mathbf{R}}$ is the largest subset of \mathbf{R}^* for which $G(\bar{\mathbf{R}}) = \bar{\mathbf{R}}$ then $\bar{\mathbf{R}} = \mathbf{R}^{SS}$. The proof proceeds by contradiction.

First, suppose that $\exists R^\omega \in \bar{\mathbf{R}}$ with $R^\omega \notin \mathbf{R}^{SS}$. Given R^ω let us find the recurrent class, R^χ , in $\bar{\mathbf{R}}$ for which $\tilde{r}(R^\chi, R^\omega) = \tilde{r}_{in}^{\min}(R^\omega | \bar{\mathbf{R}})$ and consider a minimum resistance tree for R^χ . Without loss of generality, we can construct the tree for R^χ , as shown in the left part of Figure 1. This can be convert into a R^ω tree by severing the edge $R^\psi \leftarrow R^\omega$ and adding the edge $R^\omega \leftarrow R^\chi$. The new tree, shown in the right part of Figure 1, should have a resistance of,

$$\begin{aligned}
r(T_{R^\omega}) &= r(T_{R^\chi}) - \tilde{r}(R^\omega, R^\psi) + \tilde{r}(R^\chi, R^\omega) \\
&\leq r(T_{R^\chi}) - [\tilde{r}_{out}^{\min}(R^\omega | \bar{\mathbf{R}}) - \tilde{r}_{in}^{\min}(R^\omega | \bar{\mathbf{R}})] \\
&\leq r(T_{R^\chi}),
\end{aligned}$$

given the equality $\tilde{r}(R^\chi, R^\omega) = \tilde{r}_{in}^{\min}(R^\omega | \bar{\mathbf{R}})$, the fact that $\tilde{r}(R^\omega, R^\psi) \geq \tilde{r}_{out}^{\min}(R^\omega | \bar{\mathbf{R}})$ by definition and the property that $\tilde{r}_{out}^{\min}(R^\omega | \bar{\mathbf{R}}) - \tilde{r}_{in}^{\min}(R^\omega | \bar{\mathbf{R}}) \geq 0$ due to $R^\omega \in \bar{\mathbf{R}}$.

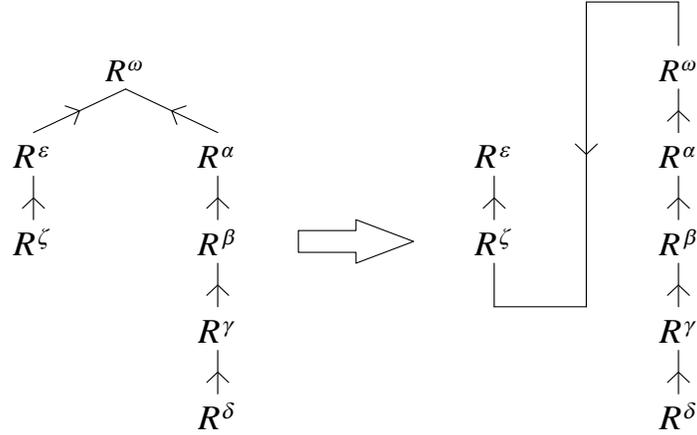


Figure 2: Tree Surgery 2

Moreover, since T_{R^x} was a minimum resistance tree for R^x , we get that $sp(R^\omega) \leq r(T_{R^\omega}) \leq r(T_{R^x}) = sp(R^x)$. If $R^x \in \mathbf{R}^{SS}$, then the only way for the above expression to hold is $sp(R^x) = sp(R^\omega) = sp^{\min}$, which contradicts the fact that $R^\omega \notin \mathbf{R}^{SS}$. If, on the other hand, $R^x \notin \mathbf{R}^{SS}$, then we can repeat the above exercise replacing R^ω with R^x , which given the finiteness of \mathbf{R}^{SS} at some point will yield the necessary contradiction.

Secondly, suppose that $\exists R^\omega \in \mathbf{R}^{SS}$ with $R^\omega \notin \bar{\mathbf{R}}$. Since $R^\omega \notin \bar{\mathbf{R}}$ we will need to consider the set $\bar{\mathbf{R}} \cup \{R^\omega\}$. For that set we should have:

$$\tilde{r}_{out}^{\min}(R^\omega | \bar{\mathbf{R}} \cup \{R^\omega\}) - \tilde{r}_{in}^{\min}(R^\omega | \bar{\mathbf{R}} \cup \{R^\omega\}) < 0$$

as otherwise $\bar{\mathbf{R}} \cup \{R^\omega\} = G(\bar{\mathbf{R}} \cup \{R^\omega\})$, contradicting the fact that $\bar{\mathbf{R}}$ is the largest subset of \mathbf{R}^* for which $G(\bar{\mathbf{R}}) = \bar{\mathbf{R}}$.

Now, without loss of generality, consider a minimum resistance tree for R^ω , as shown in the left part of Figure 2. For that tree, by construction, we have that $\min(\tilde{r}(R^\alpha, R^\omega), \tilde{r}(R^\epsilon, R^\omega)) = \tilde{r}_{in}^{\min}(R^\omega | \bar{\mathbf{R}} \cup \{R^\omega\})$. Let us now find the recurrent class, R^ζ , in $\bar{\mathbf{R}} \cup \{R^\omega\}$ for which $\tilde{r}(R^\omega, R^\zeta) = \tilde{r}_{out}^{\min}(R^\omega | \bar{\mathbf{R}} \cup \{R^\omega\})$. Then we can construct a new tree, shown in the right part of Figure 2, by severing the edge $R^\omega \leftarrow R^\epsilon$ and adding the edge $R^\zeta \leftarrow R^\omega$. Note that the resulting

tree would be a R^ε -tree while the tree's resistance would be,

$$\begin{aligned}
r(T_{R^\varepsilon}) &= r(T_{R^\omega}) + \tilde{r}(R^\omega, R^\zeta) - \tilde{r}(R^\varepsilon, R^\omega) \\
&\leq r(T_{R^\omega}) + [\tilde{r}_{out}^{\min}(R^\omega | \bar{\mathbf{R}} \cup \{R^\omega\}) - \tilde{r}_{in}^{\min}(R^\omega | \bar{\mathbf{R}} \cup \{R^\omega\})] \\
&< r(T_{R^\omega}),
\end{aligned}$$

where we have used the equality $\tilde{r}(R^\omega, R^\zeta) = \tilde{r}_{out}^{\min}(R^\omega | \bar{\mathbf{R}} \cup \{R^\omega\})$, the fact that $\tilde{r}(R^\varepsilon, R^\omega) \geq \tilde{r}_{in}^{\min}(R^\omega | \bar{\mathbf{R}} \cup \{R^\omega\})$ and the property that $\tilde{r}_{out}^{\min}(R^\omega | \bar{\mathbf{R}} \cup \{R^\omega\}) - \tilde{r}_{in}^{\min}(R^\omega | \bar{\mathbf{R}} \cup \{R^\omega\}) < 0$. Yet, since T_{R^ω} was a minimum resistance tree for R^ω , we get that $sp(R^\omega) = r(T_{R^\omega}) > r(T_{R^\varepsilon}) \geq sp(R^\varepsilon)$ contradicting the stochastic stability of R^ω . ■

Having established Proposition 3 we can now easily identify the set of stochastically stable recurrent classes using the following algorithm.

1. Consider the set of recurrent classes $\mathbf{R}^* = \{R^1, R^2, \dots, R^\Omega\}$.
2. For the each $R^\omega \in \mathbf{R}^*$ compute the expression $\tilde{r}_{out}^{\min}(R^\omega | \mathbf{R}^*) - \tilde{r}_{in}^{\min}(R^\omega | \mathbf{R}^*)$.
3. Construct a new set $\mathbf{B}_1 = \{R^\omega \in \mathbf{R}^* \mid \tilde{r}_{out}^{\min}(R^\omega | \mathbf{R}^*) - \tilde{r}_{in}^{\min}(R^\omega | \mathbf{R}^*) \geq 0\}$. If $\mathbf{B}_1 = \mathbf{R}^*$ then $\mathbf{R}^* = \mathbf{R}^{SS}$.
4. Otherwise repeat the process and for the each $R^\omega \in \mathbf{B}_1$ compute the expression $\tilde{r}_{out}^{\min}(R^\omega | \mathbf{B}_1) - \tilde{r}_{in}^{\min}(R^\omega | \mathbf{B}_1)$.
5. Construct the set $\mathbf{B}_2 = \{R^\omega \in \mathbf{B}_1 \mid \tilde{r}_{out}^{\min}(R^\omega | \mathbf{B}_1) - \tilde{r}_{in}^{\min}(R^\omega | \mathbf{B}_1) \geq 0\}$. If $\mathbf{B}_2 = \mathbf{B}_1$ then $\mathbf{B}_1 = \mathbf{R}^{SS}$.
6. Otherwise continue the process up until you find two sets $\mathbf{B}_n = \mathbf{B}_{n+1}$.

An important corollary of the above proposition is that the difference between the minimum resistances in and out should be the same for each stochastically stable recurrent class. Since this corollary will be helpful in our subsequent analysis, let us provide a complete proof.

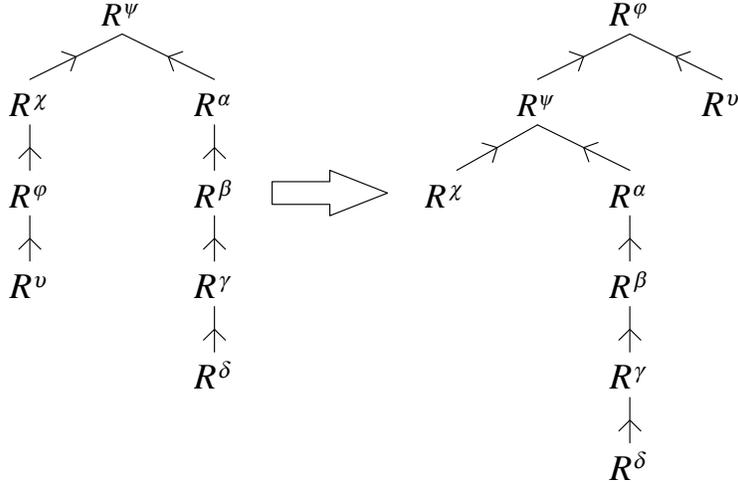


Figure 3: Tree Surgery 3

Corollary 1 (Stochastically Stable Recurrent Classes) *For any two recurrent classes $R^\psi, R^\omega \in \mathbf{R}^{SS}$, we must have:*

$$\tilde{r}_{out}^{\min}(R^\psi | \mathbf{R}^{SS}) - \tilde{r}_{in}^{\min}(R^\psi | \mathbf{R}^{SS}) = \tilde{r}_{out}^{\min}(R^\omega | \mathbf{R}^{SS}) - \tilde{r}_{in}^{\min}(R^\omega | \mathbf{R}^{SS}).$$

Proof. Suppose to the contrary that there exists a recurrent class $R^\psi \in \mathbf{R}^{SS}$ for which we have:

$$0 \leq \tilde{r}_{out}^{\min}(R^\psi | \mathbf{R}^{SS}) - \tilde{r}_{in}^{\min}(R^\psi | \mathbf{R}^{SS}) < \tilde{r}_{out}^{\min}(R^\omega | \mathbf{R}^{SS}) - \tilde{r}_{in}^{\min}(R^\omega | \mathbf{R}^{SS}),$$

where $R^\omega \in \mathbf{R}^{SS} / \{R^\psi\}$. Without loss of generality, we can construct a minimum resistance R^ψ -tree, as shown in the left part of Figure 3. Now given R^ψ let us find the recurrent class $R^\varphi \neq R^\psi$ for which $\tilde{r}(R^\psi, R^\varphi) = \tilde{r}_{in}^{\min}(R^\varphi | \mathbf{R}^{SS})$ and let us convert the above tree to a R^φ tree by severing the edge $R^z \leftarrow R^\psi$ and adding the edge $R^\varphi \leftarrow R^\psi$. The resulting tree,

shown in the right part of Figure 3, will then have a resistance of:

$$\begin{aligned}
r(T_{R^\varphi}) &= r(T_{R^\psi}) - \tilde{r}(R^\varphi, R^\chi) + \tilde{r}(R^\psi, R^\varphi) \\
&\leq r(T_{R^\psi}) - [\tilde{r}_{out}^{\min}(R^\phi | \mathbf{R}^{SS}) - \tilde{r}_{in}^{\min}(R^\phi | \mathbf{R}^{SS})] \\
&< r(T_{R^\psi}),
\end{aligned}$$

where we have used the equality $\tilde{r}(R^\psi, R^\varphi) = \tilde{r}_{in}^{\min}(R^\varphi | \mathbf{R}^{SS})$, the fact that $\tilde{r}(R^\varphi, R^\chi) \geq \tilde{r}_{out}^{\min}(R^\phi | \mathbf{R}^{SS})$ by definition and the property that $[\tilde{r}_{out}^{\min}(R^\phi | \mathbf{R}^{SS}) - \tilde{r}_{in}^{\min}(R^\phi | \mathbf{R}^{SS})] > 0$. So, we get that $sp(R^\varphi) \leq r(T_{R^\varphi}) < r(T_{R^\psi}) = sp(R^\psi)$, which contradicts the fact that $R^\psi \in \mathbf{R}^{SS}$. ■

5.4 The Stochastically Stable Distribution

Up until now, we have only dealt with the identification of the stochastically stable set, which consists of those recurrent classes that have minimum stochastic potential. In this subsection we will turn to their long-run probability distribution and establish the main theoretical result of this paper. Yet, before we get there, let us formally define this long-run probability distribution.

Definition 9 (Stochastically Stable Distribution) *Consider the recontracting process with mistakes and let the mistake probability $\varepsilon \in (0, 1)$. This process as we know is irreducible and therefore it should have a unique invariant distribution which we denote by μ^ε . Moreover, since the recontracting process with mistakes represents a regular perturbation of mistake-free process, it can be shown Young (1993) that the limit $\mu^* \equiv \lim_{\varepsilon \rightarrow 0} \mu^\varepsilon$ exists and it also constitutes a well defined probability distribution. This limit invariant distribution is what we refer to as the stochastically stable distribution of the recontracting process and it has the property that $\mu^*(z) > 0$ for each allocation $z \in A_N$ if and only if allocation is stochastically stable, namely $z \in SSA$.*

To understand now better the connection between Definitions 8 and 9 we will utilize the following lemma due to Freidlin and Wentzell (1984).

Lemma 2 (Freidlin and Wentzell) *For every finite irreducible Markov process, its unique invariant distribution μ^ε is given by*

$$\mu^\varepsilon(z) = \frac{q(z)}{\sum_{z' \in A_N} q(z')},$$

where $q(z) = \sum_{T_{R^\omega} : r(T_{R^\omega}) = sp^{\min}(z' \rightarrow z'')} \prod_{(z' \rightarrow z'') \in T_{R^\omega}} \Pr(x_{t+1} = z'' | x_t = z')$ and $z \in T_{R^\omega}$.

In order to make use of the Freidlin and Wentzell's lemma, it will be important to have a measure for the number of minimum resistance transitions in and out of each particular allocation. Hence, we need the following definition.

Definition 10 (Number of Minimum Resistance Transitions) *Let x be a feasible allocation and let us denote by $\#r_{out}^{\min}(x)$ the number of allocations $z' \in A_N$ for which $r(x, z') = r_{out}^{\min}(x)$, and by $\#r_{in}^{\min}(x)$ the number of allocations z' for which $r(z', x) = r_{in}^{\min}(x)$.*

Using these definitions, we state and prove the following claim.

Claim 3 (Factors Behind Relative Frequencies) *For any two stochastically stable allocations x, y we have $\mu^*(x) \geq \mu^*(y)$ if $\#r_{out}^{\min}(x) \leq \#r_{out}^{\min}(y)$ and $\#r_{in}^{\min}(x) \geq \#r_{in}^{\min}(y)$.*

Proof. Since $x, y \in SSA$, we should have that $\mu^*(x), \mu^*(y) > 0$. So let us consider their ratio $\frac{\mu^*(x)}{\mu^*(y)}$:

$$\begin{aligned} \frac{\mu^*(x)}{\mu^*(y)} &= \frac{\lim_{\varepsilon \rightarrow 0} \mu^\varepsilon(x)}{\lim_{\varepsilon \rightarrow 0} \mu^\varepsilon(y)} \\ &= \lim_{\varepsilon \rightarrow 0} \frac{\frac{q(x)}{\sum_{z' \in A_N} q(z')}}{\frac{q(y)}{\sum_{z' \in A_N} q(z')}} \\ &= \lim_{\varepsilon \rightarrow 0} \frac{q(x)}{q(y)} \\ &= \lim_{\varepsilon \rightarrow 0} \frac{\sum_{T_{RX} : r(T_{RX}) = sp^{\min}(z' \rightarrow z'')} \prod_{(z' \rightarrow z'') \in T_{RX}} \Pr(x_{t+1} = z'' | x_t = z')}{\sum_{T_{R^\psi} : r(T_{R^\psi}) = sp^{\min}(z' \rightarrow z'')} \prod_{(z' \rightarrow z'') \in T_{R^\psi}} \Pr(x_{t+1} = z'' | x_t = z')} \end{aligned}$$

where $x \in R^x$ and $y \in R^\psi$ are the recurrent classes to which the two allocations belong.

Let us now consider a spanning tree T_{R^x} rooted at recurrent class R^x and particularly at allocation x and look into the product $\prod_{(z' \rightarrow z'') \in T_{R^x}} \Pr(x_{t+1} = z'' | x_t = z')$.

$$\begin{aligned}
& \prod_{(z' \rightarrow z'') \in T_{R^x}} \Pr(x_{t+1} = z'' | x_t = z') \\
= & \prod_{(z' \rightarrow z'') \in T_{R^x}} \left(\sum_{\{S | n(S, z', z'') > 0\}} p_S \cdot \varepsilon^{n(S, z', z'')} + \sum_{\{S | n(S, z', z'') = 0\}} p_S \cdot \delta \right) \\
= & c_{00} \delta^{|A_N|} + c_{11} \delta^{|A_N|-1} \varepsilon + c_{12} \delta^{|A_N|-1} \varepsilon^\lambda \\
& + c_{21} \delta^{|A_N|-2} \varepsilon^2 + c_{22} \delta^{|A_N|-2} \varepsilon^{\lambda+1} + c_{23} \delta^{|A_N|-2} \varepsilon^{2\lambda} \\
& + \dots + c_{(|A_N|)(|A_N|)} \varepsilon^{(|A_N|-1)\lambda+1} + c_{(|A_N|)(|A_N|+1)} \varepsilon^{|A_N|\lambda} \\
= & c_0 \delta^{|A_N|} + \sum_{g=1}^{|A_N|} \sum_{h=1}^{g+1} c_{gh} \delta^{|A_N|-g} \varepsilon^{\lambda(h-1)+(g-h+1)},
\end{aligned}$$

where the coefficients will be given by the expression:

$$\begin{aligned}
c_{gh} = & \sum_{\substack{\tilde{S}, \bar{S}, \hat{S} \\ |\tilde{S}|=|A_N|-g, |\bar{S}|=g-h+1, |\hat{S}|=h-1}} \left(\prod_{\substack{(z' \rightarrow z'') \in T_{R^x} \\ r(z', z'')=0}} p_{\tilde{S}} |_{n(\tilde{S}, z', z'')=0} \right. \\
& \times \prod_{\substack{(z' \rightarrow z'') \in T_{R^x} \\ r(z', z'')=1}} p_{\bar{S}} |_{n(\bar{S}, z', z'')=0} \times \left. \prod_{\substack{(z' \rightarrow z'') \in T_{R^x} \\ r(z', z'')=\lambda}} p_{\hat{S}} |_{n(\hat{S}, z', z'')=0} \right)
\end{aligned}$$

and

$$c_0 = \sum_S \prod_{\substack{(z' \rightarrow z'') \in T_{R^x} \\ r(z', z'')=0}} p_S |_{n(S, z', z'')=0}.$$

Here it is important to note that for any spanning tree T_{R^x} with resistance $r(T_{R^x})$, the coefficients c_{gh} for all terms of the above polynomial $c_{gh} \delta^{|A_N|-g} \varepsilon^{\lambda(h-1)+(g-h+1)}$ which have $\lambda(h-1) + (g-h+1) = (h-1)(\lambda-1) + g < r(T_{R^x})$ must be equal to zero.

Moreover, in this case we have that $r(T_{RX}) = r(T_{R\psi}) = sp^{\min}$ hence we can factor out the term $\varepsilon^{sp^{\min}}$ from both products and since this holds for all the terms of the sum we get that:

$$\begin{aligned}
& \lim_{\varepsilon \rightarrow 0} \frac{\sum_{T_{RX}:r(T_{RX})=sp^{\min}} \prod_{(z' \rightarrow z'') \in T_{RX}} \Pr(x_{t+1} = z'' | x_t = z')}{\sum_{T_{R\psi}:r(T_{R\psi})=sp^{\min}} \prod_{(z' \rightarrow z'') \in T_{R\psi}} \Pr(x_{t+1} = z'' | x_t = z')} \\
&= \lim_{\varepsilon \rightarrow 0} \frac{\varepsilon^{sp^{\min}} \cdot \sum_{T_{RX}:r(T_{RX})=sp^{\min}} \left\{ \sum_{\{(g,h):(h-1)(\lambda-1)+g=sp^{\min}\}} c_{gh} \delta^{|A_N|-g} + \bar{K}_{RX}(\varepsilon) \right\}}{\varepsilon^{sp^{\min}} \cdot \sum_{T_{R\psi}:r(T_{R\psi})=sp^{\min}} \left\{ \sum_{\{(g,h):(h-1)(\lambda-1)+g=sp^{\min}\}} c'_{gh} \delta^{|A_N|-g} + \bar{K}_{R\psi}(\varepsilon) \right\}} \\
&= \frac{\sum_{T_{RX}:r(T_{RX})=sp^{\min}} \sum_{\{(g,h):(h-1)(\lambda-1)+g=sp^{\min}\}} c_{gh} \delta^{|A_N|-g}}{\sum_{T_{R\psi}:r(T_{R\psi})=sp^{\min}} \sum_{\{(g,h):(h-1)(\lambda-1)+g=sp^{\min}\}} c'_{gh} \delta^{|A_N|-g}}.
\end{aligned}$$

where by $\bar{K}_{RX}(\varepsilon)$ and $\bar{K}_{R\psi}(\varepsilon)$ we denote two polynomials of ε .

Let us now compare the expression in the numerator with that in the denominator. Observe first that any differences in the values of the two sums will come from edges which are not present in both sums. These edges must correspond to transitions in and out of the allocations of interest here, x and y , as any other transitions should be present in both sums. Hence, we need to compare the number of minimum resistance transitions out of y , denoted as $\#r_{out}^{\min}(y)$ and the minimum resistance transitions into x , denoted as $\#r_{in}^{\min}(x)$, which are present in the numerator sum, with the number of minimum resistance transitions out of x , denoted as $\#r_{out}^{\min}(x)$, and the minimum resistance transitions into y , denoted as $\#r_{in}^{\min}(y)$, which are present in the denominator sum. And note that it is only these minimum resistance transitions in and out of x and y that matter, as these sums correspond to minimum resistance trees. Given that the only way that $\mu^*(x) < \mu^*(y)$ is if there are more ways to get into y with resistance $r_{in}^{\min}(y)$ than into x with $r_{in}^{\min}(x)$ and also more ways to get out of x with $r_{out}^{\min}(x)$ than out of y with $r_{out}^{\min}(y)$, so that the sum in the numerator will be smaller than the sum in denominator. ■

Furthermore, a related fact that we will be using later on is the following.

Claim 4 (Equal Frequencies) *Let z and z' be two allocations that belong to the same recurrent class R^ω . Then it must be that $\mu^*(z) = \mu^*(z')$.*

Proof. From Lemma 2, we know that

$$\mu^\varepsilon(z) = \frac{q(z)}{\sum_{z' \in A_N} q(z')},$$

where $q(z) = \sum_{T_{R^\omega}: r(T_{R^\omega}) = sp^{\min}} \prod_{(z' \rightarrow z'') \in T_{R^\omega}} \Pr(x_{t+1} = z'' | x_t = z')$ and $T_{R^\omega} \ni z$. Thus, if $z, z' \in R^\omega$, then it must be that $q(z) = q(z')$ as $\tilde{r}(z, z') = 0$. Therefore, $\mu^\varepsilon(z) = \mu^\varepsilon(z'), \forall \varepsilon > 0$ and given that the perturbation is regular it must be that $\mu^*(z) = \mu^*(z')$. ■

However, before we present the proof of our main result (Proposition 1), we need to establish two auxiliary claims.

Claim 5 (Endowment Allocation) *The initial endowment allocation e is stochastically stable if and only if it is included in the core C .*

Proof. Let us begin by demonstrating that the endowment allocation cannot be part of any non-singleton recurrent class. To see why, assume for a moment that $\exists R^\omega$ with $|R^\omega| > 1$ and $e \in R^\omega$. However, given the nature of the cycle we must have that $\exists z \in R^\omega : r(z, e) = 0$, which means at least for some agent i $z_i \prec_i e_i$ so that $z \xrightarrow[0]{\{i\}} e$. Yet, this last statement requires $z \notin IR$ contradicting Lemma 1 of Serrano and Volij. Hence, if $e \in SSA \implies e \in C$. Let us now prove the converse.

Suppose that $e \in C$. Then by Proposition 2 we know that $r_{out}^{\min}(e) > 0$. Specifically we should have $r_{out}^{\min}(e) \geq 1$, as the only way for the process to get out of the singleton recurrent class $\{e\}$ is when at least one agent makes a mistake in the coalition formation process. If $\{e\}$ is the unique recurrent class of the process then it must be stochastically stable. Thus, let us assume now that there exists another recurrent class R^ω , pick an allocation $z \in R^\omega$ and consider the transition $z \rightarrow e$. Obviously $r(z, e) > 0$ since z is individually rational. Yet, $\exists i \in S_z^{\min} : z_i \sim_i e_i$. Otherwise we would have that $\forall i \in S_z^{\min} : z_i \succ_i e_i$ and the coalition

S_z^{\min} would block the endowment, contradicting the fact that $e \in C$. Therefore, it must be that $\forall R^\omega \neq \{e\} \wedge z \in R^\omega, r(z, e) = 1$. This means that $r_{in}^{\min}(e) = 1$, while $r_{out}^{\min}(z) \leq 1$, which using Claim 1 becomes $\tilde{r}_{out}^{\min}(R^\omega | \mathbf{R}^*) = 1$.

Finally, keeping in mind that for $R^\omega, R^\psi \in \mathbf{R}^* \tilde{r}(R^\psi, R^\omega) \geq 1$, we get that $r_{in}^{\min}(R^\omega | \mathbf{R}^*) \geq 1, \forall R \subseteq R^*$. Therefore, combining all the above facts we get that $\forall \mathbf{R} \subseteq \mathbf{R}^* : \{e\} \in \mathbf{R}$ we have $\tilde{r}_{out}^{\min}(R^\omega | \mathbf{R}) - r_{in}^{\min}(R^\omega | \mathbf{R}) \leq 0$, while for the endowment allocation we get that: $\tilde{r}_{out}^{\min}(\{e\} | \mathbf{R}) - r_{in}^{\min}(\{e\} | \mathbf{R}) \geq 0$. Hence, $e \in \mathbf{R}^{SS}$ and $e \in SSA$. ■

Claim 6 (Successful Coalitions) *If an allocation z is stochastically stable, then for all but one of the agents included in S_z^{\min} allocation z is strictly preferred to e .*

Proof. Suppose not. Then this means that $\exists i, j \in S_z^{\min} : z_i \succsim_i e_i \wedge z_j \succsim_j e_j$ and since S_z^{\min} is the smallest possible coalition that can achieve the direct transition $e \rightarrow z$ we should have that $r(e, z) \geq 2$. This is true also for any individually rational allocation $y_k \succsim_k e_k, \forall k \in N$ as for both $i, j \in S_z^{\min}$ we would have that $z_i \succsim_i y_i \wedge z_j \succsim_j y_j$. Hence, it must be that $r_{in}^{\min}(z) \geq 2$, a fact that precludes z from being included in any non-singleton recurrent class.

Now observe that the transition $z \rightarrow e$ can occur with at most one mistake whenever the singleton coalitions $\{i\}$ and $\{j\}$ form and $r(z, e) = 1$. Thus, if $e \in C$, this last statement is equivalent to $r_{out}^{\min}(z) = 1$, while if $e \notin C$, there must exist another core allocation x that blocks e in which case $r(e, x) = 0$. Hence, we have that $r_{out}^{\min}(z) \leq r(z, e) + r(e, x) = 1$.

Combining all the above facts we get that $\tilde{r}_{out}^{\min}(\{z\} | \mathbf{R}^*) - \tilde{r}_{in}^{\min}(\{z\} | \mathbf{R}^*) \leq -1$. Thus, $\{z\} \notin \mathbf{R}^{SS}$ meaning that z cannot be stochastically stable. ■

Now we can turn to the proof our main result.

Proposition 1 (Relative Frequencies) *Let z and z' be two house allocations that share the same degree of decentralization, while z is beneficial to more agents compared to z' , then we must have that $\mu^*(z) \geq \mu^*(z')$. Similarly, let z and z' be two house allocations where the number of agents benefiting is the same, but z is more decentralized than z' , then again we must have $\mu^*(z) \geq \mu^*(z')$.*

Proof. We restrict attention to the non-trivial case where $z \in R^\omega \wedge z' \in R^\psi$ as the statement holds trivially if $z, z' \in R^\omega$ by Claim 4. We consider two possibilities.

Possibility I: $z \in SSA$.

If $z \in SSA$ and $z' \notin SSA$, then we would obviously have $\mu^*(z) > \mu^*(z') = 0$. Thus, let us focus on the case where $z, z' \in SSA$. Suppose z and z' share the same degree of decentralization, namely $d(z) = d(z')$, but with allocation z being beneficial to more agents, namely $b(z) > b(z')$. Since z and z' are both stochastically stable, we should have by Proposition 6 that at most one agent in coalitions $S_z^{\min}, S_{z'}^{\min}$ be indifferent or worse off. Note that the latter is being precluded by the fact that both allocations are individually rational, which given that $|\{i \in N : z_i \succ_i e_i\}| > |\{i \in N : z'_i \succ_i e_i\}|$ leaves us only with the possibility: $z_i \succ_i e_i, \forall i \in S_z^{\min}$ while $z'_i \succ_i e_i, \forall i \in S_{z'}^{\min}/\{j\} \wedge z'_j \sim_j e_j$.

Similarly, suppose that for allocations z and z' we have $b(z) = b(z')$, but $d(z) > d(z')$. Again by Proposition 6 we should have, given that z and z' are both stochastically stable, that at most one agent in coalitions $S_z^{\min}, S_{z'}^{\min}$ be indifferent or worse off, with the latter being precluded by the fact that both allocations are individually rational. Yet, given our assumption of $|\{i \in N : z_i \succ_i e_i\}| = |\{i \in N : z'_i \succ_i e_i\}|$ the only possibility left here is $z_i \succ_i e_i, \forall i \in S_z^{\min}$ while $z'_i \succ_i e_i, \forall i \in S_{z'}^{\min}/\{j\} \wedge z'_j \sim_j e_j$.

Having made the above point let us now observe that if $z_i \succ_i e_i, \forall i \in S_z^{\min}$, then coalition S_z^{\min} would be able to improve upon the endowment allocation, namely $e \xrightarrow[S_z^{\min}]{0} z$ and therefore $e \notin C$. On the other hand, if $\exists j \in S_{z'}^{\min} : z'_j \sim_j e_j$, then agent j can force the economy back to the endowment allocation with just a small mistake, as $z' \xrightarrow[\{j\}]{1} e$. Combining the two facts we see that $\tilde{r}(z', z) = r(z', e) + r(e, z) = 1 + 0 = 1$, from which we can conclude that $\tilde{r}_{in}^{\min}(z) = 1$ and $\tilde{r}_{out}^{\min}(z') = 1$. Thus, given that both z and z' are stochastically stable, we can apply Corollary 1 to infer that $\tilde{r}_{in}^{\min}(z') = \tilde{r}_{out}^{\min}(z) = 1$.

Let us consider now a third recurrent allocation $s \in R^\chi$. For that allocation we claim that $s \longrightarrow z'$ is a minimum resistance transition into z' with resistance $\tilde{r}(s, z') = \tilde{r}_{in}^{\min}(z') = 1$, then this should also be the case for the transition $s \longrightarrow z$, whose resistance should be

$\tilde{r}(s, z) = \tilde{r}_{in}^{\min}(z) = 1$. Similarly, if $z \longrightarrow s$ is a minimum resistance transition out of z with resistance $\tilde{r}(z, s) = \tilde{r}_{out}^{\min}(z) = 1$, then this should also be the case for the transition $z' \longrightarrow s$, whose resistance should also be $\tilde{r}(z', s) = \tilde{r}_{out}^{\min}(z') = 1$. This is demonstrated in the following two claims.

Claim 7 *If $\tilde{r}(s, z') = \tilde{r}_{in}^{\min}(z')$, then $\tilde{r}(s, z) = \tilde{r}_{in}^{\min}(z)$.*

Proof. Suppose $s \xrightarrow[1]{S_{z'}^{\min}} z'$ and let us focus on the case of a direct transition, although the same argument applies for an indirect one. If there is one agent in the coalition $S_{z'}^{\min}$ that is indifferent between allocations s and z' , then this must be agent j for whom $z'_j \sim_j e_j$. If this was not the case and $z'_j \sim_j e_j \succ_j s_j$, then allocation s would be individually irrational and hence could not be part of any recurrent class. Given that, $s \xrightarrow[1]{\{j\}} e \xrightarrow[0]{S_z^{\min}} z$ would constitute a minimum resistance transition out of allocation s with a resistance $\tilde{r}(s, z) = r(s, e) + r(e, z) = 1 + 0 = 1 = \tilde{r}_{in}^{\min}(z)$. ■

Claim 8 *If $r(z, s) = \tilde{r}_{out}^{\min}(z)$, then $r(z', s) = \tilde{r}_{out}^{\min}(z')$.*

Proof. Suppose $z \xrightarrow[1]{S_s^{\min}} s$ and let us focus on the case of a direct transition, although the same argument applies for an indirect one. This implies that $s_i \succ_i z_i, \forall i \in S_s^{\min}/\{k\} \wedge s_k \sim_k z_k$.

Assume first that $k \in S_s^{\min} \cap S_z^{\min}$. This means that $s_i \succ_i e_i, \forall i \in S_s^{\min}/S_z^{\min}$, while $s_i \succsim_i z_i \succ_i e_i, \forall i \in S_s^{\min} \cap S_z^{\min}$. Hence, $s_i \succ_i e_i, \forall i \in S_s^{\min}$ and then coalition S_s^{\min} would block the endowment allocation. Having established that it is easy to see that the transition $z' \xrightarrow[1]{\{j\}} e \xrightarrow[0]{S_s^{\min}} s$ is a minimum resistance transition out of z' with a resistance of $\tilde{r}(z', s) = r(z', e) + r(e, s) = 1 + 0 = 1 = \tilde{r}_{out}^{\min}(z')$.

Alternatively let $k \in S_s^{\min}/S_z^{\min}$. This means that $s_i \succ_i z_i \succ_i e_i, \forall i \in S_s^{\min} \cap S_z^{\min}$, while $s_i \succ_i e_i, \forall i \in S_s^{\min}/S_z^{\min} \cup \{k\} \wedge s_k \sim_k e_k$. Consider now the transition $s \xrightarrow[1]{\{k\}} e \xrightarrow[0]{S_z^{\min}} z$ whose resistance is $\tilde{r}(s, z) = r(s, e) + r(e, z) = 1 + 0 = 1 = \tilde{r}_{out}^{\min}(s)$. This implies that allocation s should also be stochastically stable. Thus we can now apply Claim 7 for the pair of stochastically stable allocations $z, s \in SSA$ to obtain the result that if $\tilde{r}(z', z) = \tilde{r}_{in}^{\min}(z)$ then $\tilde{r}(z', s) = \tilde{r}_{in}^{\min}(s) = \tilde{r}_{out}^{\min}(z')$. ■

Let us now combine the two claims and note that from Claim 7 we can infer that whatever the number of minimum resistance paths into z' are, the number of minimum resistance paths into z should be not smaller, if not greater. Thus, we have $\#\tilde{r}_{in}^{\min}(z) \geq \#\tilde{r}_{in}^{\min}(z')$. Similarly from Claim 8 we can infer that whatever the number of minimum resistance paths out of z' are, the number of minimum resistance paths out of z should be not greater, if not smaller. Thus, again $\#\tilde{r}_{out}^{\min}(z) \leq \#\tilde{r}_{out}^{\min}(z')$. Having established that we can now invoke Claim 3 which allows us to state that $\mu^*(z) \geq \mu^*(z')$.

Possibility II: $z \notin SSA$.

If $z \notin SSA$ and $z' \in SSA$, this would mean that $\mu^*(z') > \mu^*(z) = 0$ contrary to the statement of the proposition. Hence, we need to show that if $z \notin SSA$ then $z' \notin SSA$ as well.

To begin with let us observe that both in the case where $d(z) = d(z')$ and $b(z) > b(z')$ as well as the case where $d(z) < d(z')$ and $b(z) = b(z')$ it must be that for allocation z' $|S_{z'}^{\min}| > |\{i \in S_{z'}^{\min} : z'_i \succ_i e_i\}|$ as $\exists j \in S_{z'}^{\min} : z'_j \sim_j e_j$. Hence, $r(z', e) = 1$, which implies that $\tilde{r}_{out}^{\min}(z') = 1$ as either $e \in C$ in which case $e \in SSA$, or e is blocked by some other core allocation c in which case $\tilde{r}(z', c) = r(z', e) + r(e, c) = 1$.

Hence, for $z' \in SSA$, it must be that $\tilde{r}_{in}^{\min}(z') = 1$, as otherwise $\tilde{r}_{out}^{\min}(z') - \tilde{r}_{in}^{\min}(z') < 0$. This requires that there exists at least some other recurrent individually rational allocation s for which $\tilde{r}(s, z') = 1$, which can only be true if $z'_i \succ_i s \succsim_i e_i, \forall i \in S_{z'}^{\min}$ and $z'_j \sim_j s \succsim_j e_j$.

Remember though that for allocation z we have $z_i \succ_i e_i, \forall i \in S_z^{\min}$. Thus, $\tilde{r}(z, z') = 1$, as there is always the transition $z' \xrightarrow[1]{\{j\}} e \xrightarrow[0]{S_z^{\min}} z$, and $\tilde{r}_{in}^{\min}(z) = 1$, which combined with Claim 2, yields that $z \in SSA$, a contradiction.

Finally, let us note that in the proof we have implicitly assumed that there are at least three distinct recurrent classes. This might raise the question of what happens in the case where the number of recurrent classes is just two. Yet, in such case it is easy to see that if $\mathbf{R}^* = \{R^\psi, R^\omega\}$ then $\tilde{r}_{out}^{\min}(R^\omega | \mathbf{R}^*) = \tilde{r}(R^\omega, R^\psi) = \tilde{r}_{in}^{\min}(R^\psi | \mathbf{R}^*)$ and $\tilde{r}_{out}^{\min}(R^\psi | \mathbf{R}^*) = \tilde{r}(R^\psi, R^\omega) = \tilde{r}_{in}^{\min}(R^\omega | \mathbf{R}^*)$. Given our assumptions it must be that $\tilde{r}_{out}^{\min}(R^\omega | \mathbf{R}^*) \geq \tilde{r}_{out}^{\min}(R^\psi | \mathbf{R}^*)$

and thus $\mu^*(z) \geq \mu^*(z')$. ■

6 Conclusions

The purpose of this research is to examine the emergence of cooperative allocations from decentralized exchange. Following Serrano and Volij (2008), we employ a dynamic trading process of coalitional recontracting in the housing economy, and characterize its stochastically stable distribution. The distribution provides us with a measure of the relative frequency with which each particular allocation is going to emerge in the long run. Utilizing this measure, we explain why some allocations are more likely to emerge compared to others, and then identify the underlying factors to determine the stochastic stability of any particular allocation.

Our analysis shows that there exist two main factors that affect the long-run outcome of the coalitional recontracting process. The first factor captures the incentives that each agent has to trade. The second factor is related to the degree of decentralization behind each allocation. Specifically, allocations that are beneficial for a larger fraction of agents and allocations that can be obtained in a more decentralized fashion are going to emerge at a higher frequency in the long run. Of course, the long-run outcome of the recontracting process is determined by the interplay of these two factors. Yet, the analysis of these two factors enables us to rationalize the reason why particular allocations are more likely to emerge as a result of trade or cooperation.

Furthermore, we should note that our methodology also has the potential of providing a new perspective in the literature on equilibrium refinements. This is because of our choice to focus on probabilistic assessments regarding the frequencies of emergence of various allocations included in the core, rather than aiming for a selection among those allocations. With that in mind, we believe that this research on the dynamics of cooperation will open up a whole new set of questions regarding how cooperative outcomes can emerge from repeated

interactions and hope that our contribution will encourage further work on the interesting new research topic.

References

- ABDULKADIROGLU, A., AND T. SONMEZ (1999): “House Allocation with Existing Tenants,” *Journal of Economic Theory*, 88(2), 233–260.
- ATHANASSOGLU, S., AND J. SETHURAMAN (forthcoming): “House Allocation with Fractional Endowments,” *International Journal of Game Theory*.
- BERGIN, J., AND B. L. LIPMAN (1996): “Evolution with State-Dependent Mutations,” *Econometrica*, 64(4), 943–956.
- BOGOMOLNAIA, A., AND H. MOULIN (2001): “A New Solution to the Random Assignment Problem,” *Journal of Economic Theory*, 100(2), 295–328.
- EDGEWORTH, F. Y. (1881): *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences*. C. Kegan Paul and Co., London.
- ELLISON, G. (2000): “Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution,” *Review of Economic Studies*, 67(1), 17–45.
- FELDMAN, A. M. (1974): “Recontracting Stability,” *Econometrica*, 42(1), 35–44.
- FOSTER, D., AND H. P. YOUNG (1989): “Stochastic Evolutionary Game Dynamics,” *Theoretical Population Biology*, 38, 219–232.
- FREIDLIN, M., AND A. WENTZELL (1984): *Random Perturbations of Dynamical Systems*. Springer Verlag, Berlin.
- GILLIES, D. B. (1953): “Some Theorems on N-Person Games,” Ph.D. thesis, Princeton University.
- GREEN, J. R. (1974): “The Stability of Edgeworth’s Contracting Process,” *Econometrica*, 42(1), 21–34.
- GREENWALD, A., AND J. WICKS (2005): “An Algorithm for Computing Stochastically Stable Distributions with Applications to Multiagent Learning in Repeated Games,” in *UAI ’05*, vol. Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence, pp. 623–632.
- KANDORI, M., G. J. MAILATH, AND R. ROB (1993): “Learning, Mutation, and Long Run Equilibria in Games,” *Econometrica*, 61(1), 29–56.
- KLAUS, B., O. BOCHET, AND M. WALZL (forthcoming): “A Dynamic Contracting Process for Multiple-Type Housing Markets,” *Journal of Mathematical Economics*.
- OSBORNE, M. J., AND A. RUBINSTEIN (1994): *A Course in Game Theory*. MIT Press, Cambridge, MA.
- SCARF, H. E. (1967): “The Core of an N Person Game,” *Econometrica*, 35(1), 50–69.

- SERRANO, R. (2009): “Cooperative Games: Core and Shapley Value,” in *Encyclopedia of Complexity and Systems Science*, ed. by R. Meyers. Springer, New York.
- SERRANO, R., AND O. VOLIJ (2008): “Mistakes in Cooperation: the Stochastic Stability of Edgeworth’s Recontracting,” *Economic Journal*, 118(532), 1719–1741.
- SHAPLEY, L. S. (1952): “Notes on the N-Person Game III: Some Variants of the von-Neumann-Morgenstern Definition of Solution,” *Research Memorandum, Rand Corporation*, (RM - 817).
- SHAPLEY, L. S., AND H. E. SCARF (1974): “On Cores and Indivisibility,” *Journal of Mathematical Economics*, 1(1), 23–37.
- SHUBIK, M. (1959): “Edgeworth Market Games,” in *Contributions to the Theory of Games*, ed. by A. W. Tucker, and R. D. Luce, vol. IV of *Annals of Mathematics Studies*, pp. 267–278. Princeton University Press, Princeton, NJ.
- YOUNG, P. H. (1993): “The Evolution of Conventions,” *Econometrica*, 61(1), 57–84.