

**Gale-Shapley Matching in an
Evolutionary Trade Network Game**

ISU Economic Report 43

April 1998

by

Leigh Tesfatsion

Department of Economics

375 Heady Hall

Iowa State University, Ames, IA 50011-1070

<http://www.econ.iastate.edu/tesfatsi/>

E-mail: tesfatsi@iastate.edu

ABSTRACT

This study investigates the performance of Gale-Shapley matching in an evolutionary market context. Computational experimental findings are reported for an evolutionary match-and-play trade network game in which resource-constrained traders repeatedly choose and refuse trade partners in accordance with Gale-Shapley matching, participate in risky trades modelled as two-person prisoner's dilemma games, and evolve their trade behavior over time. Particular attention is focused on correlations between ex ante market structure and the formation of trade networks, and between trade network formation and the types of trade behavior and social welfare outcomes that these trade networks support.

Keywords: Evolutionary game; Gale-Shapley matching; iterated prisoner's dilemma; trade networks; endogenous interactions; agent-based computational economics.

JEL Classification Numbers: C7, C8, C9, D8, D4.

Copyright © 1998 by Leigh Tesfatsion. All rights reserved. Readers may make verbatim copies of this document for non-commercial purposes by any means, provided this copyright notice appears on all such copies.

IOWA STATE IS AN EQUAL OPPORTUNITY EMPLOYER

1 Introduction

An *evolutionary match-and-play game* is an evolutionary game in which the concept of rational play is extended to include the choice and refusal of partners as well as the choice of strategy to play with any given partner. Such games have previously been studied by Stanley et al. (1994), Ashlock et al. (1996), Hauk (1996), and Tesfatsion (1997a). Other game theory studies that have allowed players to avoid unwanted interactions, or more generally to affect the probability of interaction with other players through their own actions, include Fogel (1995), Hirshleifer and Rasmusen (1989), Kitcher (1993), Mailath et al. (1994), and Orbell and Dawes (1993). An extensive review of this work is given in Hauk (1996).

In addition, a growing number of economists have recently begun to explore multi-agent endogenous interaction systems in which the decision (or state) of an agent depends on the decision (or state) of certain neighboring agents, where these neighbors may change over time. See, for example, Brock and Durlauf (1995), De Vany (1996), Ellison (1992), Epstein and Axtell (1996), Guriev and Shakhova (1996), Kirman (1997), Ioannides (1996), Vriend (1996), Weisbuch et al. (1997), and Young (1993).

The model developed in this study builds on the evolutionary match-and-play trade network game (TNG) developed by Tesfatsion (1997a) for studying the formation and evolution of trade networks under alternatively specified market structures. The primary objective of this study is to investigate the performance of the well-known matching mechanism developed by Gale and Shapley (1962) when implemented in an evolutionary market context.

Three types of market structures are considered: two-sided markets comprising pure buyers and pure sellers; partially fluid markets comprising pure buyers, pure sellers, and buyer-sellers capable of functioning as both buyers and sellers; and endogenous-type markets comprised entirely of buyer-sellers. Given each market structure, buyers and sellers repeatedly choose and refuse trade partners using Gale-Shapley matching, engage in risky trades modelled as two-person prisoner's dilemma games, and evolve their trade behavior over time. The model is implemented computationally by means of the TNG source code developed by McFadzean and Tesfatsion (1997), which in turn is supported by SimBioSys,

a general C++ class library for evolutionary simulations developed by McFadzean (1995).¹

To aid in this investigation, various practical and informative descriptive statistics are constructed for measuring experimentally observed correlations between exogenously given structural characteristics and the formation and evolution of trade networks, and between trade network formation and the types of trade behavior and social welfare outcomes that these trade networks support. These descriptive statistics complement and extend the descriptive statistics developed by Stanley et al. (1994), Smucker et al. (1994), and Ashlock et al. (1996) to characterize play behavior and significant play graphs.

The main conclusion drawn from this study is that the optimality criteria conventionally used to evaluate the performance of matching mechanisms in static market contexts turn out to be highly incomplete indicators of performance from an evolutionary vantage point. The static viewpoint hides the strong role played by market structure and ex ante capacity constraints in determining the types of persistent matching networks that evolve, the types of persistent interaction behaviors that these networks support, and the transactions costs and inactivity costs to agents that the achievement of these persistent networks and behaviors entails. In addition, the static viewpoint takes preference rankings over potential partners as given whereas these rankings are continually updated on the basis of past interactions in evolutionary settings. Indeed, matching behaviors and interaction behaviors evolve conjointly. This suggests the need for more comprehensive optimality criteria that take both facets into account.

More concretely, in the market experiments reported below, buyers and sellers determine their trade partners using “deferred choice and refusal” (DCR). The DCR mechanism is a version of Gale-Shapley matching that has been suitably generalized to handle two-sided, partially-fluid, and endogenous-type market structures in which buyers have arbitrary quotas on the number of trade offers they can make and sellers have arbitrary quotas on the number of trade offers they can accept. The matching outcomes generated by the DCR mechanism have been shown [Tsfatsion (1997a)] to have the usual optimality properties associated with Gale-Shapley matching: namely, pairwise stability; and Pareto optimality from the vantage

¹Source code for the TNG and for SimBioSys are both available as freeware at the author’s Web site, along with a variety of resource materials for agent-based computational economics in general. All experiments reported in the current study were run using version 105b of the TNG source code.

point of the agents who actively make offers, who in the current context are the buyers.

The evolutionary outcomes observed in these market experiments, however, include autarkic economies in which all traders are persistent wallflowers, exploitive economies in which buyers persistently defect against cooperative sellers or sellers persistently defect against cooperative buyers, and harmonious economies in which all traders are persistent cooperators. Moreover, due to transactions costs and inactivity costs, social welfare can still be low even if all active traders are persistent cooperators. These evolutionary outcomes are systematically related to market structure and to ex ante capacity constraints as captured by the buyer offer quotas and seller acceptance quotas.

Section 2 presents the basic model. Section 3 constructs an ex ante measure for excess capacity and ex post measures for the classification of trade networks, trade behaviors, and social welfare outcomes. The experimental design is presented in Section 4, and a detailed discussion of experimental findings is given in Section 5. Concluding remarks are given in Section 6.

2 The Basic Model

This section gives a brief overview of the Trade Network Game (TNG) together with the particular TNG module specifications used for the study at hand. A detailed discussion of these module specifications, together with their C++ implementation, can be found in McFadzean and Tesfatsion (1997).

The TNG consists of a collection of traders that evolves over time. As depicted in Table I, each trader in the initial generation is constructed and assigned a random trade strategy. The traders then enter into a nested pair of generation cycle and trade cycle loops during which they repeatedly determine trade partners, carry out trades, update their expectations, and evolve their trade strategies over time.

In the current study, alternative market structures are imposed through the pre-specification of buyers and sellers and through the pre-specification of quotas on trade offer submissions and acceptances. More precisely, the set of traders is taken to be the union $V = B \cup S$ of a nonempty subset B of *buyers* who can submit trade offers and a nonempty subset S of *sellers* who can receive trade offers, where B and S may be disjoint, overlapping, or coincident. A

```

int main () {
    Init();                                // Construct the initial trader generation
                                           // with random trade strategies.
    For (G = 1,...,GMax) {                 // Enter the generation cycle loop.
                                           // Generation Cycle:
        InitGen();                         // Configure traders with user-supplied
                                           // parameter values (initial expected
                                           // utility levels, quotas,...).
        For (I = 1,...,IMax) {              // Enter the trade cycle loop.
                                           // Trade Cycle:
            MatchTraders();                 // Determine trade partners,
                                           // given expected utilities,
                                           // and record refusal and
                                           // wallflower payoffs.
            Trade();                       // Implement trades and
                                           // record trade payoffs.
            UpdateExp();                   // Update expected utilities
                                           // using newly recorded payoffs.
        }                                 // Environmental Step:
        AssessFitness();                  // Assess trader fitness scores.
        Output();                         // Output trader information.
                                           // Evolution Step:
        EvolveGen();                      // Evolve a new trader generation.
    }
    Return 0;
}

```

Table I: Pseudo-Code for the TNG

trader is classified as a pure buyer, a pure seller, or a buyer-seller if he is an element of V/S , V/B , or $B \cap S$, respectively. In each trade cycle, each buyer can have no more than bq trade offers outstanding to sellers at any given time, and each seller can accept no more than sq trade offers from buyers, where the buyer offer quota bq and the seller acceptance quota sq can be any positive integers. Although highly simplified, these parametric specifications permit the study of a variety of market structures operating under different ex ante capacity constraints.

Matches between buyers and sellers in each trade cycle are determined using a modified version of the well-known Gale-Shapley (1962) deferred acceptance mechanism, hereafter referred to as the *deferred choice and refusal (DCR)* mechanism.² Under the DCR mechanism, each buyer submits up to bq trade offers to sellers he ranks as most preferable on the basis of expected utility and who he judges to be tolerable in the sense that their expected utility is not negative,³ with at most one trade offer going to any one seller. Similarly, each

²See Roth and Sotomayor (1990) for a careful detailed discussion of the basic properties of Gale-Shapley matching mechanisms.

³No buyer-seller is allowed to match with himself. This is handled by having each buyer-seller assign a

seller selects up to sq of his received offers that he finds tolerable and most preferable on the basis of expected utility and he places them on a waiting list; all other offers are refused. Buyers redirect refused offers to tolerable preferred sellers who have not yet refused them, if any such sellers exist, and the process repeats. Once sellers stop receiving new offers, they accept all trade offers currently on their waiting lists. A buyer incurs a transactions cost in the form of a negative *refusal payoff* R whenever a seller refuses one of his offers during the matching process; the seller who does the refusing is not penalized. A trader who neither submits nor accepts trade offers during the matching process receives a *wallflower payoff* 0 . These refusal and wallflower payoffs are assumed to be measured in utility terms.

If a seller accepts a trade offer from a buyer in some given trade cycle, the buyer and seller are said to be *matched* for that trade cycle. Each such match constitutes a mutually agreed upon contract stating that the buyer and seller shall engage in one trade with each other during the course of the trade cycle.⁴ This trade is risky in that outcomes are not assured.

Specifically, each trade between a buyer and seller is modelled as a two-person prisoner's dilemma game with payoffs normalized about zero. The buyer can either cooperate (fulfill his contractual obligations) or defect (fail to fulfill these obligations), and similarly for the seller. For example, the buyer could be a worker who either works productively or shirks on the job, and the seller could be an employer who provides a work environment that is either satisfactory or substandard. The range of possible utility payoffs is the same for each trade in each trade cycle: namely, L (the sucker payoff) is the lowest possible payoff, received by a cooperator whose trade partner defects; D is the payoff received by a defector whose trade partner also defects; C is the payoff received by a cooperator whose trade partner also cooperates; and H (the temptation payoff) is the highest possible payoff, received by a defector whose trade partner cooperates. More precisely, the utility payoffs are assumed to satisfy $L < D < 0 < C < H$, with $(L + H)/2 < C$.

The trade behavior of each trader in repeated trades with other traders is represented as a finite-memory pure strategy for playing a prisoner's dilemma game with an arbitrary partner

negative initial expected utility to himself, so that he judges himself to be an "intolerable" trade partner.

⁴Note this permits a pair of buyer-seller traders v and z to engage in up to two trades with each other during a trade cycle: namely, one trade in which v acts as buyer and z as seller, and a second trade in which z acts as buyer and v as seller.

an indefinite number of times, hereafter referred to as a *trade strategy*. Each trader thus has a distinct trading personality even if he engages in both buying and selling activities. At the commencement of each trade cycle loop, traders have no information about the trade strategies of other traders; they can only learn about these strategies by engaging other traders in repeated trades and observing the actions and utility payoffs that ensue. Moreover, each trader's choice of an action in a current trade with a trade partner is determined entirely on the basis of his trade history with this partner. Thus, each trader keeps separate track of the particular state he is in with regard to each of his potential trade partners.

During the course of each trade cycle loop, traders use a simple criterion filter⁵ to update their expected utility assessments for their potential trade partners as new utility payoffs are received. Each trader v starts by assigning the same exogenously given initial expected utility, U^o , to each potential trade partner z with whom he has not yet interacted. Once interactions with z take place, v calculates his current updated expected utility assessment for z by forming the average of U^o plus all utility payoffs he has received to date from interactions with z during the course of the current trade cycle loop, including both refusal payoffs and payoffs received in trades.

Because trade strategies in the TNG are implemented as finite state machines with finite memory depth, the actions undertaken by any two traders v and z in repeated trades must eventually cycle. Consequently, as the number of trades between v and z increases, and assuming neither trader refuses the trade offers of the other, the criterion filter guarantees that the expected utility that v associates with z approaches the true average utility that v would attain from infinitely repeated trades with z . On the other hand, unless offset by positive trade payoffs, repeated refusal payoffs eventually lead to a permanent cessation of trade as the expected utility associated with the trader doing the refusing drops below the minimum tolerance level, 0.

At the end of each trade cycle loop, the *fitness score* of each trader is calculated to be the average per-payoff utility level that he achieved during the course of the loop. The trade strategies of pure buyers, pure sellers, and buyer-sellers are then separately evolved

⁵As detailed in Tesfatsion (1979), a criterion filter is a learning algorithm that provides for the direct updating of an expected return function on the basis of new return outcomes without recourse to the usual interim updating of probability assessments via Bayes' Rule.

by means of standardly specified genetic algorithms involving elitism, mutation, and two-point cross-over operations. This evolution step is meant to reflect the public⁶ formation and transmission of new ideas rather than biological reproduction. Specifically, if a trade strategy successfully results in a high fitness score for a trader of a particular type, other traders of the same type are led to modify their own strategies to more closely mimic the successful strategy. Each modified strategy is also subjected to a small degree of mutation, thus permitting the traders to engage in creative experimentation as well as mimicry.

After the evolution step, the memories of the evolved traders are wiped clean with respect to their matching and trading experiences in the preceding trade cycle loop. In particular, each trader once again assigns an expected utility of U^o to each potential trade partner. The next trade cycle loop then commences, and the process repeats.

3 Descriptive Statistics

In this section, care is taken to explain the ex ante and ex post measures that have been constructed to aid in the experimental determination of correlations between ex ante market structure and trade network formation, and between trade network formation and the types of trade behavior and social welfare outcomes that these trade networks support. Trade networks depict who is trading with whom, and with what regularity. Trade behavior refers to the specific actions undertaken by a trader in trades with any given trade partner. Finally, social welfare measures the overall utility achieved by the traders from repeated trade interactions within the context of a possibly changing network of trade partners.

Let s denote the seed value for the initialization of the TNG random number generator, and let e denote a *potential economy*, i.e., an economy characterized structurally by the TNG source code together with all of the user-specified TNG parameter values apart from s . The *realized economy* generated from e , given the seed value s , is denoted by (s, e) . In the subsections below, it is first explained how capacity is measured ex ante for any potential

⁶Since the traders during the course of each trade cycle loop do not know the trade strategies of other traders, the information transmission during the evolution step must be the result of some form of public education rather than private introspection. This artificial distinction between learning about trade partners through private experience and evolving of trade strategies through public education, helpful in preliminary experimental stages for enabling some analytical treatment of model properties (e.g., criterion filter performance), should ultimately be replaced by a more realistic and seamless model of learning encompassing both matching behavior and trade strategies.

economy e . A distance measure is then constructed to differentiate among observed trade network formations for realized economies (s, e) . Finally, descriptive statistics are developed to summarize and differentiate among observed types of trade behaviors and social welfare outcomes for realized economies (s, e) .

3.1 Ex Ante Characterization of Excess Capacity

Given any potential economy e , let $N(e)$ denote the total number of sellers and let $M(e)$ denote the total number of buyers. By assumption, each seller in e has the same acceptance quota $sq(e)$ and each buyer in e has the same offer quota $bq(e)$, where $sq(e)$ and $bq(e)$ can be any positive integers. Consequently, the maximum total number of trade offers that sellers can feasibly accept from buyers during the course of a trade cycle is $N(e)sq(e)$, and the maximum total number of trade offers that buyers can feasibly have outstanding (i.e., under consideration by sellers or already accepted by sellers) at any point in time during a trade cycle is $M(e)bq(e)$.

The excess capacity of e is then measured ex ante by the extent to which potential seller acceptances outweigh potential buyer offers. More precisely, the *(relative) excess capacity* of e is defined to be

$$EC(e) = \frac{[N(e)sq(e) - M(e)bq(e)]}{M(e)bq(e)} . \quad (1)$$

If the number of buyers equals the number of sellers, this measure reduces to the following simple function of the seller acceptance quota and buyer offer quota:

$$EC(e) = \frac{[sq(e) - bq(e)]}{bq(e)} . \quad (2)$$

Moreover, if the seller acceptance quota $sq(e)$ is assumed without loss of generality to be no greater than $M(e)bq(e)$, which will be true for all experiments reported below in Section 5, then $-1 < EC(e) \leq [M(e) - 1]$. In this case, an economy e will be said to be characterized by *high excess capacity* if $EC(e)$ is close to $[M(e) - 1]$, *zero excess capacity* if $EC(e) = 0$, *tight capacity* if $EC(e) < 0$, and *extremely tight capacity* if $EC(e)$ is close to -1.

3.2 Ex Post Classification of Trade Networks by Distance

As explained in Section 2, the actions undertaken by any trader v in repeated trades with another trader z must eventually cycle. Consequently, these actions can be summarized in

the form of a *trade history* $H:P$, where the *handshake* H is a (possibly null) string of trade actions that form a non-repeated pattern and the *persistent portion* P is a (possibly null) string of trade actions that are cyclically repeated. For example, letting c denote cooperation and d denote defection, the trade history $c:c$ indicates that v cooperated in his first trade with z and continued to cooperate with z in all subsequent trades, whereas $ddd:dc$ indicates that v defected against z in his first three trades with z and thereafter alternated between defection and cooperation.

The information that is lost in this summary representation of the trade history of v with z is the exact length of the trade history. By convention, to appear in P , an action pattern must be repeated at least twice; otherwise it is assumed to form part of H . Also by convention, if any trades take place between v and z at all, then H is always taken to include the initial action of v in these trades and so will not be null. The length and member elements of P are then uniquely determined, and it is this unique determination that is used below to characterize trade network formation and trade behaviors. Note, however, that the length and member elements of H are not uniquely determined. For example, the trade history $ccddcddcddcdd$ has three possible $H:P$ forms, one for each of the three possible permutations of cdd as the representation for the persistent portion P : namely, $c:cdd$, $cc:ddc$, and $ccd:dcd$.

Two traders v and z are said to exhibit a *persistent trade relationship* during a given trade cycle loop T of a realized economy (s,e) if the following two conditions hold: (a) their trade histories with each other during the course of T take the form $H_v:P_v$ and $H_z:P_z$ with nonnull P_v and P_z ; and (b) accepted trade offers between v and z do not permanently cease during T either by choice (a permanent switch away to strictly preferred trade partners) or by refusal (one trader becoming intolerable to the other because his expected utility drops below zero).

A possible pattern of trade relationships among the traders $V(e)$ in the final generation of a potential economy e is referred to as a *trade network*, denoted generically by $K(e)$. Each trade network $K(e)$ is represented in the form of a directed graph in which the nodes of the graph represent the traders $V(e)$, the edges of the graph (directed arrows) represent trade offers directed from buyers to sellers, and the edge weight on any edge denotes the number of accepted trade offers between the buyer and seller connected by the edge.

Let $V^o(e)$ denote a *base trade pattern* that partially or fully specifies a potential pattern

of trade relationships among the traders $V(e)$ in the final generation of a potential economy e . For example, $V^o(e)$ could designate that each buyer directs trade offers to at least two sellers. Let $K^o(e)$ denote the *base trade network class* consisting of all trade networks $K(e)$ whose edges conform to the base trade pattern $V^o(e)$.

Finally, let $K(s, e)$ denote the trade network depicting the actual pattern of trade relationships among the traders $V(e)$ in the final generation of the realized economy (s, e) . The reduced form trade network $K^p(s, e)$ derived from $K(s, e)$ by setting to zero all edge weights of $K(s, e)$ that correspond to non-persistent trade relationships is referred to as the *persistent trade network* for (s, e) . An edge of $K^p(s, e)$ connecting a buyer-seller pair (v, z) is depicted as a straight-line or wavy-line directed arrow from v to z depending on whether v 's submission of trade offers to z are ultimately made continuously (in each successive trade cycle) or intermittently (randomly or recurrently across trade cycles), respectively. In the former case, v is said to be *latched* to z .

The *distance* $D^o(s, e)$ between the persistent trade network $K^p(s, e)$ and the base trade network class $K^o(e)$ for a realized economy (s, e) is then defined to be the number of nodes (traders) in $K^p(s, e)$ whose arrow patterns (persistent trade relationships) fail to conform to the base trade pattern $V^o(e)$. By construction, the distance $D^o(s, e)$ is bounded below by 0 and bounded above by the total number of traders in $V(e)$. This distance measure provides a rough way to classify the different types of persistent trade networks observed to arise for a given value of e as the seed value s is varied.

3.3 Ex Post Classification of Behaviors and Welfare Outcomes

A trader v in a realized economy (s, e) is referred to as an *unprovoked defector (UD)* if he engages in at least one defection against a trader partner who has not previously defected against him. The vector giving the separate UD percentages for pure buyers, pure sellers, and buyer-sellers in the final generation of (s, e) is referred to as the *UD profile* for (s, e) . The UD profile measures the extent to which the different types of traders behave aggressively in trades with trade partners who are either strangers or who so far have been consistently cooperative.

Also, v is referred to as a *persistent wallflower (PW)* if v constitutes an isolated node of the persistent trade network $K^p(s, e)$. Alternatively, v is referred to as a *repeat defector (RD)*

if v establishes at least one persistent trade relationship for which the persistent portion P of his trade history $H:P$ includes a defection d . If, instead, v establishes at least one persistent trade relationship and his trade history for each of his persistent trade relationships has the general form $H:c$, he is referred to as a *persistent cooperator (PC)*. By construction, each trader in a realized economy (s, e) who is not a persistent wallflower must either be a repeat defector or a persistent cooperator.

The vectors giving the separate PW, RD, and PC percentages for pure buyers, pure sellers, and buyer-sellers in the final generation of (s, e) are referred to as the *PW profile*, the *RD profile*, and the *PC profile* for (s, e) , respectively. The PW profile measures the extent to which the different types of traders fail to establish any persistent trade relationships. In contrast, the RD and PC profiles measure the extent to which the different types of traders establish persistent trade relationships characterized by predacious or fully cooperative behavior, respectively.

The vector giving the separate mean average fitness scores for pure buyers, pure sellers, and buyer-sellers in the final generation of a realized economy (s, e) is referred to as the *FIT profile* for (s, e) . The FIT profile constitutes a measure of social welfare.

4 Experimental Design

The computer experiments reported in Section 5 focus on three simple market structures: endogenous-type markets comprising 24 buyer-sellers (MS1); two-sided markets comprising 12 pure buyers and 12 pure sellers (MS2); and partially fluid markets comprising 8 pure buyers, 8 pure sellers, and 8 buyer-sellers (MS3). Within each market structure, four different configurations for the seller acceptance quota sq and buyer offer quota bq are examined: high excess capacity ($sq \gg bq$); zero excess capacity ($sq = bq = 1$); tight capacity ($sq = 1$ and $bq = 2$); and extremely tight capacity ($sq \ll bq$). The genetic algorithm elite value is automatically adjusted in each experiment to maintain the elite proportion at approximately two thirds for each nonzero trader type.⁷

⁷As detailed in McFadzean and Tesfatsion (1997), the elite value specifies how many of the most fit trade strategies in use by each type of trader in a current trader generation are inherited without modification by this same type of trader in the next trader generation. The elite value is set at 16 for endogenous-type markets, at 8 for two-sided markets, and at 6 for partially-fluid markets.

```

// PARAMETER VALUES FIXED ACROSS REPORTED EXPERIMENTS
GMax = 50           // Total number of generations.
IMax = 150          // Number of trade cycles in each trade cycle loop.
MutationRate = .005 // GA bit toggle probability.
FsmStates = 16       // Number of internal FSM states.
FsmMemory = 1        // FSM memory (in bits) allocated to past move recall.
RefusalPayoff = -0.5 // Payoff R received by a refused trader.
WallflowerPayoff = 0.0 // Payoff received by an inactive trader.
Sucker = -1.6        // Lowest possible trade payoff, L.
BothDefect = -0.6     // Mutual defection trade payoff, D.
BothCoop = +1.4       // Mutual cooperation trade payoff, C.
Temptation = +3.4     // Highest possible trade payoff, H.
InitExpPayoff = +1.4  // Initial expected utility level,  $U^o$ .
TraderCount = 24      // Total number of buyers and sellers.
// PARAMETER VALUES VARIED ACROSS REPORTED EXPERIMENTS
PureBuyers = 12       // Number of pure buyers.
PureSellers = 12       // Number of pure sellers.
BuyerSellers = 0       // Number of buyer-sellers.
Elite = 8              // Number of elite for each nonzero trader type.
BuyerQuota = 1         // Buyer offer quota bq.
SellerQuota = 12       // Seller acceptance quota sq.

```

Table II: Parameter Values for a Two-Sided Market with High Excess Capacity

The values for all remaining parameters are maintained at fixed values throughout all experiments apart from stability checks (see below). Table II lists these fixed parameter values along with the specific trader type values, elite value, and quota values for a two-sided market experiment with high excess capacity. The parameter values in Table II, together with the TNG source code, constitute a potential economy e in the sense defined in Section 3. Without loss of generality, such e will be represented in Section 5 as a function $e(\text{MS}, bq, sq)$ of market structure MS and the quota values bq and sq only, with dependence on all fixed parameter values and the automatically adjusted elite value suppressed. For example, the potential economy e corresponding to Table II will be referred to as $e(\text{MS2}, 1, 12)$.

As indicated in Table II, the number of generations was set at 50 for each tested potential economy e . Twenty realized economies (s, e) were experimentally generated for each e using twenty arbitrarily selected seed values s for the TNG pseudo-random number generator.⁸ Detailed information about the final (fiftieth) generation of traders in (s, e) was recorded for each run s . Using this information, the persistent trade network $K^p(s, e)$ was graphically depicted and the UD, PW, RD, PC, and FIT profiles were determined and recorded for each

⁸These twenty seed values are as follows: 5, 10, 15, 20, 25, 30, 45, 65, 63, 31, 11, 64, 41, 66, 13, 54, 641, 413, 425, and 212. The final fourteen values were determined by random throws of two and three die.

run s . By construction, for each (s, e) , the RD profile can be derived by subtracting the sum of the PW and PC profiles from the exogenously specified profile of trader types for the potential economy e . Consequently, only the UD, PW, PC, and FIT profiles are reported in Section 5.

A base trade pattern $V^o(e)$ was then specified for each tested potential economy e . Although the choice of this base trade pattern is simply a normalization determining a 0 point for the distance measure D^o , and hence intrinsically arbitrary, the degree of specificity of this pattern governs the dispersion of the resulting distance values and the extent to which these distance values display useful correlations with trade behaviors and social welfare outcomes as measured by the UD, PW, PC, and FIT profiles. For example, a base trade pattern specifying vacuously that each buyer is either a wallflower or directs trade offers to at least one seller results in all trade networks having a distance value zero whereas a highly specific base trade pattern could result in all trade networks having a distance value 24.

In practice, then, the choice of the base trade pattern $V^o(e)$ for each tested potential economy e was fine-tuned so that the resulting distance values provided a meaningful informative classification of trade network types. Given $V^o(e)$, the distance $D^o(s, e)$ of the persistent trade network $K^p(s, e)$ from the base trade network class $K^o(e)$ was recorded for each run s , and a histogram for the distance values $D^o(s, e)$ was constructed giving the percentage of runs s corresponding to each possible distance value.

Finally, as a rough stability check, the number of generations was also increased to 100 for each tested potential economy e and the minimum, maximum, and average fitness scores for the traders in each of the 100 generations were graphically generated for each realized economy (s, e) . In most cases, these fitness score values exhibited remarkable stability over generations 25 through 100. Cases in which instabilities were detected are noted in Section 5. In addition, as a rough check on small sample size, the number of traders was increased from 24 to 36 with proportionate increases in the number of traders of each type for each market structure, e.g., an increase from 12 to 18 pure sellers and from 12 to 18 pure buyers for two-sided market experiments. For each experiment conducted with 24 traders, six economies were also generated and examined with 36 traders, keeping all other parameter values and the base trade pattern the same. In each case, the resulting pattern of distance values and associated UD, PW, PC, and FIT profiles closely resembled those found for the smaller

population size.

5 Experimental Findings

5.1 Endogenous-Type Market Experiments

Consider, first, an endogenous-type potential economy $e = e(\text{MS1}, 1, 24)$ comprising 24 buyer-seller traders, each having an offer quota of $bq = 1$ and an acceptance quota of $sq = 24$. These quota values imply that the maximum number of trade offers that sellers can accept in any given trade cycle greatly exceeds the maximum number of trade offers that buyers can make. Consequently, this e is characterized by high excess capacity $EC(e)$ in the sense defined in Section 3.1. Moreover, since each seller is individually capable of accepting all of the trade offers that buyers as a whole can make, the buyers face a zero structural risk of having their trade offers refused on the basis of limited acceptance capacity.

As depicted in Figure 1(a), the base trade pattern $V^o(e)$ for this economy is as follows: Each buyer-seller trader directs trade offers to other buyer-seller traders without latching. This base trade pattern implies that no trader functions as a pure seller, no trader directs his trade offers exclusively and continuously to only one trade partner, and no trader is a wallflower.

— Insert Figure 1 About Here —

As seen in Table III, 90% of the twenty realized economies (s, e) experimentally generated for this e were observed to lie in the distance cluster 0–3. This means that, for each such realized economy, at most three of the twenty-four traders in the final (fiftieth) trader generation deviate from the base trade pattern. The mean UD profile for this distance cluster is 3%, meaning that the average percentage of traders exhibiting UD behavior in the eighteen realized economies lying in distance cluster 0–3 is 3%. Similarly, the mean PW profile is 1%, the mean PC profile is 96%, and the mean FIT value is 1.36.⁹ Two outlier cases were also

⁹In Table III, and in all subsequent tables below, the standard deviations for the UD, PW, and PC profiles (measured in percentages) appear in parentheses beneath the mean values for these profiles and are rounded off to the nearest integer value. Also, the standard deviations for the FIT profiles appear in parentheses below the mean FIT profiles and are rounded off to two decimal places. The calculation of these standard deviations is not applicable (NA) for distance clusters encompassing only one run, i.e., for distance clusters encompassing only 5% of the total sample of twenty realized economies.

observed at distance values 11 and 23. Both outliers exhibit more frequent UD behavior and less frequent PC behavior, and one has a substantially lower social welfare outcome as well.

D° Cluster	% Runs	Mean UD	Mean PW	Mean PC	Mean FIT
0–3	90%	3% (3%)	1% (2%)	96% (4%)	1.36 (.05)
11	5%	8% (NA)	0% (NA)	38% (NA)	1.34 (NA)
23	5%	96% (NA)	0% (NA)	88% (NA)	1.11 (NA)

Table III: Endogenous-Type Market Experiments with High Excess Capacity

A rough stability check was conducted for each of the twenty realized economies (s, e) for this high excess capacity economy e to check whether the information recorded in Table III for the final (fiftieth) generations appeared to be informative for other generations as well. Specifically, holding all other parameter values and the seed value s fixed, the number of generations was increased to 100 and the minimum, maximum, and average fitness scores attained by the traders in each of these 100 generations were recorded and graphically printed out.

Figure 2 depicts the stability results obtained for the realized economy $(413, e)$ with distance value 0; these results are typical of the stability results obtained for all economies in distance cluster 0–3. The average fitness score for traders fluctuates closely around the mutual cooperation payoff level 1.40 over generations 10 through 100. The occasional downward spikes in the minimum fitness score correspond to generations in which one or more mutant UD traders induce retaliatory RD and refusal behavior in other traders which results in below-average fitness scores for the mutants. The single upward spike in the maximum fitness score corresponds to a generation in which a mutant UD trader successfully manages to establish a parasitical RD trade relationship with a less predacious trade partner; note the dip in the minimum fitness score mirroring the upward spike in the maximum fitness score.

— Insert Figure 2 About Here —

The dispersion in the distance values $D^o(s, e)$ in Table III for the cross-sectional sampling of twenty realized economies (s, e) arises for exactly the same reason that episodes of upward and downward spiking arise in the time series fitness data depicted in Figure 2: namely, the presence of mutant UD traders. Moreover, the tight clustering of most of the distance values about 0 in Table III corresponds to the rarity and brevity of the spiking episodes in Figure 2, and the high percentage of PC behavior seen in Table III corresponds to the fact that the average fitness score in Figure 2 tends to fluctuate closely around the mutual cooperation payoff level 1.40.

All of these observations would appear to have a simple structural explanation. In endogenous-type economies, all traders evolve together in the evolution step; hence any trade strategies garnering below-average fitness scores are soon eliminated and replaced with variants of more successful strategies. Consequently, there is a strong evolutionary inducement towards uniform expressed trade behavior and in particular towards mutual cooperation, the uniform expressed trade behavior that generates the highest trader fitness scores.

On the other hand, the outlier economy $(5, e)$ realized at distance value 11 demonstrates that these evolutionary pressures towards mutual cooperation do not always dominate. Eleven traders in the final generation for this economy end up latched to other traders. Specifically, one mutant UD trader ultimately latches on to another trader in a mutual PC relationship, and a second mutant trader engages in UD and RD behavior that provokes retaliatory RD behavior on the part of 14 other traders. Ten of these 14 RD traders end up latched to the mutant RD trader, with the latter getting the worst of the bargain: he achieves a fitness score of only 0.62 whereas the traders who latch on to him achieve fitness scores between 1.33 and 1.43. The mean FIT value of 1.34 attained by this economy thus hides a high degree of intra-generational variation in both trade behavior and fitness scores.

The outlier economy $(65, e)$ realized at distance value 23 is also of interest. Traders in this economy exhibit a high degree of PC behavior in the first few generations. By generation 17, however, most traders are engaging in UD and RD behavior. In generation 18 the economy suffers a wallflower crash; only three trades take place in each of the three final trade cycles in generation 18. By generation 50 the economy is still in an unsettled state; 23 of the 24 traders engage in UD behavior although they ultimately end up in latched PC relationships. By generation 64, however, the economy has recovered from the wallflower crash; UD behavior is

infrequent and PC behavior is predominant. Indeed, the stability check for this case reveals that the average fitness score attained by the traders steadily increases towards the mutual cooperation payoff level 1.40 over generations 19 through 64 and fluctuates closely around 1.40 over generations 64 through 100.

Consider, next, an endogenous-type potential economy $e = e(\text{MS1}, 1, 1)$ comprising 24 buyer-seller traders with an offer quota bq again set equal to 1 but an acceptance quota sq reduced from 24 to 1. This economy e thus has zero excess capacity $EC(e)$, implying that a trader who randomly or recurrently directs his trade offers among other traders would now risk having his offers refused on the basis of limited acceptance capacity. As depicted in Figure 1(b), the base trade pattern $V^o(e)$ for this economy is as follows: Each trader is in a doubly-latched relationship with another trader. This base trade pattern implies that no trader functions either as a pure seller or as a pure seller. Moreover, since $bq = sq = 1$, each trader must be in an exclusive, continuous, two-way trade relationship with just one other trader.

As seen in Table IV, 100% of the twenty realized economies (s, e) generated for this e were observed to lie in the distance cluster 0–4, meaning that at most 4 of the 24 traders in the final generation of each realized economy deviate from the base trade pattern. These deviations take two forms: wallflowers; and traders who rotate their offers between two traders instead of latching on to one. The wallflowers are either UD traders who receive retaliatory refusal payoffs or cooperative traders who by chance receive so many refusal payoffs from all potential trade partners that the expected utility they assign to each of these potential trade partners drops below zero.

D^o Cluster	% Runs	Mean UD	Mean PW	Mean PC	Mean FIT
0–4	100%	8% (22%)	2% (3%)	96% (5%)	1.19 (.06)

Table IV: Endogenous-Type Market Experiments with Zero Excess Capacity

Note that the mean FIT value for distance cluster 0–4 in Table IV is only 1.19, substantially below the mutual cooperation payoff level 1.40, despite the fact that 96% of the traders in this distance cluster exhibit PC behavior. The explanation for this relatively low social welfare outcome is instructive.

It is not aggression or predation in the form of UD or RD behavior that results in low fitness scores for the traders but rather the large accumulations of refusal payoffs that the traders incur in their attempts to find trade partners. In particular, refusal payoffs reflect the amount of “wasted” effort that traders expend in the process of attaining a persistent trade network and hence can be interpreted as transactions costs. These transactions costs are structurally determined by the form of the matching process and the assumption of zero excess capacity. They do not arise from trade behavior per se and so cannot be eliminated by evolutionary selection pressures on trade strategies. This finding cautions that the optimality of the trade networks that arise in evolutionary economic contexts cannot be directly assessed in terms of social welfare (fitness) outcomes, for these outcomes reflect the particular historical processes by which the trade networks were attained. Rather, care must be taken to distinguish between the costs of attaining a given trade network starting from some specified status quo and the benefits of the trade network once attained.

The structural risk of incurring refusal payoffs that arises in the current zero excess capacity context also explains the very tight clustering of distance values about 0 in Table IV, reflecting a very high degree of double latching. Traders who fail to latch invariably accumulate high refusal payoffs. Moreover, once a PC trader v latches on to another PC trader w , the resulting steady stream of high trade payoffs that w receives from v encourages w to direct trade offers back to v rather than to other traders who are randomly, recurrently, or steadily refusing w ’s trade offers.

Consider, next, an endogenous-type potential economy $e = e(\text{MS1,2,1})$ comprising 24 buyer-seller traders with an acceptance quota sq again set equal to 1 but an offer quota bq increased from 1 to 2. Given these quota values, e exhibits tight capacity $EC(e)$ in the sense of Section 3.1. In particular, the traders as a whole can place at most half their potential trade offers in any given trade cycle. As depicted in Figure 1(c), the base trade pattern $V^\circ(e)$ for this economy is specified to be the same as in the previous case: Each trader is in a doubly-latched relationship with another trader.

As indicated in Table V, the twenty realized economies (s, e) generated for this e were observed to lie in two distinct distance clusters. The first distance cluster 2–11 includes 85% of the realized economies and is characterized by very infrequent UD behavior, very frequent PC behavior, and a low mean FIT value. The second distance cluster at 24 includes 15% of

the realized economies and is characterized by 100% UD and PW behavior and an extremely low mean FIT value.

D° Cluster	% Runs	Mean UD	Mean PW	Mean PC	Mean FIT
2-11	85%	1% (2%)	1% (2%)	98% (4%)	0.95 (.04)
24	15%	100% (0%)	100% (0%)	0% (0%)	-0.17 (.00)

Table V: Endogenous-Type Market Experiments with Tight Capacity

The first distance cluster in Table V illustrates how highly asymmetric intra-generational fitness scores can arise and persist among identically behaving PC traders due entirely to ex ante capacity constraints. In particular, given tight capacity, the high structural risk of refusal is a strong inducement to PC behavior among traders when functioning in a buyer capacity. Nevertheless, tight capacity implies that some traders will fail to place all of their trade offers regardless of their desirability as trade partners. In particular, even if all traders exhibit PC behavior, it can happen by chance that some traders repeatedly place both their trade offers while other traders either place no trade offers or have their trade offers accepted only on a random or recurrent basis.

Specifically, the non-zero distance values in distance cluster 2-11 in Table V arise from two kinds of deviant behavior: (i) PC traders in a doubly-latched relationship who also latch on to a third PC trader (possible since $bq = 2$) but who refuse the latter trader's offers in return (since $sq = 1$); and (ii) PC traders who latch on to two other PC traders but who only recurrently accept their trade offers in return. In case (i), the PC trader latched to two PC traders achieves a relatively high fitness score of about 1.25; his fitness score is lower than the mutual cooperation level 1.40 due to refusal payoffs incurred during the process of finding placements for his two trade offers. However, the third PC trader who fails to place any of his trade offers achieves a very low fitness score of about 0.27 (due to his very large accumulation of refusal payoffs) while the PC trader in the one doubly-latched relationship receives a relatively low fitness score of about 1.03 due to the refusal payoffs he incurs in trying unsuccessfully to place his second trade offer. In case (ii), the PC trader latched to two other PC traders achieves a relatively high fitness score of about 1.32 whereas the two

PC traders to whom he is latched achieve a relatively low fitness score of about 0.66 due to recurrent refusal payoffs. For each type of deviation the relatively high fitness scores balance out the relatively low fitness scores. Consequently, the economies in this distance cluster all achieve an average fitness score of about 0.98 regardless of the number of traders who exhibit these deviations.

The second distance cluster at 24 in Table V consists of three realized economies (s, e) characterized by 100% UD and PW behavior. Indeed, for each of these economies, all traders in the final (fiftieth) generation are initial defectors who rapidly degenerate into PWs as refusal payoffs accumulate due both to retaliation and to capacity constraints. Economies for which all traders degenerate into persistent wallflowers will hereafter be referred to as *wallflower economies*. The stability checks for these three economies indicate that each economy actually becomes a wallflower economy very early on, by generation 10 at the latest, and remains a wallflower economy through generation 100. In short, there is complete coordination failure in each of these three economies.

Finally, consider an endogenous-type potential economy $e = e(\text{MS1}, 24, 1)$ comprising 24 buyer-seller traders with an acceptance quota sq again set equal to 1 but an offer quota bq increased all the way to 24. In this case, capacity is extremely tight and traders face an extraordinarily high structural risk of incurring refusal payoffs if they randomly or recurrently distribute their trade offers among other traders. As depicted in Figure 1(d), the base trade pattern $V^\circ(e)$ for this economy is specified to be the same as in the previous two cases: Each trader is in a doubly-latched relationship with another trader.

The distance clusters and corresponding behavioral profiles determined experimentally for this extremely tight capacity case, presented in Table VI, are similar to those for the tight capacity case depicted in Table V. A simple structural reason can be given for this similarity. Once a trader receives enough refusal payoffs from another trader to drop his expected utility assessment for that trader below zero, he stops directing trade offers to that trader. Consequently, there is an inherent upper bound to the transactions costs that traders can sustain from the increased risk of refusals resulting from an additional tightening of capacity. As Tables V and VI also indicate, however, moving from tight to extremely tight capacity results in the realization of twice as many wallflower economies.

D° Cluster	% Runs	Mean UD	Mean PW	Mean PC	Mean FIT
2–12	70%	3% (5%)	3% (7%)	94% (9%)	0.89 (.07)
24	30%	100% (0%)	100% (0%)	0% (0%)	-0.16 (.00)

Table VI: Endogenous-Type Market Experiments with Extremely Tight Capacity

5.2 Two-Sided Market Experiments

Consider the case of a two-sided potential economy $e = e(\text{MS2}, 1, 12)$ comprising 12 pure buyers and 12 pure sellers with each buyer having an offer quota $bq = 1$ and each seller having an acceptance quota $sq = 12$. These quota values indicate that the economy e is characterized by high excess capacity $EC(e)$ in the sense defined in Section 3.1. In particular, the structural risk to buyers of having their offers refused by sellers on the basis of limited acceptance capacity is zero. On the other hand, unlike buyer-sellers, pure sellers are forced to be inactive unless buyers happen to direct trade offers their way. Consequently, in contrast to endogenous-type economies with high excess capacity, the sellers in this two-sided economy with high excess capacity face a substantial structural risk of incurring wallflower payoffs. The economy e thus represents a “buyers’ market.” As depicted in Figure 3(a), the base trade pattern $V^\circ(e)$ for this economy e is as follows: Each buyer is latched to a seller, and no seller is a wallflower.

— Insert Figure 3 About Here —

As seen in Table VII, 75% of the twenty realized economies (s, e) generated for this e were observed to lie in the distance cluster 3–9 and 25% were observed to lie in the distance cluster 23–24. In the first distance cluster, sellers achieve a very low mean FIT value of only 0.35. This results from two factors: the high accumulation of wallflower payoffs by sellers due to high excess capacity; and the high percentage of UD and RD behavior exhibited by buyers. Indeed, the persistent trade networks for this first distance cluster reveal that UD and RD buyers typically latch on to a selected subset of PC or less predacious sellers and drive down their fitness scores to small positive values, causing the remaining sellers to become PWs with fitness scores close to 0; the nonzero distance values in distance cluster

3–9 are essentially a count of the sellers who become PWs. This buyer behavior ensures that the parasitized subset of sellers fare relatively well in the evolution step, due to the separate evolution of pure buyers and pure sellers, and so reproduce into the next generation. This in turn ensures a continual source of hosts for the buyers to prey upon.

D° Cluster	% Runs	Mean UD		Mean PW		Mean PC		Mean FIT	
		PB	PS	PB	PS	PB	PS	PB	PS
3–9	75%	97% (5%)	16% (34%)	2% (3%)	40% (12%)	3% (5%)	39% (28%)	1.74 (.27)	0.35 (.14)
23–24	25%	2% (3%)	5% (7%)	2% (3%)	5% (7%)	98% (3%)	95% (7%)	1.39 (.02)	1.02 (.03)

Table VII: Two-Sided Market Experiments with High Excess Capacity

In the second distance cluster in Table VII, sellers achieve a mean FIT value of 1.02 that is substantially below the mutual cooperation payoff level 1.40 despite the high percentage of PC behavior exhibited by both buyers and sellers. This low mean FIT value results from the large number of wallflower payoffs that sellers accumulate due to high excess capacity, a structural cause that is independent of how cooperatively the sellers behave in their trade interactions. The typical trade pattern exhibited in this distance cluster is buyers directing trade offers among sellers without latching.

Consider, instead, the case of a two-sided potential economy $e = e(\text{MS2}, 1, 1)$ comprising 12 pure buyers and 12 pure sellers in which each buyer again has an offer quota $bq = 1$ but the acceptance quota sq of each seller is reduced from 12 to 1. This economy e is characterized by zero excess capacity $EC(e)$, implying that buyers who randomly or recurrently distribute trade offers among sellers now face a structural risk of refusal due to limited acceptance capacity. In contrast, the structural risk to sellers of incurring wallflower payoffs by chance is now zero. Consequently, e is tilted towards being a “sellers’ market.” As depicted in Figure 3(b), the base trade pattern $V^\circ(e)$ for this economy e is as follows: Each seller is in a latched relationship with one buyer, and no buyer is a wallflower.

As seen in Table VIII, 80% of the twenty realized economies (s, e) for this e were observed to lie in the distance cluster 0–2. Also, one realized economy was observed at distance value 4 and three realized economies were observed at distance value 24.

D° Cluster	% Runs	Mean UD		Mean PW		Mean PC		Mean FIT	
		PB	PS	PB	PS	PB	PS	PB	PS
0–2	80%	15% (32%)	22% (38%)	1% (3%)	1% (3%)	94% (6%)	86% (25%)	1.07 (.20)	1.34 (.21)
4	5%	100% (NA)	100% (NA)	17% (NA)	17% (NA)	0% (NA)	0% (NA)	0.62 (NA)	0.29 (NA)
24	15%	0% (0%)	22% (20%)	0% (0%)	8% (0%)	89% (16%)	78% (20%)	0.24 (.08)	1.42 (.05)

Table VIII: Two-Sided Market Experiments with Zero Excess Capacity

The mean FIT value for buyers in the distance cluster 0–2 in Table VIII is low relative to the mean FIT value for sellers; indeed, buyers attain a lower mean FIT value than sellers in all but 1 of the 16 realized economies in this distance cluster. The reason for this is the rather heavy accumulation of refusal payoffs that buyers incur in the course of finding their trade partners. Even though the economies in this distance cluster end up in a highly coordinated state, with almost all traders engaging in PC behavior in doubly-latched pairs, each buyer typically accumulates 2 or 3 refusal payoffs from a wide range of sellers on the way to achieving this coordinated state.

Strong evolutionary inducements exist for traders of a particular type to exhibit similar trade behavior since they are evolved together in the evolution step. Given the structural risk of refusal already faced by buyers, it is therefore not surprising to see in Table VIII that buyers in the distance cluster 0–2 largely evolve into PCs to avoid retaliatory refusal payoffs. The question arises, however, why sellers largely coordinate on PC behavior rather than engaging in more frequent UD/RD behavior to take advantage of the relatively weak position of buyers in this zero excess capacity context. For example, in the realized economy (641, e) with distance value 0, all sellers coordinate on the UD behavior $cd:c$ against buyers uniformly engaged in the PC behavior $c:c$, resulting in a FIT value for sellers of 1.47 that is higher than the mean FIT value of 1.34 that sellers achieve in distance cluster 0–2. However, seller coordination on UD/RD behavior is observed in only 3 of the 16 realized economies in distance cluster 0–2.

The realized economy at distance value 4 in Table VIII has several unusual features: for example, trade histories with exceedingly long handshakes and persistent portions (e.g., 26 actions in H and 28 actions in P), a great variation in fitness scores among traders of the

same type, and a mean FIT value for pure buyers that exceeds the mean FIT value for sellers. The three realized economies at distance value 24 also exhibit unusual features, such as exceedingly high refusal payoffs for buyers and a lack of any latching.

The stability checks conducted for the realized economies at distance values 4 and 24 reveal that the average fitness scores achieved by successive trader generations in these economies over generations 1 through 100 display persistent drifting (distance value 4), cycling (distance value 24, $s = 413$), or bubbles (distance value 24, $s = 63$ and $s = 31$). On the other hand, many of the realized economies in the distance cluster 0–2 also exhibit unsettled fitness score behavior over generations 1 through 100 in the form of persistent drifting, bubbles, or regime shifts. The reason for this appears to be that trade network formations are particularly fragile in this zero excess capacity context since they form in response to refusal payoffs and yet support largely PC or even $c:c$ trade behavior. For example, a $c:c$ economy with distance value 0 (complete latching with no current refusals) can suddenly revert to an economy with distance value 24 (no latching and a high number of refusals) if a mutant $c:c$ buyer enters the economy and disrupts current latched relationships by randomly scattering trade offers across sellers in successive trade cycles. The latter scenario arises, for example, in the final generation of the realized economy $(413, e)$ with distance value 24.

Consider, next, the case of a two-sided potential economy $e = e(\text{MS2}, 2, 1)$ comprising 12 pure buyers and 12 pure sellers in which each seller again has an acceptance quota $sq = 1$ but the buyer offer quota bq is increased from 1 to 2. This economy e is characterized by tight capacity $EC(e)$ in the sense of Section 3.1. In particular, buyers can now place at most half their potential trade offers in each trade cycle, implying they face a high structural risk of refusal, whereas the structural risk to sellers of incurring wallflower payoffs by chance remains at zero. Consequently, e is a “sellers’ market.” As depicted in Figure 3(c), the base trade pattern $V^o(e)$ for this economy e is as follows: Each buyer directs trade offers to sellers without latching, and no seller is a wallflower.

As seen in Table IX, 55% of the twenty realized economies (s, e) generated for this e were observed to lie in the distance cluster 0–7. The typical trade behavior exhibited in this distance cluster is $c:c$ buyers directing their two trade offers randomly among $c:c$ sellers, who are forced by the acceptance quota $sq = 1$ to refuse all but one trade offer in each trade cycle.

Consequently, buyers accumulate large numbers of refusal payoffs in each trade cycle and achieve a very low mean FIT value of only 0.30 despite mutually cooperative trade behavior. Deviations from the base trade pattern consist largely of buyers who become wallflowers by chance when they receive so many refusal payoffs from sellers that all sellers become intolerable to them. Two other forms of deviation also infrequently occur: *c:c* buyers who manage to latch on to a *c:c* seller when refusals (due to limited capacity) make the seller intolerable to all other buyers; and UD sellers who become wallflowers when all buyers direct their offers elsewhere. The latter deviations explain why the mean FIT value of 1.35 for sellers in distance cluster 0–7 is below the mutual cooperation payoff level 1.40.

D° Cluster	% Runs	Mean UD		Mean PW		Mean PC		Mean FIT	
		PB	PS	PB	PS	PB	PS	PB	PS
0–7	55%	2% (3%)	5% (9%)	19% (10%)	4% (7%)	81% (10%)	96% (6%)	0.30 (.06)	1.35 (.09)
24	45%	100% (0%)	90% (28%)	82% (26%)	77% (34%)	3% (8%)	5% (13%)	0.04 (.20)	0.22 (.39)

Table IX: Two-Sided Market Experiments with Tight Capacity

Table IX also reports that nearly all traders in the nine realized economies lying in distance cluster 24 exhibit UD behavior, a dramatically different outcome from the predominantly *c:c* trade behavior observed for distance cluster 0–7. In three of these realized economies, all traders end up either in latched RD relationships or as PWs, with buyers attaining a mean FIT value of 0.32 and sellers attaining a mean FIT value of 0.76. The stability checks for these three realized economies reveal that two of them exhibit unsettled behavior over generations 1 through 100 in the form of a drifting average fitness score and the third suffers a wallflower crash in generation 60 that persists through generation 100. In the remaining six realized economies lying in distance cluster 24, all traders are initial defectors and end up as PWs. Indeed, the stability checks for these economies reveal that they become wallflower economies by at most the tenth generation and remain wallflower economies through generation 100.

Finally, consider the case of a two-sided potential economy $e = e(\text{MS2}, 12, 1)$ comprising 12 pure buyers and 12 pure sellers in which each seller again has an acceptance quota $sq = 1$

but the buyer offer quota bq is increased all the way to 12. This economy e is characterized by extremely tight capacity $EC(e)$ in the sense of Section 3.1. In particular, any one buyer is now potentially capable of exhausting all seller capacity in the economy, hence buyers face an extraordinarily high structural risk of refusal. On the other hand, sellers still face zero structural risk of becoming wallflowers by chance. This e is therefore a “sellers’ market.” As depicted in Figure 3(d), the base trade pattern $V^o(e)$ for this e is as follows: Each seller accepts trade offers randomly from exactly two buyers, and no buyer is a wallflower.

As seen in Table X, the twenty realized economies (s, e) generated for this e form three distinct distance clusters. In distance cluster 0–6, active traders are predominantly engaging in $c:c$ trade behavior. The most frequent deviations from the base trade pattern are nine $c:c$ buyers who become PWs by chance when they receive so many refusals from sellers due to limited capacity that all sellers become intolerable to them, and seven $c:c$ sellers who either have one buyer latched on to them or who accept offers randomly from more than two buyers. Other less frequently observed deviations consist of one UD buyer who becomes a PW by retaliatory refusals, two UD sellers who induce retaliatory RD behavior in a buyer that leads the buyer to latch on to the seller, and one $c:c$ seller who becomes a PW by chance.

In contrast, in the four realized economies in distance cluster 15–17, nearly all sellers are UDs and all traders end up either in latched relationships (predominantly RD) or as PWs. The stability checks for these economies indicate that all exhibit an unsettled drift in average fitness score over generations 1 through 100. Finally, the 9 economies with distance value 24 are wallflower economies in which all traders are initial defectors and end up as PWs. The stability checks for the latter economies reveal that these economies become wallflower economies by the eighteenth generation at the latest and remain wallflower economies through generation 100.

D° Cluster	% Runs	Mean UD		Mean PW		Mean PC		Mean FIT	
		PB	PS	PB	PS	PB	PS	PB	PS
0–6	35%	1% (3%)	1% (3%)	12% (4%)	1% (3%)	86% (7%)	96% (6%)	0.31 (.03)	1.37 (.06)
15–17	20%	10% (14%)	92% (14%)	35% (7%)	2% (4%)	17% (20%)	25% (34%)	0.35 (.17)	1.22 (.20)
24	45%	100% (0%)	100% (0%)	100% (0%)	100% (0%)	0% (0%)	0% (0%)	-0.10 (.00)	-0.01 (.00)

Table X: Two-Sided Market Experiments with Extremely Tight Capacity

5.3 Partially-Fluid Market Experiments

Consider the case of a partially fluid potential economy $e = e(\text{MS3}, 1, 16)$ with 8 pure buyers, 8 pure sellers, and 8 buyer-sellers, and with a buyer offer quota $bq = 1$ and a seller acceptance quota $sq = 16$. This economy e has high excess capacity in the sense of Section 3.1 and constitutes a “buyers’ market.” On the one hand, each seller is individually capable of accepting all potential trade offers by buyers, implying that buyers face zero structural risk of refusal. On the other hand, pure sellers face a substantial structural risk of receiving wallflower payoffs because of the relative scarcity of trade offers. As depicted in Figure 4(a), the base trade pattern $V^\circ(e)$ for this economy e is as follows: Each pure buyer and buyer-seller directs trade offers to buyer-sellers and/or pure sellers without latching, and no pure seller is a wallflower.

— Insert Figure 4 About Here —

As seen in Table XI, the twenty economies (s, e) generated for this partially fluid economy e with high excess capacity are fairly evenly divided among three distinct distance clusters. In distance cluster 0–2, the relatively low mean FIT value for pure sellers is due primarily to large accumulations of wallflower payoffs resulting from high excess capacity; UD behavior is relatively infrequent and almost all traders are PCs.

D° C1st.	% of Runs	Mean UD			Mean PW			Mean PC			Mean FIT		
		PB	PS	BS	PB	PS	BS	PB	PS	BS	PB	PS	BS
0–2	30%	2% (5%)	2% (5%)	6% (10%)	2% (5%)	2% (5%)	0% (0%)	98% (5%)	98% (5%)	81% (32%)	1.38 (.01)	1.03 (.03)	1.38 (.03)
6–9	35%	25% (40%)	41% (48%)	39% (41%)	14% (35%)	41% (48%)	0% (0%)	75% (40%)	48% (41%)	64% (38%)	1.16 (.11)	0.73 (.24)	1.25 (.09)
16–21	35%	98% (4%)	23% (38%)	98% (4%)	16% (35%)	30% (30%)	2% (4%)	18% (34%)	40% (37%)	21% (35%)	1.13 (.63)	0.55 (.25)	1.43 (.30)

Table XI: Partially-Fluid Market Experiments with High Excess Capacity

In distance cluster 6–9, the mean FIT value for pure sellers is further depressed by the increased frequency of their own UD behavior, which provokes retaliatory RD behavior or refusals from buyers, and by the more frequent UD behavior of buyers which provokes retaliatory RD behavior in some pure sellers. The UD and retaliatory RD behavior by buyers who interact with UD pure sellers and with UD buyer-sellers in their capacity as sellers tends to lower the mean FIT value of buyers as well. Interestingly, in two of the seven realized economies in this distance cluster, all pure sellers exhibit UD behavior and degenerate rapidly into PWs, leaving the pure buyers and buyer-sellers to form their own persistent trade network. In another realized economy, it is the pure buyers who exhibit UD behavior and soon degenerate into PWs, leaving the buyer-sellers and pure sellers to form their own persistent trade network. Finally, for two other realized economies, UD buyer-seller behavior leads to the near-complete ostracism of buyer-sellers by other traders, with the buyer-sellers then essentially forming one isolated persistent trade network and the pure buyers and pure sellers forming another. Note that buyer-sellers are the only trader type able to thrive on their own.

Finally, in the distance cluster 16–21, the mean FIT value for pure sellers is further reduced due to the now frequent UD and RD behavior exhibited both by pure buyers and by buyer-sellers, nearly all of whom are latched to the pure sellers. Indeed, in six of the seven realized economies in this distance cluster, buyer-sellers act as pure buyers only, and all but one buyer-seller is latched to a pure seller. Buyers latched to the same pure seller tend to display similar trade behavior, whether they are pure buyers or buyer-sellers, despite the fact that pure buyers are evolved separately from buyer-sellers. The seventh realized economy in this distance cluster is also of interest: all pure sellers are initial defectors who degenerate into PWs; all pure buyers are UD who latch on to a single buyer-seller, with all but one exhibiting RD behavior in this latched relationship; and all buyer-sellers direct their trade

offers to other buyer-sellers and engage in PC trade behavior with each other.

The stability checks for the twenty realized economies (s, e) generated for this partially fluid economy e with high excess capacity reveal unsettled fitness score behavior over generations 1 through 100 in the form of a wallflower collapse (1 case), bubbles (2 cases), regime shift (6 cases), and persistent drifting (4 cases). It was at first conjectured that this observed instability might be due to the small population size of 8 for each trader type. Surprisingly, however, when six realized economies were re-run with an increased population size of 12 for each trader type, keeping all other parameter values fixed, the resulting distance values, trade behaviors, and social welfare outcomes closely resembled those obtained for the smaller population size. It therefore appears that the observed instabilities may instead be due to the fluid role played by buyer-sellers. In particular, the ability of buyer-sellers to function either as buyers or as sellers permits them to crowd out the pure sellers or the pure buyers, thus causing them to degenerate into PWs. In addition, buyer-sellers have the unique ability to form a self-sufficient network of trade relationships without the participation of either pure buyers or pure sellers.

Next consider the case of a partially fluid potential economy $e = e(\text{MS3}, 1, 1)$ with a buyer offer quota bq again set equal to 1 but a seller acceptance quota sq decreased from 16 to 1. This economy e is characterized by zero excess capacity $EC(e)$ in the sense of Section 3.1, which tilts the economy back towards being a “sellers’ market.” In particular, buyers who distribute offers recurrently or randomly among sellers now face a structural risk of incurring refusal payoffs from sellers on the basis of limited acceptance capacities, whereas the structural risk to pure sellers of receiving wallflower payoffs due to excess capacity has been eliminated. As depicted in Figure 4(b), the base trade pattern $V^o(e)$ for this economy e is as follows: Each pure buyer is latched to one buyer-seller, each buyer-seller is latched to one pure seller, and no pure seller is a wallflower.

Table XII reports the distance values, trade behaviors, and social welfare outcomes observed for this partially fluid economy e with zero excess capacity. A comparison of these findings to the much more diffuse findings reported in Table XI for high excess capacity indicates the extraordinarily strong disciplinary role played by ex ante capacity constraints in the determination of evolutionary outcomes for partially fluid economies.

D° C1st.	% of Runs	Mean UD			Mean PW			Mean PC			Mean FIT		
		PB	PS	BS	PB	PS	BS	PB	PS	BS	PB	PS	BS
0–6	80%	6% (24%)	20% (10%)	9% (24%)	0% (0%)	1% (3%)	2% (4%)	98% (4%)	92% (24%)	97% (5%)	1.16 (.06)	1.42 (.14)	1.11 (.13)
16–24	20%	100% (0%)	28% (42%)	47% (47%)	28% (42%)	25% (31%)	6% (11%)	0% (0%)	19% (14%)	53% (47%)	0.20 (.17)	0.87 (.33)	0.87 (.22)

Table XII: Partially-Fluid Market Experiments with Zero Excess Capacity

Specifically, as seen in Table XII, 80% of the twenty realized economies (s, ϵ) generated for this zero excess capacity ϵ were observed to lie in the distance cluster 0–6. Indeed, 40% of these economies actually have distance value $D^\circ = 0$, meaning that 100% of the traders in the final generations for these economies are coordinated into disjoint trading triads consisting of one pure buyer latched to one buyer-seller who in turn is latched to one pure seller. Moreover, nearly all traders in the distance cluster 0–6 exhibit PC behavior. The relatively low mean FIT values for pure buyers and buyer-sellers in this distance cluster are largely the result of refusal payoffs accumulated by chance. The mean FIT value for buyer-sellers is also reduced by the UD and RD behavior engaged in by some pure sellers against largely PC buyer-sellers. This UD and RD behavior by pure sellers also explains why the mean FIT value of 1.42 for pure sellers is slightly higher than the mutual cooperation payoff level 1.40. Deviations from the base trade pattern in this distance cluster consist of occasionally broken triads, e.g., a pure buyer latched directly to a pure seller with the extra buyer-seller appearing either as a PW or as an extra link inserted between a pure buyer and a buyer-seller in what otherwise would be a base trade pattern triad.

The four realized economies in distance cluster 16–24 are dominated by the following types of deviations from the base trade pattern: widespread failure of pure buyers to latch with buyer-sellers and of buyer-sellers to latch with pure sellers (one case); widespread direct matching between pure buyers and pure sellers, with buyer-sellers latched largely among themselves (two cases); and pure buyers degenerated into PWs by refusal, pure sellers degenerated into PWs by chance, and buyer-sellers latched among themselves (one case). The reason for the relatively high frequency of PC behavior among buyer-sellers in this distance cluster results from the fact that, in three of the four realized economies in this distance cluster, the buyer-sellers end up matched largely among themselves. As seen for endogenous-type economies, self-matching generally induces mutual PC behavior among buyer-sellers.

Now consider the case of a partially fluid potential economy $e = e(\text{MS3}, 2, 1)$ with a seller acceptance quota sq again set equal to 1 but a buyer offer quota bq increased from 1 to 2. This economy e is characterized by tight capacity $EC(e)$ in the sense of Section 3.1, and constitutes a “sellers’ market.” In particular, buyers who distribute trade offers recurrently or randomly among sellers now face a high structural risk of incurring refusals from sellers on the basis of limited acceptance capacities whereas the structural risk to sellers of incurring wallflower payoffs by chance is zero. As depicted in Figure 4(c), the base trade pattern $V^o(e)$ for this economy e is as follows: Each pure buyer directs trade offers to at least two sellers (bs and/or ps), each buyer-seller functions as a pure seller (i.e., he does not make any trade offers) and no trader is a wallflower.

As seen in Table XIII, 75% of the twenty realized economies (s, e) generated for this tight capacity economy e lie in distance cluster 0–7, 10% are at distance value 10, and 15% are at distance value 24. In distance cluster 0–7, over three-fourths of the traders exhibit PC behavior. Nevertheless, pure buyers and buyer-sellers achieve lower mean FIT values than pure sellers due primarily to large accumulations of refusal payoffs resulting from tight capacity. Buyer-sellers do better than pure buyers since almost all of them end up functioning as pure sellers in this sellers’ market; but they typically still incur many refusal payoffs in the process of attaining this persistent state.

D^o Clst.	% of Runs	Mean UD			Mean PW			Mean PC			Mean FIT		
		PB	PS	BS	PB	PS	BS	PB	PS	BS	PB	PS	BS
0–7	75%	2% (4%)	19% (38%)	17% (34%)	6% (8%)	4% (7%)	1% (3%)	75% (28%)	80% (38%)	83% (34%)	0.63 (.19)	1.30 (.19)	0.90 (.07)
10	10%	0% (0%)	50% (50%)	50% (50%)	19% (6%)	44% (44%)	50% (50%)	81% (6%)	56% (44%)	50% (50%)	0.32 (.10)	0.85 (.53)	0.46 (.49)
24	15%	100% (0%)	100% (0%)	100% (0%)	100% (0%)	100% (0%)	100% (0%)	0% (0%)	0% (0%)	0% (0%)	-0.13 (.00)	-0.02 (.00)	-0.13 (.00)

Table XIII: Partially-Fluid Market Experiments with Tight Capacity

In one of the two economies at distance value 10, buyer-sellers exhibit UD behavior and degenerate rapidly into PWs. In the second economy at distance value 10, it is the pure sellers who exhibit UD behavior and seven out of eight degenerate rapidly into PWs. Consequently, UD behavior appears to be risky for sellers even in this sellers’ market since buyer-sellers and pure sellers must compete against each other for the attention of the pure buyers. Finally, the three economies at distance value 24 are wallflower economies in which each trader is an initial defector. The structural checks for the latter economies reveal that

they become wallflower economies by generation 5 at the latest, and they remain wallflower economies through generation 100.

Finally, consider the case of a partially fluid potential economy $e = e(\text{MS3}, 1, 16)$ with a buyer offer quota bq again set equal to 1 but a seller acceptance quota sq now increased all the way to 16. This economy e is characterized by extremely tight capacity $EC(e)$ in the sense of Section 3.1. In particular, buyers who distribute offers recurrently or randomly among sellers now accumulate exceedingly high numbers of refusal payoffs. As depicted in Figure 4(d), the base trade pattern $V^o(e)$ for this economy e is as follows: Each trader is a wallflower.

D^o Clst.	% of Runs	Mean UD			Mean PW			Mean PC			Mean FIT		
		PB	PS	BS	PB	PS	BS	PB	PS	BS	PB	PS	BS
0	75%	100% (0%)	100% (0%)	100% (0%)	100% (0%)	100% (0%)	100% (0%)	0% (0%)	0% (0%)	0% (0%)	-0.13 (.00)	-0.01 (.00)	-0.13 (.00)
16	5%	0% (NA)	0% (NA)	100% (NA)	25% (NA)	0% (NA)	25% (NA)	38% (NA)	88% (NA)	0% (NA)	0.52 (NA)	1.33 (NA)	0.57 (NA)
23–24	20%	0% (0%)	6% (11%)	3% (5%)	0% (0%)	3% (5%)	0% (0%)	97% (5%)	97% (5%)	97% (5%)	0.39 (.07)	1.36 (.07)	0.79 (.04)

Table XIV: Partially-Fluid Market Experiments with Extremely Tight Capacity

As seen in Table XIV, 75% of the twenty realized economies (s, e) for this economy e with extremely tight capacity have distance value 0, implying that they are wallflower economies. In each of these wallflower economies, all traders in the final generation engage in initial defections. Since the chance of forming persistent trade relationships is essentially zero due to the exceedingly high structural risk of refusal, traders who do best are those who defect from the very first trade. As the stability checks for these economies indicate, however, the resulting retaliatory refusals on top of structurally generated refusals causes these economies to degenerate almost immediately into wallflower economies and to persist as wallflower economies through generation 100. As seen in Table XIV, the few realized economies (s, e) that manage to achieve positive distance values, and hence to escape becoming wallflower economies, are those in which traders largely avoid UD and RD behavior. The typical trade pattern observed for these latter economies is as follows: Each buyer is directing offers to sellers without latching, and each pure seller is accepting offers recurrently from exactly two buyers.

6 Concluding Remarks

Among the more interesting experimental findings detailed in Section 5, two stand out.

First, for many of the tested potential economies e , the distance values for the persistent trade networks $K^p(s, e)$ tend to cluster around two or three isolated distance values, and the mean distance of each distance cluster tends to be strongly correlated with the mean unprovoked defector (UD), persistent wallflower (PW), persistent cooperator (PC), and fitness (FIT) profiles calculated for the distance cluster. For such economies, then, there does not appear to be any central-tendency network in the sense defined by Banks and Carley (1994) but rather a number of different local basins of attraction. One possible explanation for these distinct distance clusters is that they correspond to multiple Nash equilibria for the underlying evolutionary match-and-play game in which the traders are participating. On the other hand, the distinct distance clusters could be artifacts of the relatively small sample size of 20 that was used in the experiments in order to keep the graphical determination and analysis of trade network formations manageable. More testing is needed here.

Second, the optimality criteria commonly used to evaluate the performance of Gale-Shapley matching mechanisms in static market contexts—namely, Pareto optimality and pairwise stability—proved to be highly incomplete indicators of performance in the evolutionary market context of the current study. As repeatedly seen in the computational experiments reported in Section 5, all of which rely on Gale-Shapley matching, trade network formation and trade behavior are systematically related to market structure and to ex ante capacity constraints as reflected in the relative and absolute sizes of the buyer offer quota bq and the seller acceptance quota sq . Moreover, transactions costs (negative R payoffs) and inactivity costs (0 wallflower payoffs) are important determining factors in the formation of trade networks; and these costs can cause social welfare to be low even if all active traders are persistent cooperators.

These observations suggest that matching behavior and trade behavior should be studied conjointly in dynamic market contexts. Yet even this may not provide sufficient flexibility for the understanding of real-world markets. In the current study, for example, traders adapt both their matching behavior and their trade behavior in response to trade interactions in each trade cycle, but only trade strategies are permitted to evolve structurally. The

structure of the market and the rules by which matching takes place are assumed to be given exogenously. In real-world market contexts, however, market structure, matching mechanisms, and trade strategies simultaneously coevolve, and for good reason: all three interdependent facets strongly affect market performance.


References

- Ashlock, D., Smucker, M. D., Stanley, E. A., Tesfatsion, L. (1996). “Preferential Partner Selection in an Evolutionary Study of Prisoner’s Dilemma,” *BioSystems* **37**, 99–125.
- Arthur, W. Brian, Durlauf, S. N., and Lane, D. A., Eds. (1997). *The Economy as an Evolving Complex System, II*, Santa Fe Institute Studies in the Sciences of Complexity, Volume XXVII. Reading, MA: Addison-Wesley.
- Banks, D., and Carley, K. (1994). “Metric Inference for Social Networks,” *Journal of Classification* **11**, 121–149.
- Brock, W., and Durlauf, S. N. (1995). “Discrete Choice with Social Interactions,” Working Paper No. 95-10-084, Santa Fe Institute, Santa Fe, NM.
- De Vany, A. (1996). “The Emergence and Evolution of Self-Organized Coalitions,” in *Computational Economic Systems: Models, Methods, and Econometrics* (M. Gilli, Ed.), pp. 25–50. New York: Kluwer Scientific Publications.
- Durlauf, S. (1996). “Neighborhood Feedbacks, Endogenous Stratification, and Income Inequality,” in *Dynamic Disequilibrium Modelling* (W. A. Barnett, G. Gandolfo, and C. Hillinger, Eds.), pp. 505–534. Cambridge: Cambridge University Press.
- Ellison, G. (1992). “Learning, Local Interaction, and Coordination,” *Econometrica* **61**, pp. 1047–1072.
- Epstein, J. M., and Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*, Cambridge, MA: MIT Press/Brookings.
- Fogel, D. B. (1995). “On the Relationship Between the Duration of an Encounter and the Evolution of Cooperation in the Iterated Prisoner’s Dilemma,” *Evolutionary Computation* **3**, 349–363.
- Gale, D. and Shapley, L. (1962). “College Admissions and the Stability of Marriage,” *American Mathematical Monthly* **69**, 9–15.
- Guriev, S., and Shakhova, M. (1996). “Self-Organization of Trade Networks in an Economy with Imperfect Infrastructure,” in *Self-Organization of Complex Structures: From Individual to Collective Dynamics* (F. Schweitzer, Ed.), London: Gordon and Breach Scientific Publishers.


- Hauk, E. (1996), “Leaving the Prison: A Discussion of the Iterated Prisoner’s Dilemma Under Preference Partner Selection,” Thesis, European University Institute, Florence.
- Hirshleifer, D., and Rasmusen, E. (1989). “Cooperation in a Repeated Prisoners’ Dilemma with Ostracism,” *Journal of Economic Behavior and Organization* **12**, 87–106.
- Ioannides, Y. M. (1997). “Evolution of Trading Structure,” pp. 129–167 in Arthur et al., *The Economy as an Evolving Complex System, II, op. cit.*.
- Kirman, A. P. (1997). “The Economy as an Interactive System,” pp. 491–531 in Arthur et al., *The Economy as an Evolving Complex System, II, op. cit.*.
- Kitcher, P. (1993). “The Evolution of Altruism,” *The Journal of Philosophy* **90**, 497–516.
- Mailath, G., Samuelson, L., and Shaked, A. (1994). “Evolution and Endogenous Interactions,” SSRI Working Paper 9426, University of Wisconsin, Madison.
- McFadzean, D. (1995). “SimBioSys: A Class Framework for Evolutionary Simulations,” Master’s Thesis, Department of Computer Science, University of Calgary, Alberta, Canada.
- McFadzean, D., and Tesfatsion, L. (1997). “A C++ Platform for the Evolution of Trade Networks,” Economic Report No. 39, Iowa State University, to appear in *Computational Economics*.
- Orbell, J. M., and Dawes, R. M. (1993). “Social Welfare, Cooperators’ Advantage, and the Option of Not Playing the Game,” *American Sociological Review* **58**, 787–800.
- Roth, A., and Sotomayor, M. A. O. (1990). Two-sided matching: A study in game-theoretic modeling and analysis. Cambridge: Cambridge Univ. Press.
- Smucker, M. D., Stanley, E. A., and Ashlock, D. (1994). “Analyzing Social Network Structures in the Iterated Prisoner’s Dilemma with Choice and Refusal,” Technical Report CS-TR-94-1259, Computer Science Department, University of Wisconsin, Madison.
- Stanley, E. A., Ashlock, D. and Tesfatsion, L. (1994). “Iterated Prisoner’s Dilemma with Choice and Refusal of Partners,” in *Artificial Life III* (C. Langton, Ed.), pp. 131–175. Santa Fe Institute Studies in the Sciences of Complexity, Volume XVII. Reading, MA: Addison-Wesley.
- Tesfatsion, L. (1979). “Direct Updating of Intertemporal Criterion Functions for a Class of Adaptive Control Problems,” *IEEE Transactions on Systems, Man, and Cybernetics* **SMC-9**, 143–151.

- Tesfatsion, L. (1997a). “A Trade Network Game with Endogenous Partner Selection,” in *Computational Approaches to Economic Problems* (H. Amman, B. Rustem, A. B. Whinston, Eds.), pp. 249–269. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Tesfatsion, L. (1997b). “How Economists Can Get Alive,” pp. 533–564 in Arthur et al., *The Economy as an Evolving Complex System, II, op. cit.*.
- Vriend, N. J. (1995). “Self-Organization of Markets: An Example of a Computational Approach,” *Computational Economics* **8**, 205–231.
- Weisbuch, G., Kirman, A., and Herreiner, D. (1997). “Market Organization,” in *Simulating Social Phenomena* (R. Conte, R. Hegselmann, and P. Terna, Eds.), pp. 221–240. Lecture Notes in Economics and Mathematics, Volume 456. Berlin: Springer-Verlag.
- Young, P. (1993). “The Evolution of Conventions,” *Econometrica* **61**, 57–84.


Copyright © 1998 Leigh Tesfatsion. All Rights Reserved.




(a) High Excess Capacity
($EC = 23$)



(b) Zero Excess Capacity
($EC = 0$)



(c) Tight Capacity
($EC = -1/2$)



(d) Extremely Tight Capacity
($EC = -23/24$)

Figure 1: Base Trade Patterns for Endogenous-Type Economies with Different Ex Ante Capacities. A relatively larger box for a trader type indicates that this trader type achieves a relatively higher FIT value in the realized economies whose trade patterns approximate the base trade pattern.

THIS FIGURE IS AVAILABLE FROM THE AUTHOR UPON REQUEST

Figure 2: Stability Check for an Endogenous-Type Economy. The maximum, minimum, and average fitness scores are graphed for trader generations 1 through 100.

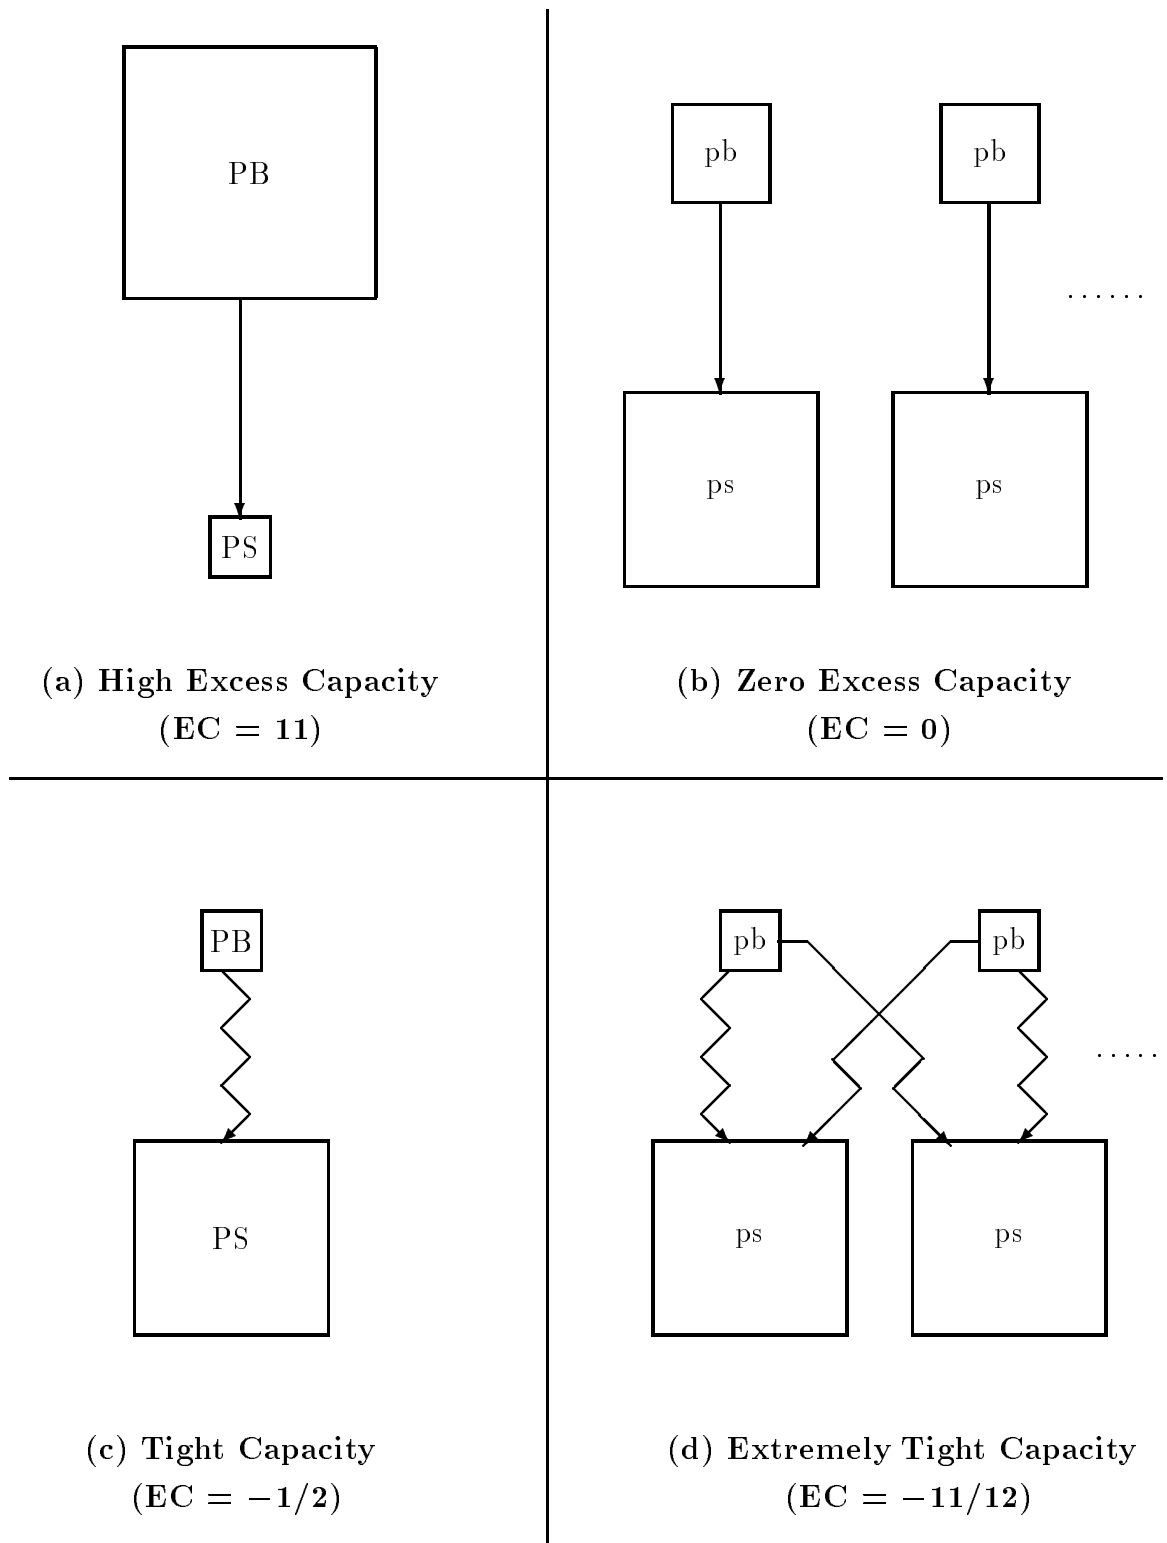


Figure 3: Base Trade Patterns for Two-Sided Economies with Different Ex Ante Capacities. A relatively larger box for a trader type indicates that this trader type achieves a relatively higher FIT value in the realized economies whose trade patterns approximate the base trade pattern.

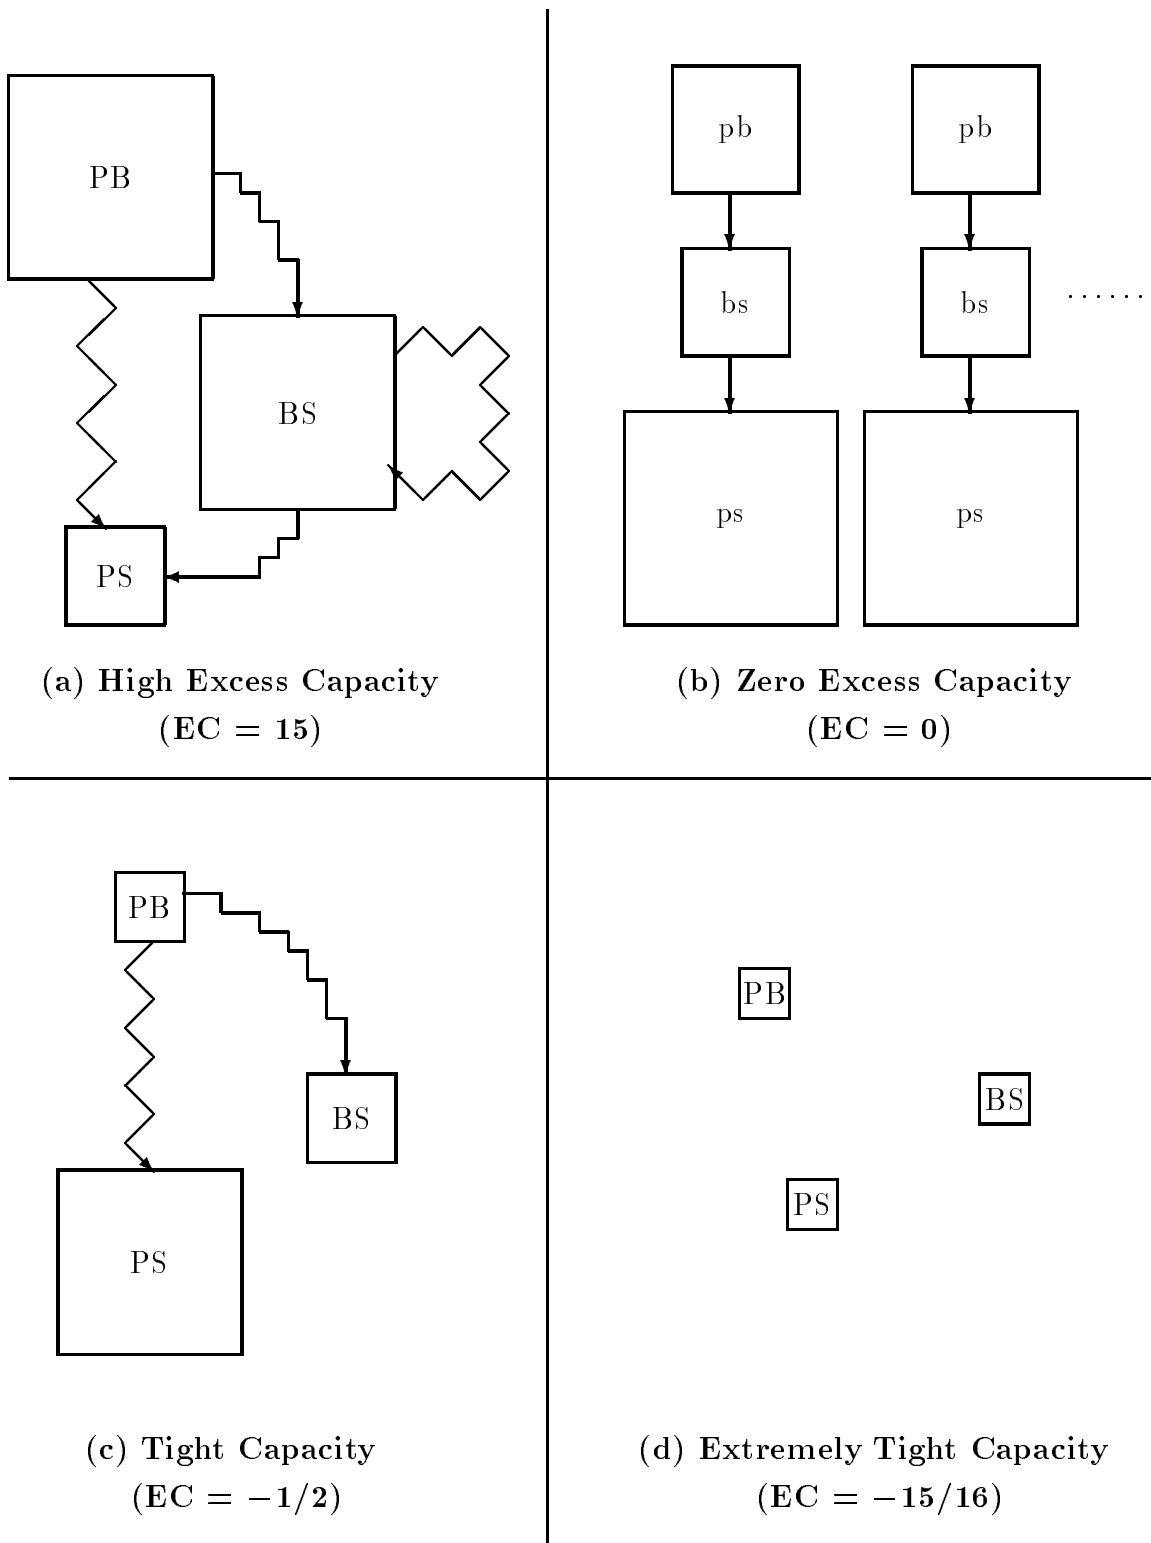


Figure 4: Base Trade Patterns for Partially-Fluid Economies with Different Ex Ante Capacities. A relatively larger box for a trader type indicates that this trader type achieves a relatively higher FIT value in the realized economies whose trade patterns approximate the base trade pattern.