

Individual Rationality and Market Efficiency

Steven Gjerstad, *Economic Science Institute, Chapman University, Orange, California.*

Jason Shachat, University of Durham, *Durham, U.K.* and Wuhan University, *Wuhan, China*

RUNNING HEAD: [Individual Rationality and Market Efficiency]

Correspondence address: Steven Gjerstad, 1 University Drive, Orange, CA 92866,
gjerstad@chapman.edu, 1-765-409-5310.

Abstract: Smith's (1962) demonstration that prices and allocations quickly converge to the competitive equilibrium in the continuous double auction (CDA) remains one of the most important results in experimental economics. Market experiments and exchange models have added considerably to our knowledge of how markets reach equilibrium, and how they respond to disruptions. Perhaps the best-known model of exchange in CDA market experiments is the random behavior "zero-intelligence" (ZI) model by Gode and Sunder (1993). They argue that the CDA generates efficient allocations and "convergence of transaction prices to the proximity of the theoretical equilibrium price," provided only that agents meet their budget constraints. We demonstrate that prices do not converge in their simulations. Their budget constraint requires that a buyer's currency never exceeds her commodity value, which is an unnatural restriction. Their conclusion that market efficiency results from the structure of the CDA independent of traders' profit seeking behavior rests on their claim that the constraints that they impose are a part of the market institution. We show that actually they impose *individual rationality*. Misinterpretation of this behavioural constraint has led to unproductive debate on market adjustment processes.

Key Words: Bounded rationality; double auction; exchange economy; market experiment; "zero intelligence" model

INTRODUCTION

Hayek (1945) asked how the information about aggregate market conditions that is distributed among many buyers and sellers, who all have private information about their costs of production or their values of consumption, can be coordinated through the trading process to reach an efficient allocation at stable prices. Chamberlin (1948) found that when buyers and sellers randomly encounter one another, negotiate prices, and possibly trade, prices are volatile and extracted only a fraction of available surplus. Smith (1962) had a key insight into the convergence process. He replaced Chamberlin's decentralized negotiation and exchange with the continuous double auction (CDA), in which ask prices, bid prices, and trades are publicly disclosed, and found that trade prices quickly approached the competitive equilibrium price and generated surplus approximated the competitive equilibrium surplus. His experiment provided a framework for analysis of Hayek's question. In Smith's CDA experiment, the experimenter induces costs and values that are private information in naturally occurring markets. Induced costs and values combined with public data on bids, asks, and trade prices are crucial for development of models of behavior in the CDA.

Several such models have been developed. These include the strategic model in Wilson (1986), three models that depart from full rationality – Friedman (1991), Easley and Ledyard (1993), and Gjerstad and Dickhaut (1998) – and the model of random behavior by Gode and Sunder (1993). Cason and Friedman (1996) compare predictions of the first two and the last of these five models to outcomes from market experiments with human subjects. Gjerstad (2007) briefly summarizes all five models. Gode and

Sunder reach the most striking conclusions. They assert that “a double auction . . . can sustain high levels of allocative efficiency even if agents do not . . . seek profits” (p. 120) and “the convergence of transaction prices to the proximity of the theoretical equilibrium price in ZI-C markets is a consequence of the market discipline” (p.131). The purpose of this paper is to assess these two claims: neither claim withstands scrutiny. We demonstrate that their view that trader rationality is largely unnecessary in models of market dynamics lacks foundation.

Our paper is organized as follows. In the next section, we summarize the model by Gode and Sunder, indicate its influence among researchers on learning in markets, and evaluate their two main claims. Prices from their simulations belie their claim that prices converge. Their claim that allocations are efficient in the absence of trader rationality or profit seeking behavior relies on their peculiar “budget constraint,” whereby no trader can come to the market with currency that exceeds her value for the commodity. In the section that follows the next one, we demonstrate that markets with induced costs and values are equivalent to exchange economies, and we use our exchange economy representation to demonstrate that the constraints imposed by Gode and Sunder are in fact individual rationality constraints. In the final section we consider examples that suggest the fundamental economic significance of the market adjustment problem and indicate several challenges that serious models of adjustment should confront.

GODE AND SUNDER’S ZERO-INTELLIGENCE MODEL

Gode and Sunder populate a market with simple bidding agents – which they call zero-intelligence (ZI) traders – in order to assess the forces that generate efficient

outcomes in CDA markets. In a CDA market, a seller may submit an ask at any time during a trading period. Similarly, at any time a buyer may submit a bid. An ask placed at or below the current high bid results in a trade at the bid price; a bid that meets or exceeds the current low ask results in a trade at the ask price. In the simulations that Gode and Sunder conduct, their ZI-constrained (ZI-C) traders have restricted bids or offers: a ZI-C buyer bids between zero and her budget; a ZI-C seller submits ask prices between his cost and some upper bound. They call this budget constraint “the market discipline” (p. 123) because the market prevents agents from violating their budget constraints. They compare outcomes of markets populated with these ZI-C traders to markets populated with ZI unconstrained (ZI-U) traders that are not subject to these constraints. Gode and Sunder (1993) summarize the conclusions that they draw from this comparison.

The primary cause of the high allocative efficiency of double auctions is the market discipline imposed on traders; learning, intelligence, or profit motivation is not necessary. The same market discipline also plays an important role in the convergence of transaction prices to equilibrium levels. (p. 134)

We find these two claims are erroneous. Nevertheless, their argument has been widely accepted and has given many researchers problematic starting points for their analyses. The influence of their paper can be seen through both the quantitative assessment of its influence and examples of its approval by prominent researchers.

As of April 2021, there have been 569 citations of the paper in the Social Sciences Citation Index. Thirty-six of these citations were between 1994 and 2000, eighty-five between 2001 and 2008, two hundred thirty-seven between 2008 and 2014, and two hundred eleven between 2015 and March 2021. The forty-two papers published in the JPE in 1993 average 189.3 citations; only two have more than 569 citations.

It is not difficult to find statements by prominent authors that accept the principle claims by Gode and Sunder. “The rules for competition, if well designed, can ensure that a market produces an allocation that is close to efficient even with traders who are *incapable of calculating what is in their interest*, according to experiments by Gode and Sunder (1993). The wisdom of the market compensates for the market participants’ lack of rationality.” (McAfee and McMillan (1996), emphasis added.) “Gode and Sunder (1993) give examples of double auction markets in which “zero-intelligence” traders (computers which bid randomly *subject only to budget constraints*) may achieve near perfect market efficiency.” (Conlisk (1996), emphasis added.) “Gode and Sunder . . . find that market efficiency levels close to 100% are attained even when their traders have “zero intelligence” in the sense that they submit random bids and asks that are *subject only to a budget constraint*.” (Tsfatsion (2002), emphasis added.) “[T]he continuous flow of offers, coupled with traders’ *budget constraints*, generates a mechanical but powerful push in the direction of efficient outcomes.” (Samuelson (2005), emphasis added.) “Gode and Sunder (1993) . . . suggest that efficiency of the double-auction institution *derives largely from its structure* rather than from individual learning.” (List (2005), emphasis added.) “[I]t has been repeatedly observed that, in laboratory experiments, double-auction institutions consistently produce allocations and *prices*

close to the predictions of price-taking models. These predictions are even sustained when subjects are random decision-makers as in Gode and Sunder (1993).” (Bosch-Domenech and Silvestre (1997), emphasis added.) “Gode and Sunder (1993) supported market “rationality” even when composed of “irrational” economic units. . . . Efficiency derived mainly from the structure of the market itself rather than from the advantages typically ascribed to human economic agents (i.e., motivation, intelligence, and learning ability). (Evans, 1997). “Contrary to one of the shibboleths of neoclassical economics on the necessity for explicit optimization by agents at a micro level for allocative efficiency in markets, they [Gode and Sunder] found that it is the *operation of the rules of double auction* that is instrumental for over 90% of allocative efficiency even with so-called zero intelligence traders.” (Markose, Arifovic, and Sunder (2007), emphasis added.)

Gode and Sunder claim that when the market institution enforces the traders’ budget constraints (or imposes “the market discipline”), that leads to price convergence and market efficiency. We dispute this conclusion. We show that their constraint cannot be meaningfully interpreted as a budget constraint. Rather, the constraint they impose can only be interpreted as a restriction on traders’ behavior, namely individual rationality. Contrary to their claims that “profit motivation is not necessary” to attain efficient outcomes (p. 134) or that their ZI “agents do not . . . seek profits” (p. 120), we show that their agents only accept offers that increase their payoff. Since their ZI-C simulations reach efficient outcomes and their ZI-U simulations do not, it follows that some restriction on trader behavior is crucial for market efficiency. Moreover, price convergence requires some form of learning or adaptation.

Price Convergence

Numerous CDA experiments beginning with Smith (1962) demonstrate that trades by human subjects quickly and reliably converge to the equilibrium price. Gode and Sunder claim that they obtain “convergence of transaction prices to equilibrium levels” in their markets with budget constrained traders. Fig. 1 (a) and Fig. 1 (b) compare a representative outcome from a simulation of the ZI model to a representative outcome from an experiment with human subjects. These two price paths demonstrate that prices converge in markets with human subjects, whereas they oscillate wildly in ZI-C simulations. In these two graphs, the standard deviations of trade prices in the selected periods are closest to the median of this statistic from simulations of 150 periods with the ZI-C model (median standard deviation 8.53) and for periods 6 – 15 from each of five sessions with human subjects (median standard deviation 1.70). Further comparison with a model that includes simple learning demonstrates that something more than individual rationality – but far less than full strategic rationality – approximates the competitive equilibrium outcome.

Figure 1 about here.

Although ZI-C simulations fail to converge, the simple heuristic belief learning (HBL) model from Gjerstad (2007) (which is based on the Gjerstad and Dickhaut (1998) model) leads to price convergence in a standard market experiment with induced values and costs. In their model, a seller forms a belief that his ask will be accepted by some

buyer. He bases his belief on observed market data, including frequencies of asks, bids, accepted asks, and accepted bids. Then he chooses an ask that maximizes his expected surplus. A seller submits his ask at a random time, with an exponential distribution that depends on both his expected surplus and on the fraction of time remaining in the period. The belief of a buyer, the buyer's bid choice, and the timing of her bid are similar.

Figure 1 (c) demonstrates that price standard deviations with this model are similar to markets with human subjects. This figure shows the price sequence from the HBL model simulation that is closest to the median standard deviation from all 15 periods in ten simulations (median standard deviation 1.74). These simulations are described in detail in Gjerstad (2007). Price standard deviations from the ZI model exceed those from human subject experiments and from the HBL model by a factor of about five. Figure 1 (d) reinforces the point that the HBL model captures price convergence. This figure depicts the outcome from a market with HBL model sellers and human buyers reported in Gjerstad (2007). Among five fifteen period sessions with HBL model agents on one side of the market and human subjects on the other side of the market, in this trading period the price standard deviation was lowest. The median of the price standard deviation is nearly 20 times as large in ZI simulations as it is in this period with sharp price convergence.

Prices in the Gode and Sunder model fluctuate in a range that is bounded above by the highest unit value of the buyers and below by the lowest unit cost of the sellers. These ranges tend to narrow as trade proceeds within a period, but in each new period values and costs are renewed. Their definition of price convergence appears particularly

contrived when it is interpreted in a practical context. Suppose that grain distributors and food producers arrive at the Minneapolis Grain Exchange to trade Hard Red Spring Wheat, with quotes in single contracts for 5000 bushels. Would we say that the market price had converged if at 9:35 a.m. – just after the market opens – the first buyer paid \$624, at 9:45 a.m. the next buyer paid \$371, later the third paid \$736, and so on all day, until the last buyer paid \$514 at 1:13 p.m., where \$514 is close to the price at the intersection of supply and demand? In the Gode and Sunder story not only does this happen on the first day; something similar happens day after day. We suggest, the benchmark for price convergence should be more substantial than either (1) trade of the last unit in the period near the competitive equilibrium price or (2) convergence in mean. In CDA experiments price convergence is much more robust.

A budget or an individual rationality constraint?

The interpretation of the budget constraint in Gode and Sunder is inconsistent with economic theory and reasonable assumptions on decentralized private information. Their budget constrained buyer always arrives at a market with an amount of money exactly equal to her value for the commodity she intends to purchase. This constraint cannot be interpreted as a budget constraint or as a form of market discipline. Surely, if a buyer arrives at the market with currency M , there is an economic rationale for the market organizers to restrict her bids b to the interval $[0, M]$, otherwise she won't be able to make payment if she bids more than M and ends up trading at a price greater than M . The inequality $b \leq M$ in the budget constraint is perfectly sensible. What happens when M exceeds the buyer's value? Why should the market organizers restrict her currency to

$M \leq v$? Of course, they have no reason nor capacity to do so. The market organizers do not care if a buyer brings more money to the market than her value for her intended purchase. Even if they did have a reason to impose this restriction, they do not know the value(s) of a buyer, and typically they will not know the amount of money that a buyer brings to the market, so they couldn't carry out this restriction if they had the desire to do so. The restriction that $M \leq v$ is entirely unnatural. Surely this constraint is not a part of any conceivable free market: its enforcement would require interrogation (to determine a buyer's values), unrestricted searches (to verify that the buyer has no more currency than her value), and confiscation of assets (when the buyer's currency exceeds her value); even so it could be evaded if a buyer simply misrepresents her value. The combination of the legitimate constraint $b \leq M$ and the fallacious constraint $M \leq v$ leads to their restriction $b \leq M \leq v$: bids must be no more than the buyer's value. Most importantly, this artefactual constraint is the basis of their argument that its enforcement is a form of market discipline, which in turn leads to their conclusions that "the convergence of price to equilibrium and the extraction of almost all the total surplus seem to be consequences of the double auction rules" (p. 135).

The discussion above demonstrates the nature of the problem with the Gode and Sunder model in the case that they consider, in which each agent only seeks to transact one unit. Problems multiply if a buyer has demand for more than one unit, or a seller can produce more than one unit. In the natural extension of their model, the budget constraint would be the sum of the buyer's unit values. Does the buyer's budget constraint restrict her bid on her first unit to a value below her first unit value, or only to the total budget? If it is the latter, bids on the earlier units are no longer constrained to be less than the

buyer's unit value, so we have moved into the realm of the ZI-U (unconstrained) agents. But these simulations do not achieve efficiency. So it must be the former case: the bid on each unit is restricted to the buyer's value on that unit. In the next section, we demonstrate that with a single unit, the constraint that Gode and Sunder impose is an individual rationality constraint.

Individual Rationality and Market Efficiency

We have now established that the constraints Gode and Sunder impose are neither budget constraints nor a legitimate form of market discipline. The next section demonstrates that their constraints conform to individual rationality constraints. Following Luce and Raiffa (1957, pp. 192 – 193), by *individual rationality* we mean that an agent only attempts to take part in a trade that increases, or at least leaves constant, his own utility. Their ZI-C simulations – in which their agents exhibit individual rationality – achieve Pareto optimal allocations; their unconstrained (ZI-U) traders do not reach Pareto optimal outcomes. The different performance of the ZI-C and ZI-U traders – combined with the fact that their constraints can only be meaningfully interpreted as aspects of trader behavior rather than budget constraints or as a form of market discipline – demonstrates that it is trader behavior that generates price convergence and efficiency in the CDA, not just the structure of the market rules.

AN EXCHANGE ECONOMY FORMULATION OF INDUCED COSTS AND VALUES

In this section we first construct a buyer's quasi-linear utility function from a vector of values and show that constrained maximization of the utility function yields a demand function that is dual to the vector of values. After that construction, we follow a similar procedure for sellers. In the concluding subsection we use these constructions to demonstrate that the constraints imposed by Gode and Sunder are in fact individual rationality constraints.

Induced values and quasi-linear utility

Each buyer $j \in J$ has a vector of values $v_j = (v_j^1, v_j^1, v_j^1, \dots, v_j^{n_j})$ for units of the commodity Y , where $v_j^1 \geq v_j^1 \geq v_j^1 \geq \dots \geq v_j^{n_j} \geq 0$. The total redemption value to buyer j when she purchases y units is

$$r_j(y) = \begin{cases} 0, & y = 0; \\ \sum_{\gamma=1}^{\min\{y, n_j\}} v_j^\gamma, & y = 1, 2, 3, \dots \end{cases}$$

We use the buyer's redemption value function $r_j(\cdot)$ to develop the buyer's quasi-linear utility function. Define the consumption space of buyer j as $X \times Y$, and let $(x_j, v_j) \in X \times Y$ denote the units of currency X and commodity Y held by buyer j . The utility function of buyer j is

$$u_j(x, y) = x + r_j(y) + M_j, \tag{1}$$

where M_j is a constant. (The constant M_j has no theoretical implications but is relevant to experimental studies. If buyer j has the initial endowment $(x_j^0, 0)$ and $M_j = -x_j^0$ then $u_j(x_j^0, 0) = 0$ so that the autarky outcome has payoff 0.) Equation (1) is linear in the currency (X) and additively separable in the currency and commodity, i.e., it is quasi-linear. The currency endowment $x_j^0 \geq \sum_{\gamma=1}^{n_j} v_j^\gamma$ is sufficient to guarantee that buyer j would be able to purchase each unit at any price at or below the value of the unit.

Our rationalization of the induced demand schedule as the solution to the constrained maximization of a quasilinear utility function is similar to the construction by Smith [1982, p. 932]. Smith derives the induced demand curve by maximizing the utility function $u_j(x, y) = x + r_j(y)$ subject to the budget constraint $x + py \leq 0$ where $x \geq 0$ and $y \geq 0$. In contrast, we define finite positive endowments of X for buyers and of Y for sellers that are consistent with the typical specification of consumer choice problems.

Theorem 1 The demand for Y by buyer j – derived from maximization of equation (1) for a sufficiently large endowment – is dual to v_j .

Proof The vector v_j of values is non-increasing, so that the total value function $r_j(y)$ is (weakly) concave for $y \in Y$. Therefore the utility function $u_j(x, y) = x + r_j(y) + M_j$ is (weakly) quasi-concave. The theorem of the maximum implies that for any given price p of good Y , the set of values that maximize $u_j(\cdot)$ is convex.

Let $y_j(p)$ be the demand of buyer j at price p , i.e., the solution to the maximization problem for $u_j(x, y)$. We complete the proof by showing that the demand $y_j(p)$ has the same graph as the vector v_j of value. If $p = v_j^k$, then $y_j(p) = \{k - 1, k\}$. If $p \in (v_j^{k-1}, v_j^k)$,

then $y_j(p) = k$. ■

Induced costs and quasi-linear utility

Seller i has a marginal cost schedule that is represented as a vector $c = (c^1, c^2, \dots, c^{m_j})$, where the commodity endowment for seller i is $y_i^0 = m_i$. Element c_i^k is the marginal cost incurred by seller i when he produces the k^{th} unit. The marginal cost for any unit beyond m_i is infinite. For $k \in \{0, 1, 2, \dots, m_j\}$ the redemption value for seller i when he sells k units is

$$r_i(k) = \begin{cases} 0, & k = 0; \\ -\sum_{l=1}^k c_l^i, & k = 1, 2, \dots, m_i. \end{cases}$$

We use the redemption value function of seller i to define his quasi-linear utility function as

$$u_i(x, y) = x = r_i(m_i - y), 0 \leq y \leq m_j. \quad (2)$$

Theorem 2 Seller i 's supply of Y – derived from maximization of equation (2) – is dual to c_i .

Proof The vector c_i of costs is non-decreasing, so that the total value function $r_i(k)$ is (weakly) concave for $k \in \{1, 2, \dots, m_i\}$, as is $r_i(m_i - y)$ for $y \in \{1, 2, \dots, m_i\}$. Therefore the utility function $u_i(x, y) = x = r_i(m_i - y)$ is (weakly) quasi-concave. The theorem of the maximum implies that for any given price p of good Y , the set of values that maximize $u_i(\cdot)$ is convex.

Let $y_i(p)$ be the supply of seller i at price p , i.e., the solution to the maximization problem for $u_i(x, y)$. We complete the proof by showing that the supply $y_i(p)$ has the

same graph as the vector c_i of costs. If $p = c^k$, then $y_i(p) \in \{k-1, k\}$. If $p \in (c^k, c^{k+1})$, then $y_i(p) = k$. ■

Individual rationality from the exchange economy perspective

The final step in our argument demonstrates that a buyer who only proposes bids or accepts asks that are at or below his unit value exhibits individual rationality, as does a seller who only proposes asks or accepts bids that are above his unit cost. We've carried out our utility function construction for buyers and sellers with one or more units to trade at the equilibrium price. In the ZI simulations each buyer has a positive value for one unit and each seller has only one unit available, so we apply our construction to this case.

If buyer j only submits bids $b_j \leq v_j$, this is equivalent to $x_j^0 - b_j + v_j + M_j \geq x_j^0 + M_j$.

For buyer j with the utility function $u_j(x, y)$ in equation (1) and endowment $(x_j^0, 0)$ of the currency and the consumption good, $r_j(1) = v_j$ and $r_j(0) = 0$ so

$x_j^0 - b_j + r_j(1) + M_j \geq x_j^0 + r_j(0) + M_j$. From the definition of the utility function in equation (1), this is equivalent to $u_j(x_j^0 - b_j, 1) \geq u_j(x_j^0, 0)$. With indivisible units of the consumption good as in the experiments, this is the individual rationality constraint.

The argument for seller i is similar. Suppose that seller i only considers asks $a_i \geq c_i$.

Assume that seller i has the utility function in equation (2). Since $r_i(0) = 0$ and $r_i(1) = -c_i$, the condition $a_i \geq c_i$ is equivalent to the condition $a_i + r_i(1 - 0) \geq r_i(1 - 1)$.

This is equivalent to $u_i(a_i, 0) \geq u_i(0, 1)$, which is the individual rationality constraint with indivisible units.

CONCLUSIONS

In Section 2.2 we showed that the constraints that Gode and Sunder impose cannot be interpreted as budget constraints; in Section 3.3 we showed that they are individual rationality constraints. Gode and Sunder demonstrate that (ZI-C) traders reach approximate Pareto optimal allocations, and that unconstrained (ZI-U) traders do not reach a Pareto optimal allocation. This establishes that individual rationality is both necessary and sufficient to reach a Pareto optimal allocation.

Hurwicz, Radner, and Reiter [1975] show that in any general equilibrium economy without externalities, random individually rational behavior leads to Pareto optimal allocations in a simple trading institution called the *B*-process. The *B*-process is defined generally enough so that the double auction with discrete units is a special case. With a discrete commodity space, as in a market experiment, random sequences of proposed trades submitted from each agent result in a sequence of net trades. An element of the trade sequence is non-zero if submitted proposals include a compatible trade (i.e., there is at least one trade proposal for which the net trade sums to zero). Hurwicz, Radner, and Reiter show under weak conditions on preferences and technologies that if at every iteration of the bargaining process, each individual only submits individually rational trade proposals, then the process converges to a Pareto optimal allocation in finite time. In effect, Gode and Sunder show that in a special case the *B*-process converges to a Pareto optimal allocation, although they either did not recognize or did not acknowledge that they had imposed an individual rationality constraint.

Price adjustment and convergence is a fundamental economic problem. Economists almost exclusively formulate their price theories as equilibrium models, but

price adjustment is by its nature a disequilibrium phenomenon. Substantial market disruptions, such as shocks to demand or supply, frequently disconnect price expectations from the actual market equilibrium price. Experiments have provided much insight into how trade activity realigns expectations and trading strategies of buyers and sellers with the market conditions that prevail after the disruption.

Practical problems of this sort are common. At a national scale, the closures of nine refineries in the aftermath of hurricane Katrina in August 2005 led to a 27% reduction in domestic oil production and a 21% reduction in U.S. oil refining capacity (Yergin [2006]), which destabilized prices of petroleum products for months. At a regional scale, almost every market was disrupted as infrastructure was destroyed and hundreds of thousands of people relocated. In neighboring cities, housing and labor market conditions were both greatly affected by the influx of displaced persons. Economists do not have theories that indicate how prices will stabilize after such a disruption. Even in more prosaic situations, such as a large expansion of an auto manufacturing plant in a medium sized city, equilibrium models do not offer much insight into the path of price adjustment in labor markets, housing markets, or even markets for goods and services.

Price paths during these adjustment periods will affect the profits that accrue to sellers, which in turn will affect the rate of capacity investment by firms. Changes to capacity will affect market supply and further impact price adjustment processes. This long-run market adjustment problem has been addressed experimentally, see Shachat and Zhang (2017), but the problem of dynamics with capital adjustment is an unsolved problem. This says nothing of the interactions across markets, that is, the general

equilibrium adjustment issues. ~~These are serious economic problems, and they call for serious economic models, even if those models require some work to develop and to understand.~~ Experimental economics and learning models are particularly well positioned to make substantive contributions to these issues, but first we must adopt models that demonstrate how prices converge in markets and how markets respond to disruptions.

REFERENCES

- Bosch-Domenech, A. & Silvestre, J. (1997). Credit Constraints in General Equilibrium: Experimental Results. *Economic Journal*, **107**, 1445 – 1464.
- Cason, T. N. & Friedman, D. (1996). Price formation in double auction markets. *Journal of Economic Dynamics and Control*, **20**, 1307 – 1337.
- Chamberlin, E. H. (1948). An Experimental Imperfect Market. *Journal of Political Economy*, **56**, 95 – 108.
- Conlisk, J. (1996). Why Bounded Rationality? *Journal of Economic Literature*, **XXXIV**, 669 – 700.
- Easley, D. & Ledyard, J. O. (1993). Theories of Price Formation and Exchange in Double Oral Auctions. In D. Friedman & J. Rust (Eds.), *The Double Auction Market: Institutions, Theories, and Evidence* (pp. 63 – 97). Reading, MA: Addison-Wesley.
- Friedman, D. (1991). A Simple Testable Model of Price Formation in the Double Auction Market. *Journal of Economic Behavior and Organization*, **15**, 47 – 70.
- Evans, D. A. (1997). The Role of Markets in Reducing Expected Utility Violations. *Journal of Political Economy*, **105**, 622 – 636.
- Gjerstad, S. (2007). The competitive market paradox. *Journal of Economic Dynamics and Control*, **31**, pp. 1753 – 1780.
- Gjerstad, S. and Dickhaut, J. (1998). Price Formation in Double Auctions. *Games and Economic Behavior*, **22**, 1 – 29.
- Gode, D. K. & Sunder, S. (1993). Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy*, **101**, 119 – 137.
- Hayek, F. A. (1945). The Use of Knowledge in Society. *American Economic Review*, **35**, 519 – 530.
- Hurwicz, L., Radner, R. & Reiter, S. (1975). A Stochastic Decentralized Resource Allocation Process: Part I. *Econometrica*, **43**, 187 – 221.
- List, J. (2004). Testing Neoclassical Competitive Theory in Multilateral Decentralized Markets. *Journal of Political Economy*, **112**, 1131 – 1156.
- Luce, R. D. & Raiffa, H. (1957). *Games and Decisions*. New York: Wiley.

- Markose, S., Arifovic, J. & Sunder, S. (2007). Editorial: Advances in experimental and agent-based modelling: Asset markets, economic networks, computational mechanism design and evolutionary game dynamics. *Journal of Economic Dynamics and Control*, **31**, 1801 – 1807.
- McAfee, R. P. & McMillan, J. (1996). Competition and Game Theory. *Journal of Marketing Research*, **33**, 263 – 267.
- Samuelson, L. (2005). Economic Theory and Experimental Economics. *Journal of Economic Literature*, **XLIII**, 65 – 107.
- Shachat, J., & Zhang, Z. (2017). The Hayek hypothesis and long-run competitive equilibrium: An experimental investigation. *The Economic Journal*, **127**, 199-228.
- Smith, V. L. (1962). An Experimental Study of Competitive Market Behavior. *Journal of Political Economy*, **LXX**, 111 – 137.
- Smith, V. L. (1982). Microeconomic Systems as an Experimental Science. *American Economic Review*, **72**, pp. 923 – 955.
- Tesfatsion, L. (2002). Agent-Based Computational Economics: Growing Economies from the Bottom Up. *Artificial Life*, **8**, 55 – 82.
- Wilson, R. B. (1987). On Equilibria of Bid-Ask Markets. In G. W. Feiwel (Ed.), *Arrow and the Ascent of Modern Economic Theory* (pp. 375 – 414). New York: New York University Press.
- Yergin, D. (2006). Ensuring Energy Security. *Foreign Affairs* **85**, 69 – 82.

FIGURE CAPTIONS

Fig. 1: Supply, demand, and trade prices from models and experiments.