



UvA-DARE (Digital Academic Repository)

Understanding trading at all prices

Ruiter, A.G.J.M.

Publication date
2019

[Link to publication](#)

Citation for published version (APA):

Ruiter, A. G. J. M. (Accepted/In press). *Understanding trading at all prices*. Paper presented at Barcelona GSE Summer Forum, Barcelona, Spain.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Understanding trading at all prices, v6

Aad Ruiter^{*†}

April 23, 2019

Abstract

This paper argues that we should understand human trading behavior as an equilibrium of heuristics. Demonstrating such an equilibrium requires that we apply agent-based modeling to experimental data. The steps include (i) simulating the successive sets of feasible actions from which human subjects select preferred actions; (ii) determining which heuristics predict the data best (both individual actions and aggregate patterns); and, (iii) determining whether these heuristics are ecologically rational. For trading at all prices in experimental Scarf economies, no configuration of heuristics has yet been identified that can explain all observed aggregate patterns. Issues and a research agenda are discussed.

KEYWORDS: equilibrium of heuristics, experimental economics, agent-based modeling, Scarf

JEL codes: C63, D44, D58, D83, D90

1 Introduction

How do economies achieve equilibrium (if at all)? Experimental economics provides valuable data with respect to price discovery and equilibration (e.g. Anderson et al. (2004); Crockett et al. (2011); Goeree and Lindsay (2016)). In order to make the most of these data, we need agent-based modeling. Duffy (2006) urges agent-based modeling to do more than just reproduce aggregate patterns that can be observed in the laboratory or in reality. Indeed, we should try to “get into the heads” of individual decision-makers. For that, agent-based modeling needs a notion of equilibrium. I will argue that an equilibrium of heuristics is the appropriate choice. From this perspective, it is possible to align economic theory, experiments and agent-based modeling. This paper describes how such an approach can be applied, using results of my thesis as an illustration, c.f. Ruiter (2017).

The analysis centers on the examples of Scarf (1960). These economies are famous for showing that the so-called tâtonnement process does not necessarily

^{*}University of Amsterdam, CenDef, e-mail: a.g.j.m.ruiter@uva.nl.

[†]I would like to thank Cars Hommes for helpful comments.

converge. Since then, they have become a benchmark for proposals regarding price adjustment processes and for market mechanism design (c.f. Oehmke and Oehmke (1991); Kumar and Shubik (2004); Gintis (2007); Goeree and Lindsay (2016); Ruiter (2017)).

Anderson et al. (2004) reports results of experimental trading in these Scarf economies. Its experiments are a major step forward compared to earlier experiments with a single financial market, in which all traders have exogenously given reservation prices (e.g. Smith (1962), Gode and Sunder (1993), Cliff and Bruten (1997)). Anderson et al. consider two markets, and they allow trading at non-equilibrium prices. As a result, the equilibrium is able to shift. Furthermore, they designate one commodity as money, which (theoretically at least) makes it more difficult to achieve convergence. Nevertheless, the results of Anderson et al. (2004) are in line with tâtonnement theory. Achieving convergence in the stable Scarf economy is the most significant result: prices are close to their equilibrium values, and so is the allocation.

Ruiter (2017) analyzes these results with data that were kindly provided by prof. Anderson.¹ My simulation platform, called FACTS (short for Fallible Agents' Commodity Trading System), is especially geared to analyzing individual behavior.²

This paper is organized as follows. Section 2 adapts a proposal of Hahn (1984) so as to become an equilibrium of heuristics. The experiments of Anderson et al. are described in section 3. Next, section 4 presents the core of a stylized model of human trading that is implemented in FACTS (section 5). The analysis of the data of Anderson et al. can be found in section 6. Finally, section 7 sketches an agenda for improving our understanding of human trading behavior.

2 On the notion of equilibrium

We cannot characterize an economy as being out of equilibrium, unless we have a good notion of equilibrium. Consider Arthur's discussion of the El Farol problem, c.f. Arthur (2006).³ After showing that homogenous expectations fail to solve this problem, Arthur claims that a generative approach (i.e. agent-based

¹The raw data of a stable and of a counter-clockwise treatment together comprise 9799 valid individual decisions.

²Fallibility refers to traders not being the perfectly rational agents of economic textbooks. The data of Anderson et al. (2004) provide nice illustrations of fallibility. For instance, in the unstable markets sophisticated rational agents could have deduced the equilibrium prices based on introspection alone, i.e. without having to trade. Nevertheless, in none of the sessions with an unstable Scarf economy did human subjects achieve the equilibrium. There are also good indications that the subjects ignored the excess of buyers over sellers in the experimental stable Scarf economy. See, Ruiter (2017).

³The El Farol problem can be described as follows. One hundred people decide each week whether or not to go the El Farol bar on Thursday evening. Since this bar is rather small, they prefer to stay at home if they expect more than 60 people to show up. If they expect that there will be fewer people then they prefer to go to the bar rather than stay at home. Decisions are taken independently and simultaneously, c.f. Arthur (1994).

modeling) resolves the anomaly. His view is worth quoting in full:

'Putting this into practice in the case of El Farol means assuming that agents individually form a number of predictive hypotheses or models, and each week act on their currently most accurate one. (Call this their active predictor.) In this way beliefs or hypotheses compete for use in an ecology these beliefs create. Computer simulation then shows that the mean attendance quickly converges to 60. In fact, the predictors self-organize into an equilibrium pattern or "ecology" in which, on average, 40% of the active predictors are forecasting above 60 and 60% below 60. And while the population of active predictors splits into this 60/40 average ratio, it keeps changing in membership forever. There is a strong equilibrium here, but it emerges ecologically and it is not the outcome of deductive reasoning.', Arthur (2006, p. 1561).

The mean attendance may well be the phenomenon we seek to explain, but economic theorists would not necessarily agree with Arthur that the generative solution is an equilibrium (as distinct from a stable state). If people change their plans, they are not in equilibrium (i.e. able to implement their intended actions); and the economy as a whole cannot be in equilibrium unless all individual agents are in equilibrium. If we would forsake equilibria and focus on stable states instead, then a lot will be lost. As a case in point, the whole edifice of general equilibrium theory is obtained through the characterization of equilibria. Hence, we should look for an appropriate definition of equilibrium.

As early as 1974, in his inaugural lecture, Frank Hahn proposed an alternative direction that begs to be embraced by agent-based modeling, c.f. Hahn (1984).^{4,5} Ignoring technicalities, Hahn conceived of economies as systems that generate signals. Agents map these messages to acts, and this mapping (or policy, as Hahn calls it) depends on their theory of the economy. Hahn proposed to define equilibrium as states in which the economy 'generates messages which do not cause agents to change the theories which they hold or the policies which they pursue.', c.f. Hahn (1984). In such an equilibrium 'agents have adapted themselves to their economic environment and where their expectations in the widest sense are in the proper meaning not falsified'. In such an equilibrium all theories and policies are independent of the date (i.e. agents no longer learn). An agent may be wrong, 'but as long as (...) the method by which he makes his forecasts is the same at all dates he will not be learning in my sense', c.f. Hahn (1984).

The complexity of the goal set by Hahn (1984) depends on the level of detail of policies. Hahn, continuing in the tradition of general equilibrium theory,

⁴This is ironic, because Hahn was a pure theorist. Years later, he lamented 'Instead of theorems we shall need simulations, instead of simple transparent axioms there looms the likelihood of psychological, sociological and historical postulates. These new roads will find willing and happy travellers, but it is unlikely that those with the temperament and facilities of mid-twentieth-century theorists will find this a congenial road', Hahn (1991).

⁵One can argue that Alchian (1950) and Becker (1962) anticipate a similar type of equilibrium.

aimed for comprehensive policies. For instance, consumers are supposed to plan which combinations of commodities to buy. Here, a combination of commodities refers to the complete list of *all* different commodities available in different places at different times and under different circumstances. In Hahn's proposal these comprehensive plans must be updated continuously, conditional on new messages received. Hahn acknowledged that his approach makes excessive demands on the rationality and the computational ability of the agent.

While agent-based modeling may go a long way in implementing his idea, I propose to apply one critical change to Hahn's concept, based on the following consideration. How do economic agents evaluate their past performance? Do they consider each and every single trade, and ask themselves if (with hindsight) they could have done better at that level? If so, then they should maximize utility (or profits) as is usual in general equilibrium theory. However, I doubt that this is how people decide and learn. More likely, they choose by applying (combinations of) heuristics, and they improve themselves by determining whether or not they should have used better heuristics. These heuristics may refer to expectation formation, searching for feasible options, choice from a set of alternative actions, bargaining, et cetera. If people are satisfied with the heuristics they did apply, then they will accept the economic outcome, even it falls short of the comprehensive standard.

Generally speaking, economists assume that people try to further their own interest. That translates to people using heuristics that are ecologically rational, i.e. well-adapted to the (institutional) environment, c.f. Smith (2008); Gigerenzer and Pachur (2011). Therefore, rather than an equilibrium of plans as in Debreu (1959), or an equilibrium of plans and price expectations as in Radner (1972), we need an equilibrium of ecologically rational heuristics.⁶ In this sense, indeed, the generative solution of the El Farol problem can be an equilibrium.

It will be clear that the heuristics of individual agents will determine aggregate patterns (e.g. in prices), which in turn will determine the success of specific heuristics. In order to get a handle on the deep dependencies and on heterogeneity we need agent-based modeling. For a survey of heterogeneity in agent-based models, see Arifovic and Duffy (2018). Discriminating between alternative behavioral hypotheses is greatly facilitated by the sensitivity of simulated aggregate patterns to the specification of heuristics at the micro level. Moreover, if we adopt the notion of an equilibrium of ecologically rational heuristics, then this constrains heterogeneity at the micro level.⁷ Indeed, emergent aggregate

⁶I will not attempt to put forward a positive, formal definition of an equilibrium of heuristics - maybe it will emerge over time. Admittedly, the concept raises serious methodological issues: Establishing ecological rationality requires a universe of heuristics; what if that universe is too small? Different aspiration levels may play a role when people try to assess whether they could have done better. Less demanding individuals may settle for other heuristics. How do people learn heuristics? How do they learn what is ecologically rational? When is the similarity between simulated and observed aggregate patterns close enough? However this may be, I believe that the notion of an equilibrium of heuristics can be beneficial for economic theory, experimental economics, and for agent-based modeling.

⁷Ecological rationality provides the best check on the "wilderness of bounded rationality", c.f. Sargent (1993); and similarly on heterogeneity.

patterns and sustainable diversity can be considered as the macro foundations of microeconomics. Agent-based modeling can expose the micro-macro relations that prevail in a given context.

3 Experimental trading in the Scarf economies

The law of demand and supply maintains that prices must rise (fall) if there is excess demand (supply). The so-called tâtonnement price adjustment process implements this dynamics, while postponing trade until the equilibrium prices have been determined.⁸ Scarf (1960) provides three examples of small exchange economies with three traders and three commodities. These examples differ in one respect only: how the commodities are initially allocated. This is sufficient for generating strikingly different price dynamics with tâtonnement. In the stable example, prices converge to an equilibrium. The other two examples are unstable; here, prices orbit around their equilibrium values, either in a clockwise or counter-clockwise direction.

Anderson et al. (2004) reports results of experimental trading in these Scarf economies. The original economies were replicated five times for having fifteen traders. In order to make trading more challenging for human subjects, the total amounts of the goods were altered so as to avoid symmetry in the equilibrium prices. Furthermore, one commodity was designated money; the other two commodities had to be traded in exchange for money. Trading took place in a continuous double auction. Here, both buyers and sellers can submit offers at any time and in any market they like. An exchange occurs if a buyer accepts the best ask, or if a seller accepts the best bid. That is, in these experiments trading took place at non-equilibrium (or 'all') prices. This means that traders ran the risk of paying too much or receiving too little. However, without trading they could not increase their utility level and they could not learn.

Anderson et al. found results similar to Scarf (1960). For instance, in the stable economy the prices were close to their respective equilibrium values, c.f. figure 1. In the unstable examples, there was low frequency orbiting of sorts, in opposite directions. This can be illustrated by the so-called "clock-hand model". Let (p_2^t, p_3^t) be the last known trading prices at time t in the markets for commodities 2 and 3, and let (p_2^*, p_3^*) be the equilibrium prices.^{9,10} Furthermore, let the angle between successive observations be defined as $\angle (p_2^t, p_3^t) - (p_2^*, p_3^*) - (p_2^{t+1}, p_3^{t+1})$. The sides of these angles can be

⁸The reason for postponing trade is preventing that the equilibrium will shift.

⁹Commodity 1 is the numéraire; hence $p_1 = 1$.

¹⁰Transaction prices in the two markets arrive in an arbitrary order, e.g. $p_2^1, p_2^2, p_3^3, p_2^4, \dots$ and so on (superscripts refer to the passage of time). These data can be converted into a time series of last known prices, $(\mathbf{p}^t)_t$, by successively updating elements:

$$\begin{pmatrix} 1 \\ n.a. \\ n.a. \end{pmatrix}, \begin{pmatrix} 1 \\ p_2^1 \\ n.a. \end{pmatrix}, \begin{pmatrix} 1 \\ p_2^2 \\ n.a. \end{pmatrix}, \begin{pmatrix} 1 \\ p_2^2 \\ p_3^3 \end{pmatrix}, \begin{pmatrix} 1 \\ p_2^4 \\ p_3^3 \end{pmatrix}, \dots$$

where n.a. means 'not available'.




Figure 1: *Experimental trading prices in the stable Scarf economy are near their Walrasian equilibrium values (commodity 2: 40, commodity 3: 20; commodity 1 is the numéraire) for some time.*

interpreted as the position of the hand of a clock at times t and $t + 1$, with the hand centered on the equilibrium prices. If there is systematic orbiting in the clockwise (counter-clockwise) direction, then these angles will tend to be positive (negative) and so will their sum. The evolution of this sum is shown in figure 2.

Apart from convergence in the stable economy, and orbiting in the unstable examples, there are other aggregate patterns that can be used to test behavioral hypotheses against. A single experimental session was divided into a number of periods. Remarkably, end-of-period allocations were quite close to the equilibrium allocation (even in the unstable economies, see figure 6 below). Trading prices were close to expected prices at the time that the trades occurred. This suggests that expected prices served as the basis for reservation prices. Typically, small quantities were traded. One market (the one with the greater turnover) had less uncertainty than the other (as measured by the entropy of the distribution of prices). This translated into greater quantities proposed and accepted. The size of the quantities seems to reflect the confidence in price expectations. There were strategic offers, (e.g. buyers submitting asks aimed at lowering the selling price), and arbitrage offers (e.g. traders buying commodities that did not contribute to their utility).¹¹ Identifiable strategic and arbitrage offers amounted to about 10% of all offers; during the practice period there

¹¹For instance, traders of type 1 only derive utility from commodities 1 and 3. Traders of type 2 (3) prefer commodities 1 and 2 (2 and 3).

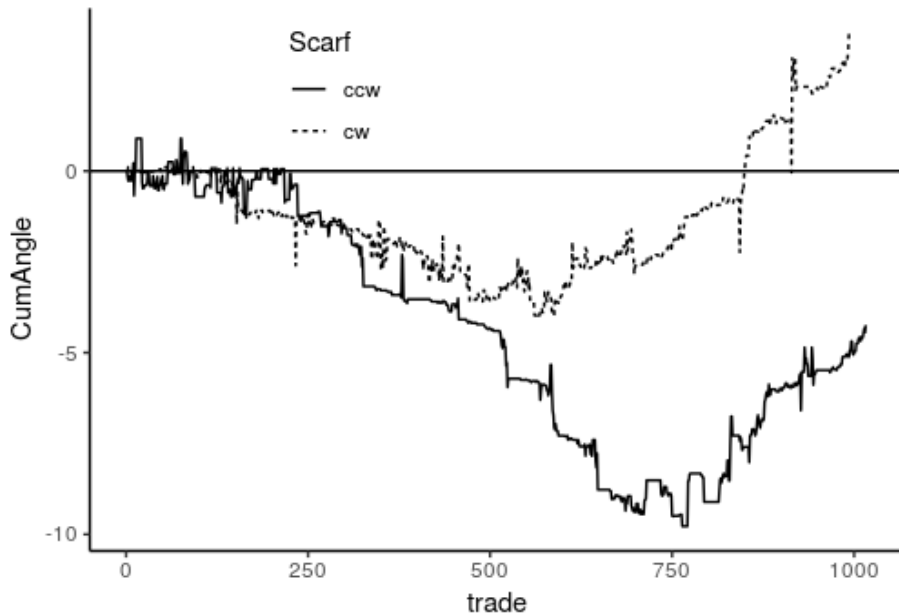


Figure 2: Cumulative angles in a counter-clockwise ('ccw') and in a clockwise session ('cw'). Orbiting in these sessions (labeled by Anderson et al. as 511 and 420 respectively) is not as articulate as it is in Hirota et al. (2005) or Goeree and Lindsay (2016), but eventually the directions agree with the predictions of tâtonnement theory.

was substantially more arbitrage than in the periods that earned the subjects financial rewards.

4 A stylized model of trading behavior

Following choice theory, traders are assumed to choose the best alternative from a set of perceived, feasible actions.

- This set depends on (i) the type of market (the market mechanism largely determines transparency; the market protocol may preclude short-selling); (ii) the currently available best offers (which may be accepted, or possibly improved upon); (iii) price expectations (they determine whether pending offers are favorable or not, and hence whether they should be canceled or accepted and whether they can be improved); (iv) endowments (in the experiments of Anderson et al., traders need money before they can buy other commodities; also, if a trader has acquired some of the commodities he prefers, then he can submit strategic asks to move the selling price); risk attitude (if a trader wants his utility-level to be non-decreasing, then he should ignore opportunities that involve a commodity that is not currently

constraining his level of utility;¹² if a trader only wants to buy what he needs and sell what he can miss, then this precludes arbitrage). Of these factors, pending offers and price expectations are the most important, because they vary the most during the course of trading.

- Ranking the feasible opportunities can be done in different ways. If price expectations are beliefs (i.e. probability distributions), then opportunities can be represented by lotteries and traders can maximize expected utility, or expected value (as in cumulative prospect theory).¹³ Alternatively, rules of thumb can be applied: e.g. cancel a pending offer with an unfavorable price before accepting an offer with a favorable price; accept a favorable offer before submitting another proposal; submit a regular offer before a strategic one.¹⁴ These rules can be applied in conjunction with beliefs and with price expectations as point expectations.

This stylized model needs further elaboration before it can be implemented in an agent-based model. The point here, however, is that such an agent-based model, in combination with experimental data, allows us to simulate the sets of perceived, feasible actions at every stage of the experiment, giving us a very deep insight into individual decisions.

5 FACTS

In order to analyze human trading behavior, I have designed a platform for simulating exchange, called FACTS. It consists of a standard exchange economy, with an embedded market mechanism (a continuous double auction or a clearing house). Individual agents propose prices, while an auctioneer facilitates the trading process (by managing the orderbook, communicating changes therein, keeping track of ownership of the commodities, and by calculating statistics). An offer is a message with which an agent signals his willingness to buy or sell a specific quantity of a particular commodity at a certain price. Each message has a date, that refers to the time at which the contents of the message will reach the recipient. A message queue automatically sorts incoming messages according to this date, c.f. figure 3.

FACTS can run in two modes. It can be conditioned on a specific experiment. This means that beforehand, a human subject experiment is loaded into the message queue. In this mode, FACTS compares the moves of human subjects and of the robot players that are matched to them. After that, the human

¹²This, by the way, is not a very plausible behavioral hypothesis: in the stable Scarf economy of Anderson et al., 14% of the submitted offers would have lowered the current utility level of the proposer, if accepted.

¹³The valuation certain actions can be difficult. For instance, submitting an offer in one market while owning a pending offer with an unfavorable price in another market. Or, submitting a strategic offer (which can have the unintended consequence of being accepted by another trader).

¹⁴These are the rules of thumb reported on in table 1. Traders choose the market that has the most preferred type of action, with ties broken randomly. See Ruiter (2017) for details.

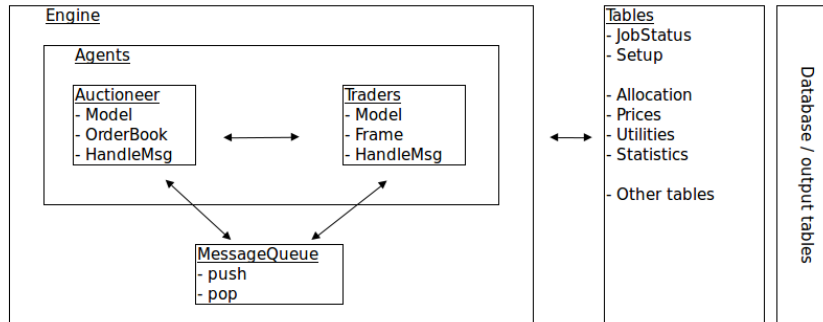


Figure 3: *Design of FACTS. The exchange process, which consists of the auctioneer publishing bid / ask spreads and traders responding with new offers, is communicated through the message queue. When the auctioneer calculates statistics he may contact traders directly. Logical tables hide the database and / or output files.*

move is implemented; this is to assure that robot traders at every point in the game have the same public information as human players did in the original experiment. In unconditional mode, trading follows the moves proposed by robot instead of human traders. In this case, a message from the auctioneer about the latest bid / ask spreads triggers each trader to take a new decision. Agents, who want to propose a new offer, send a message to themselves through the message queue. These messages serve as the actual trigger for submitting a new offer. The trader who is quickest to respond is selected for submitting the next offer; the "notes to self" from other, slower traders are simply discarded. This mechanism allows FACTS to emulate asynchronous trading

The heuristics that determine trading behavior are collected in frames. Individual traders can have multiple frames, and hence display different types of trading behavior. In addition to frames, traders have a model of their environment and an interface, c.f. figure 4.

6 Analysis of the experimental data

In order to understand trading at all prices, we seek configurations of heuristics that best predict the actual offers of individual human traders, reproduce aggregate patterns, and that are ecologically rational. For this we evaluate whether simulated sets of perceived, feasible actions cover actual decisions of human traders; whether the proposed course of action of robot and human traders coincide; whether robot traders achieve convergence and orbiting in the stable and unstable Scarf economies respectively; whether simulated and observed end-of-period allocations are close; and, whether learning favors configurations that can explain observed patterns.

With respect to choice from a set of feasible actions, three candidates are




Figure 4: Architecture of a trader. When agents receive a message, they interpret it: queries from the auctioneer or informative messages are dispatched to their internal model; messages that trigger a decision are assigned to the frame they currently deem best.

considered: expected utility maximization and expected value maximization (as in cumulative prospect theory) against beliefs, and also rules of thumb.¹⁵ In both the stable and the counter-clockwise economy, the rules of thumb predict human actions best, c.f. table 1. As it turns out, they are also ecologically rational compared to expected utility and expected value maximization, c.f. table 2.¹⁶ It is interesting that expected utility maximization survives in the unstable treatments while cumulative prospect theory does not. The latter better predicts human actions in the counter clockwise economy, c.f. table 1.

FACTS implements eleven algorithms (plus variations) for generating reservation prices. Among them are Zero Intelligence (ZI, random actions and/or random reservation prices, c.f. Gode and Sunder (1993)), Zero Intelligence Plus (ZIP, updating reservation prices so as to remain competitive, c.f. Cliff and Bruten (1997)), Adaptive-Aggressive (AA, managing a trade-off between adaptive and aggressive reservation prices, c.f. Vytelingum (2006)) and different applications of Gjerstad-Dickhaut beliefs (GDW, beliefs expressing the proba-

¹⁵Ruiter (2017) also studies choice that is sensitive to complexity. Entropy-sensitive utility functions can be described as $v(\mathcal{L}) = \sum_k u(x_k) p_k + \rho H(\mathbf{p})$, with $\mathcal{L} = (\mathbf{x}; \mathbf{p})$ a lottery, u a von Neuman-Morgenstern utility function, and $H(\mathbf{p})$ the Shannon entropy of probability distribution \mathbf{p} . This functional form satisfies Savage's axioms, provided that lotteries are simplified to the point where all prizes x_i are different. Entropy-sensitive preferences express that agents care about *both* the expected utility *and* the complexity of the lotteries they can choose from. This provides an alternative explanation for the Allais paradox, the common ratio paradox, probabilistic insurance aversion and for the Ellsberg paradox. Maximization of expected entropy-sensitive utility performs slightly better in predicting observed decisions than maximization of regular von Neumann-Morgenstern expected utility.

¹⁶Learning in FACTS is a mixture of reinforcement learning and replicator dynamics. After each simulation run, traders can switch between heuristics.

Table 1: Correct prediction of human decisions

action	stable				counter-clockwise			
	<i>N</i>	RoTh	EU	CPT	<i>N</i>	RoTh	EU	CPT
propose bid good 2	732	65	56	68	696	71	55	57
propose bid good 3	554	64	48	63	681	49	41	54
propose ask good 2	761	57	34	40	731	37	25	26
propose ask good 3	966	68	32	46	1043	69	40	53
accept bid good 2	188	100	85	55	113	95	84	67
accept bid good 3	80	94	70	42	203	97	98	98
accept ask good 2	407	94	88	50	251	100	100	91
accept ask good 3	1011	93	76	24	474	80	86	72
cancel bid good 2	26	100	5	31	55	100	62	79
cancel bid good 3	7	100	50	50	16	100	40	60
cancel ask good 2	26	100	0	67	32	100	60	60
cancel ask good 3	78	100	94	98	44	100	63	88
total	4836	70	49	49	4339	66	49	56

Simulated prediction rates (%), by algorithm, conditional on human actions recognized as being feasible (and on eGD-expectations, see below). A value of 100 indicates perfect predictions. Figures are averages over 1,000 runs per heuristic. The total number of actions is listed in column 'N'; actions that can be interpreted as arbitrage have been excluded. Rules of thumb (column 'RoTh', see section 4) predict human actions best. Columns 'EU' and 'CPT' refer to expected utility maximization and expected value maximization (as in cumulative prospect theory) respectively. Expected value maximization better predicts than expected utility maximization in both the stable and the counter-clockwise treatment. This is mainly due to transforming probabilities into decision weights: replacing utility by the value function while retaining probabilities would give a score similar to EU; combining the utility function with decision weights, however, causes the score to increase.

bility that a proposed price will be accepted; eGD, a point expectation based on these beliefs, c.f. Gjerstad and Dickhaut (1998)).^{17,18} Beliefs not only admit maximization of expected utility or expected value when choosing from a given

¹⁷Gjerstad and Dickhaut (1998) defines beliefs based on a short history of the last five bids and asks. The data of Anderson et al., however, have a peculiar feature. Apparently, it was easier for traders to submit n acceptances of one unit rather than one acceptance of n units. This can trick robots with a limited memory into believing that the spread of prices is zero. That is why price expectations as distributions have been constructed by tabulating all observations. This has the additional advantage that Gjerstad-Dickhaut traders can consider a trade-off between proposing and waiting. The drawback, however, is that expectations become increasingly less sensitive to new observations, see also section 7.

¹⁸Another algorithm, that derives reservation prices from a utility target, has the distinct

Table 2: Learning a heuristic for selecting the best action

treatment	RoTh	EU	CPT	length
stable	15	0	0	750
ccw	13	2	0	854
cw	13	2	0	1178

Final distributions of traders (15 in total) over rules for selecting a best alternative and the length of the simulations. The initial distribution was the same for all treatments; care has been taken to vary behavior across trader types. The rules of thumb clearly emerge as dominant.

Table 3: Recognition and prediction of human moves (%)

Description	stable		counter-clockwise	
	recognized	predicted	recognized	predicted
ZI	37.7	53.7	41.8	57.8
ZIP	76.5	67.1	76.9	69.8
AA	74.5	65.0	75.2	67.8
eGD	75.5	68.0	73.2	62.5
GDW	77.0	51.8	68.9	51.7

Average percentages of human actions that are recognized as a feasible option and that are correctly predicted, as simulated over 1,000 runs. Actions exclude arbitrage. Rates are fairly stable across the stable and counter-clockwise treatments. ZIP is consistently good. Utility maximization against Gjerstad-Dickhaut beliefs poorly predicts human moves.

set of feasible options. Traders with beliefs can also select an optimal buying or a selling price by maximizing expected utility or expected value against their beliefs.

Table 3 shows that this approach performs poorly: GDW does worse than ZI-trading. This finding is due to the fact that expected utility mainly increases because of the likelihood that an offer will be accepted; hence buyers propose high prices and sellers propose low prices. This is not descriptive of human trading (nor is ZI-trading). Prediction percentages in table 3 were calculated assuming that all algorithms (except ZI) use rules of thumb to select a best option; decisions identified as arbitrage have been excluded. Here, the combination eGD / rules of thumb predicts 68% of all decisions correctly.¹⁹

Convergence can be measured by the concentration of prices: there is convergence if the percentage of prices close to the average of a run is sufficiently

advantage that traders perceive markets as being interconnected.

¹⁹In table 1 this combination achieves 70%, but that is relative to decisions in which the human action was recognized as feasible.

Table 4: Assessing the similarity of price formation

description	avg distance (money)			confidence intervals (%)					
	stable	ccw	total	stable		ccw		cw	
				good 2	good 3	good 2	good 3	good 2	good 3
humans				74	78	17	17	17	26
ZI	9.7	9.5	9.5	6-9	6-9	6-10	6-10	6-10	6-10
eGD	5.8	6.4	6.2	87-93	95-98	3-7	14-83	3-84	3-76
ZIP	6.2	6.7	6.5	16-99	14-73	5-28	3-84	4-95	6-38
AA	6.1	6.7	6.4	11-49	10-57	2-6	0-6	0-35	2-8
GDW	8.2	8.8	8.5	8-38	11-39	22-45	2-44	3-43	16-46

Statistics are averages, calculated over 1,000 runs. Here, eGD prices are on average closer to the observed experimental prices. The confidence interval for concentration of eGD ranges from 87% to 93% in the stable Scarf economy, implying a very high probability of convergence. As a matter of fact, convergence is too strong compared with human traders. Other algorithms have a lower likelihood of achieving convergence in the stable Scarf economy. In the unstable economies (counter-clockwise, 'ccw' and clockwise, 'cw') confidence intervals are wide or small at low levels of concentration, indicating that price formation is either not very robust or not converging. Here, eGD is likely to show too much convergence. The ZIP-algorithm covers most human concentration statistics, but its intervals are generally quite wide.

high. Table 4 defines “close” as within a margin of $\pm 10\%$ of the average of a run. It compares simulated confidence intervals for the concentration measure with the observed levels of concentration in the experiments of Anderson et al.. Only eGD consistently achieves convergence in the stable Scarf economy (because the confidence interval is small and situated in the high percentages).²⁰ However, eGD also exhibits too much convergence in the unstable examples. This is due to reservation prices becoming increasingly insensitive to new observations, c.f. footnote 17.

As explained above, orbiting can be measured by accumulating the angles of between successive observations of last known trading prices with respect to the equilibrium prices. If there is systematic orbiting, the sum of these angles will be positive (clockwise orbiting) or negative (counter-clockwise orbiting). Figure 5 shows several simulated densities of these angles. ZIP-traders reproduce the observed orbiting of prices in human trading: not only is the direction correct, but the prices exhibit a pattern with a low frequency fluctuation. As a matter of fact, the density is due to one price steadily increasing over the course of a session. This suggests that orbiting in robot trading is due to the absence of a negative feedback mechanism.²¹ In experimental human trading there is

²⁰The number of 1,000 simulations is more than sufficient for having stable confidence intervals, c.f. Ruiter (2017)

²¹In the examples of Herbert Scarf, on the other hand, orbiting is the expression of a negative




Figure 5: *Simulated probability densities of cumulative angles in the unstable Scarf economies. The ZI-densities (left) are symmetric around zero, showing that (i) there is no inherent bias in the system and (ii) that orbiting is due to expectation formation. ZIP-trading (middle) exhibits systematic orbiting because the densities have shifted away from zero; furthermore, they each have shifted in the right direction; the eGD-densities (right) do not agree the experiments of Anderson et al.: in the clockwise economy values are highly concentrated around zero, implying absence of orbiting; if there is orbiting in the counter-clockwise economy then it occurs in the wrong direction.*

also one price that increases steadily over several periods, but at some point its climb is reversed.

Figure 6 shows the relative shares of different types of traders in the total amount of each commodity at the end of a period. In every triangle, the distances from each side to an arbitrary point in the interior of the triangle always add up to one. These distances represent the shares of the different types of traders in the total amount of a particular commodity. If all is owned by traders of the same type then the allocation is in a corner of the triangle. In the equilibrium, dots lie on midpoints of the sides: commodity 1 on the left side of the triangle; commodity 2 on the bottom and commodity 3 on the right side of the triangle.

Human traders achieve allocations that are quite close to the equilibrium, even in the unstable counter-clockwise economy. The allocations of eGD-traders, on the other hand, are further removed from the equilibrium. This is a reflection of there currently being not enough transactions in robot trading. Here, eGD-trading is taken as an example, because its allocations are more concentrated than the allocations of ZI-, ZIP-, AA- and GWD-traders.

There have been multiple experiments in order to determine which algorithms for setting reservation prices are preferred by traders that can learn. Table 5 gives the results of a pairwise competition of ZIP and AA against eGD.²² The latter is relatively strong compared to alternative algorithms (also to others not discussed here), but in the stable Scarf economy it is dominated by ZIP.

feedback mechanism.

²²Here, I do not discuss heterogeneity, because I want to emphasize that no configuration has yet been found that is consistent with observable aggregate patterns regarding convergence. Ruiters (2017) demonstrates heterogeneity at the micro level in terms of entropy-sensitivity (c.f. footnote 15) and also that diversity in this sense can be ecologically rational.




Figure 6: *Shares per trader type in commodities. Each dot represents an allocation at the end of a period. Rows 1 / 2 show human trading and rows 3 / 4 show eGD-trading in the stable and counter-clockwise economy respectively. Human trading comes quite close to the equilibrium allocation per trader type. Commodity 1 (money) is more dispersed than the other commodities, because these can only be traded in exchange for money.*

Table 5: Pairwise competition with eGD

expectation	final frequency			length		
	stable	ccw	cw	stable	ccw	cw
ZIP	15	5	9	750	753	1,000
AA	0	5	4	750	765	1,000

Final frequencies of traders adopting a heuristics other than eGD, plus the lengths of the simulations (max = 1,000 iterations). The initial distribution is the same for all three examples; care has been taken to vary behavior across trader types. If five traders adopt ZIP in the counter-clockwise example, this means that 10 traders adopt eGD reservation prices.

The fact that ZIP dominates eGD in the stable Scarf economy implies that human trading has not yet been properly explained. ZIP-traders do not achieve convergence in the stable example, while eGD-traders do. The dominance of ZIP seems to be due to ZIP traders generating more transactions, which raises their level of utility.²³ If all robot traders in FACTS would generate more transactions, this comparative advantage of ZIP-traders may diminish. The problem with explaining convergence, however, goes deeper: the eGD heuristic converges for the wrong reason: it becomes insensitive to new observations. A proper explanation requires finding the right balance between anchoring and sensitivity to economic forces.²⁴

To summarize, (i) individual choices from sets of perceived, feasible options are best explained by simple rules of thumb, instead of by utility maximization or cumulative prospect theory; not only do rules of thumb yield better predictions, they are also preferred by agents that can learn; (ii) utility maximization is not descriptive of human trading behavior, neither as way to obtain prices for proposing nor when it comes to selecting an option from a set of feasible alternatives; (iii) so far, no configuration has been identified that is consistent with

²³This result may be artificial, insofar as robot traders currently do not generate enough transactions. The current design of FACTS does not capture traders accepting n times one unit instead of accepting n units once, c.f. footnote 17. In unconditional simulations, after an acceptance, a robot trader has to compete with all others for submitting another offer. This causes different trading patterns, different price dynamics and not enough transactions. The extent to which price expectations are affected by $n \times 1$ acceptances is reduced by weighing prices with quantities. See also section 7.

²⁴FACTS can probably be improved when it comes to the timing of offers. Traders who want to cancel a pending offer are quicker to submit an offer than others, and marginal traders are the slowest to respond, but otherwise timing now is random. Gjerstad and Dickhaut (1998) suggests that traders who can accept a pending offer react more quickly than others. In a separate run, I found that convergence indeed slightly improved if acceptance was treated as an urgent option, similar to canceling an offer with unfavorable price. But there may be a stronger mechanism at work. Suppose that a pending bid is closer to a trader's expected equilibrium price than the pending ask. To the extent that his expectation is correct, this trader may expect there is a greater excess supply at the current ask than there is excess demand at the pending bid. So, he also should anticipate that sellers experience more pressure to act quickly than buyers. If the trader in question is a buyer, then he should be more patient than he would have been if he were a seller.

all observed aggregate patterns of human trading; (iv) orbiting in experimental trading in the unstable Scarf economies, as measured by cumulative angles between successive observations of last known prices, is an expression of one price steadily increasing over the course of multiple periods; the ZIP-algorithm can reproduce this; (vi) robot trading in FACTS currently does not generate enough transactions; mitigating this issue (and improving the timing of offers) may help explaining convergence.

7 Current research

An equilibrium of heuristics provides a good framework for analyzing human trading at all prices. It forces us to look at the details of individual behavior, and it suggests natural bounds on heterogeneity. Most of all, it tells us what needs to be done in order to arrive at a satisfactory explanation of trading at all prices.

With respect to improving FACTS the most important issues are (i) having a sufficient number of transactions, by emulating trading patterns in which traders can submit n acceptances of one unit and by improving quantity setting; (ii) improving the timing of offers; and, (iii) by letting learning be based on private information alone. I am also working on the interface and the documentation, so that FACTS can be made available to other students of out-of-equilibrium economics.

There are a lot of additional analyses that can be undertaken. In Ruiter (2017), the evolutionary competition between algorithms is the only form of calibration involving heterogeneity. It would be very interesting to also calibrate FACTS by letting individual robot traders learn from their own human alter egos. That would allow a more penetrating analysis of actual choices of human traders that are not perceived as feasible options by the matching robot traders. Furthermore, there are pending questions: how to form fast and frugal expectations that are as good (but not as rigid) as eGD-expectations?²⁵ How to filter implausible prices? Professor Goeree has recently provided me with the data underlying Goeree and Lindsay (2016). This also creates new opportunities for analysis and validation.

If aggregate patterns are sufficiently well explained then it would be interesting to simulate the allocations at which trading will terminate (without resetting endowments during trading). In the presence of money, these allocations may lie in the core or represent market failure. Presumably, end states have different probabilities of being formed.²⁶

²⁵Perhaps by deriving reservation prices as those prices at which bids and asks have an equal chance of being accepted.

²⁶Arthur (2006) rightly prefers equilibrium *formation* over equilibrium *selection*.

References

- Alchian, A. A. (1950). Uncertainty, evolution, and economic theory. *The Journal of Political Economy* 58(3), 211–221.
- Anderson, C. M., C. R. Plott, K.-I. Shimomura, and S. Granat (2004). Global instability in experimental equilibrium: The Scarf example. *Journal of Economic Theory* 115, 209–249.
- Arifovic, J. and J. Duffy (2018). Chapter 9 - heterogeneous agent modeling: Experimental evidence. In C. Hommes and B. LeBaron (Eds.), *Handbook of Computational Economics*, Volume 4, pp. 491–540. Elsevier.
- Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *The American Economic Review* 84(2), 406–411.
- Arthur, W. B. (2006). Chapter 32 out-of-equilibrium economics and agent-based modeling. In L. Tesfatsion and K. L. Judd (Eds.), *Handbook of Computational Economics*, Volume 2, pp. 1551–1564. Elsevier.
- Becker, G. S. (1962). Irrational behavior and economic theory. *The Journal of Political Economy* 70(1), 1–13.
- Cliff, D. and J. Bruten (1997). More than zero intelligence needed for continuous double-auction trading. Report HPL-97-157, Hewlett-Packard Laboratories.
- Crockett, S., R. Oprea, and C. R. Plott (2011). Extreme walrasian dynamics: The gale example in the lab. *American Economic Review* 101(December 2011), 3196–3220.
- Debreu, G. (1959). *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*. Cowles Foundation For Research In Economics At Yale University. Monograph. John Wiley & Sons, Inc.
- Duffy, J. (2006). Agent-based models and human subject experiments. In L. Tesfatsion and K. L. Judd (Eds.), *Agent-Based Computational Economics*, Volume 2, pp. 949–1011.
- Gigerenzer, G. and T. Pachur (2011). *Heuristics: The Foundations of Adaptive Behavior*. Oxford University Press.
- Gintis, H. (2007). The dynamics of general equilibrium. *Economic Journal* 117(523), 1280–1309.
- Gjerstad, S. and J. Dickhaut (1998). Price formation in double auctions. *Games and Economic Behavior* 22(1), 1–29.
- Gode, D. K. and S. Sunder (1993). Allocative efficiency of markets with zero-intelligence traders : Market as a partial substitute for individual rationality. *The Journal of Political Economy* 101(1), 119–137.

- Goeree, J. K. and L. Lindsay (2016). Market design and the stability of general equilibrium. *Journal of Economic Theory* 165, 37–68.
- Hahn, F. H. (1984). On the notion of equilibrium in economics. In F. Hahn (Ed.), *Equilibrium and Macroeconomics*. Basil Blackwell.
- Hahn, F. H. (1991). The next hundred years. *The Economic Journal* 101(404), 47–50.
- Hirota, M., M. Hse, C. R. Plott, and B. W. Rogers (2005). Divergence , closed cycles and convergence in scarf environments : Experiments in the dynamics of general equilibrium systems.
- Kumar, A. and M. Shubik (2004). Variations on the theme of Scarf’s counterexample. *Computational Economics* 24(1), 1–19.
- Oehmke, J. F. and R. H. Oehmke (1991). The global stability of Walrasian tatonnement. *Applied Mathematical Letters* 4(4), 33–35.
- Radner, R. (1972). Existence of equilibrium of plans, prices, and price expectations in a sequence of markets. *Econometrica* 40(2), 289–303.
- Ruiter, A. (2017). *Price discovery with fallible choice*. Thesis.
- Sargent, T. J. (1993). *Bounded rationality in macroeconomics: The Arne Ryde memorial lectures*. Clarendon Press.
- Scarf, H. (1960). Some examples of global instability of the competitive equilibrium. *International Economic Review* 1(3), 157–172.
- Smith, V. L. (1962). An experimental study of competitive market behavior. *The Journal of political Economy* 70(2), 111–137.
- Smith, V. L. (2008). *Rationality in Economics*. Cambridge University Press.
- Vytelingum, P. (2006). *The Structure and Behavior of the Continuous Double Auction*. Thesis.