Price discovery with fallible choice

Ruiter, A.G.J.M.

Chapter 4

Robot trading

How do traders propose prices? Do they set prices by optimizing expected utility against their beliefs that a proposed price will be accepted? Or do they haggle towards reservation prices, which in their turn depend on expected prices? Either way, price expectations are fundamental. This chapter tests how well different approaches to expectation formation capture human trading behavior.

The tests use FACTS, a simulation platform that has been developed as a part of this thesis (c.f. section 4.1). A detailed description of robot traders can be found in appendix B; section 4.1.2 presents a global overview. Before we turn to the results of the calibration of expectation formation in section 4.3, we discuss some methodological issues (section 4.2). Section 4.3.5 considers the impact of trading on end of period allocations and section 4.4 reflects on capturing human trading behavior and on several topics of disequilibrium theory. Finally, section 4.5 concludes.

4.1 FACTS

4.1.1 Design

FACTS (short for Fallible Agents' Commodity Trading System) is a platform designed for simulating exchange, dedicated to investigating trading behavior.[1] FACTS consists of a standard exchange economy, with an embedded market mechanism. FACTS implements a Continuous Double Auction (CDA), in which both buyers and sellers can submit offers at any time.[2] An offer is a message in which an agent indicates his willingness to buy or sell a specific quantity of a particular commodity at a certain price. Each message has a date, which refers to the time at which the contents of the message will reach the recipient. The message queue automatically sorts incoming messages according to this date. After the auctioneer informs all traders simultaneously about the latest bid / ask spreads, each trader uses this info to learn about prices and opportunities. Agents, who want to propose a new offer, send a message

[1] For more detailed information about FACTS, contact the author.

[2] It also implements a Clearing House, which clears received orders at discrete times. All simulation experiments discussed here, however, were done with the CDA.

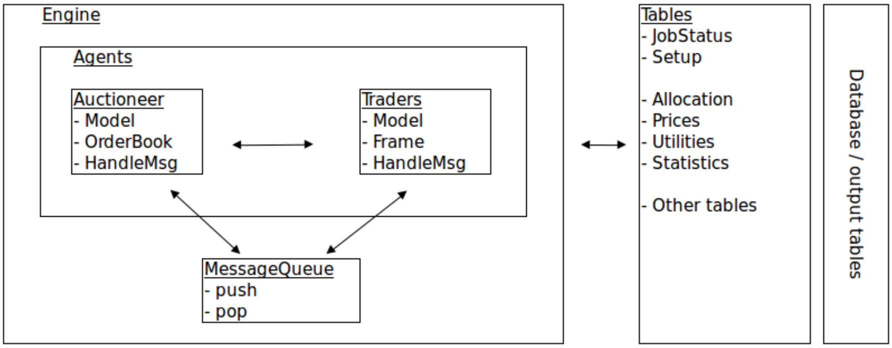

Figure 4.1 – *Design of FACTS. The exchange process, which consists of the auctioneer publishing bid / ask spreads and traders responding with new offers, is communicated through the message queue. When the auctioneer calculates statistics he may contact traders directly. Logical tables hide the database and / or output files.*

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: while slower traders are still contemplating their response to the latest bid / ask spread, they see their environment continually changing.[3,4] Simulations in FACTS, basically consist of popping and pushing messages from / into the message queue, until either time has reached its horizon or the message queue is depleted, whichever comes first, c.f. figure 4.1.

FACTS is able to run in different modes. It can let robot traders interact freely, subject to the market protocol.[5] However, it is also able to compare simulated offers with offers that were proposed during another experiment. For instance, a laboratory experiment in which human subjects trade in a particular exchange economy. For having a meaningful comparison, robot traders at each iteration must observe the same history as the participants of the laboratory experiment. To achieve this, FACTS accepts scripts, which describe experiments as a sequence of offers. Before the simulation begins, the script is loaded into the message queue. The simulation starts with the auctioneer publishing the bid / ask spread (which initially is saying that there are no pending offers). This triggers robot traders to submit an offer. In a scripted run, however, these simulated offers are not implemented. Instead, they

[3] An alternative approach is to randomly enable certain agents to submit an offer, c.f. Tesauro and Das (2001).

[4] In early experiments with Continuous Double Auctions (e.g. Smith (1962); Friedman and Rust (1993); Cliff (1997)), trading was synchronized by offering each trader the chance to either accept or reject each proposal that had been submitted.

[5] The market protocol specifies rules, which detail the exchange process, e.g. whether offers should be feasible (i.e. no short-selling), and / or improving; whether or not a new offer cancels a similar offer of the same trader, whether or not an acceptance also clears the pending rival offer.

Table 4.1 – Message flows, conditional on ScriptMode

ScriptMode	in	out	result of the inbound message
unscripted	3	13	auctioneer publishes spreads
	13	14	trader learns, posts a response time
	14	4	trader submits an offer
	4	(10),13	auctioneer processes offer; (reports trades); publishes spreads
	10	-	trader processes a trade
scripted	3	13	auctioneer publishes spreads
	13	14	trader learns, posts a response time
	14	7	trader submits an offer
	7	-	auctioneer keeps the offer for analysis
	25	4	trader submits a scripted (human) offer
	4	(10),13	auctioneer processes offer; (reports trades); publishes spreads
	10	-	trader processes a trade
semi-scripted	3	13	auctioneer publishes spreads
	13	14	trader learns, posts a response time
	14	4	trader submits an offer
	4	(10)	auctioneer processes offer; (reports trades)
	10	-	trader processes a trade

The table lists responses to inbound messages. The numbers refer to the message subject; for instance, a message with subject 3 instructs the auctioneer to publish spreads. As a result of this, the auctioneer sends each trader a message with subject 13, which contains the current floor offers. If there is a trade, then the auctioneer informs the agents involved with a message 10. In response to this event, they update their internal state but they do not reply with another message. If the run is unscripted or scripted, a message with subject 13 indirectly triggers the next message 13. In semi-scripted runs, all scripted offers are mapped to messages 3 beforehand, when the script is loaded into the message queue.

are stored for comparison with the next, scripted, offer. In a semi-scripted run, the auctioneer elicits new offers every time that a human trader in a particular experiment did propose an offer. This way, the length of the simulation experiment (and of sub-periods, if any) is exactly the same as in the laboratory experiment.[6] Table 4.1 shows message flows in different setups.

Traders assign incoming messages to their internal model or to a frame, c.f. figure 4.2. Frames are responsible for learning, perceiving opportunities and for submitting offers. A frame is a context for decision making. This includes an understanding of what has to be decided, a definition of relevant and available information, methods, et cetera. Frames implement behaviors; the next section gives an overview.

[6] A new sub-period starts after endowments have been reset.

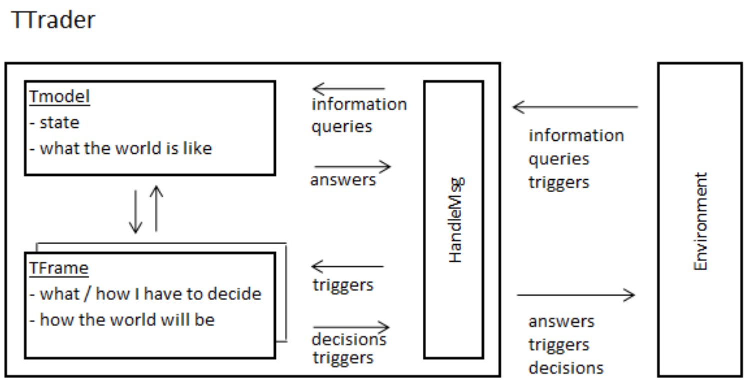

Figure 4.2 – *Architecture of a trader. When agents receive a message, they interpret it: queries from the auctioneer or informative messages (e.g. with subject 10) are dispatched to their internal model; messages which trigger a decision (i.e. with subject 13 or 14) are assigned to a frame. Traders can have multiple frames, which allow them to display different modes of behavior. If so, they assign an incoming message to the frame, which they currently consider to be the best.*

4.1.2 Overview of robot behavior

Generally speaking, behavior is how agents respond to incoming messages. In robot traders, the following process connects input to output: (i) receive market information; (ii) learn from the message; (iii) choose a best action; (iv) determine a response time; (v) haggle (if applicable) and (vi) submit the offer. Receiving market information and submitting an offer means interfacing with the message queue. Learning from a message consists of updating price expectations and perceiving opportunities. The perception of alternatives and valuation of opportunities are described in appendix B. There, it is also explained how best actions determine response times. The above steps are the same for all algorithms and where possible they use the same code. As a result, the perception of opportunities by robot traders varies with their price expectations, i.e. with the way they learn.

Tables 4.2 and 4.3 present a global overview of different algorithms and their characteristics; more details can be found in appendix B, where the robot traders, to the extent that it is necessary, have been pre-calibrated.

Let price expectations be point-expectations, and let beliefs be distributions, which express the probability that an offer at a certain price will be accepted. Method (a) of expectation formation / initialization is described in appendix A and is designed to avoid premature market failure. The belief-based algorithms use the same initialization, but here the individual point-expectations are combined to form a belief-distribution (method (b)).[7]

There are different ways of updating price expectations and beliefs. Method (a)

---

[7]This can be considered as an advantage for the belief-based algorithms.

of expectations / updating refers to drawing prices from an appropriate random distribution, and method (b) to calculating a moving average. Traders can derive expected prices from beliefs as "no arbitrage" prices, at which bids and asks have an equal probability of being accepted: method (d). The other methods consist of applying specific rules, see appendix B for details.

Table 4.3 – Differences in trading behavior

(a) Overview of characteristics per algorithm

description	ZI	eRnd	eEMA	eBAS	eGD	eME
expectations						
- initialization	(a)	(a)	(a)	(a)	(a)	(a)
- updating	(a)	(a)	(b)	(c)	(d)	(d)
choice of action						
- reservation price	(a)	(a)	(a)	(a)	(a)	(a)
- markup	0 / 0	0 / 0	0 / 0	0 / 0	0.01 / 0.01	0 .01 / 0.01
- perceived options	(a)	(a)	(a)	(a)	(a)	(a)
- best action	(a)	(b)	(b)	(b)	(b)	(b)
- quantity setting	(a)	(a)	(a)	(a)	(a)	(a)
haggling	no	yes	yes	yes	yes	yes

(b) Overview of characteristics per algorithm (ctd)

description	ZIP	AA	TU	GDA	GDW	MEA	MEW
expectations							
- initialization	(a)	(a)	(a)	(b)	(b)	(b)	(b)
- updating	(e)	(f)	(g)	(h)	(h)	(i)	(i)
choice of action							
- reservation price	(a)	(a)	(b)	(c)	(d)	(c)	(d)
- markup	0 / 0	0 / 0	0 / 0	0 / 0	0 / 0	0 / 0	0 / 0
- perceived options	(a)	(a)	(a)	(a)	(a)	(a)	(a)
- best action	(b)	(b)	(b)	(b)	(b)	(b)	(b)
- quantity setting	(a)	(a)	(a)	(a)	(a)	(a)	(a)
haggling	no	yes	yes	no	no	no	no

Different letters indicate different methods (see text). Values x/y for the markup refer to the stable / unstable economies respectively. Observe that expectation formation is the main distinguishing feature between different algorithms.

All traders observe the same history, but what they actually "see" depends on how they understand their decision problem. This explains differences in information sets, even though they receive the same messages. For appreciating some of these differences, consider the following. Most robot traders simply observe proposals and acceptances, but some interpret certain offers as a rejection of another proposal. In addition to differentiating between bids and asks, some algorithms also distinguish

Table 4.2 – Overview of algorithms

category	name	short description
shouts	ZI	Random choice of a feasible action with random shouts (i.e. offers)
	ZIP	Rule-based updating of shouts
expected prices as reservation prices	eRnd	Random expected prices
	eBAS	Expectations based on bid / ask spreads
	eEMA	Expected prices as a moving average of observed prices
	eGD	Expected prices derived from Gjerstad-Dickhaut beliefs
	eME	Expected prices derived from MaxEnt beliefs
	AA	Adaptive-Aggressive modifications of expected prices
target utility	TU	Reservation prices based on a utility target and plausibility window
beliefs	GDA	Price setting based on Gjerstad-Dickhaut beliefs with admissible prices
	GDW	As GDA, traders wait if the preferred price is not admissible
	MEA	Price setting based on MaxEnt beliefs with admissible prices
	MEW	As MEA, traders wait if the preferred price is not admissible

Algorithms have been arranged according to their central concept. For instance, the ZIP algorithm of Cliff and Bruten consists of rules which maintain shouts (i.e. offered bid and ask prices). Alternatively, traders can derive shouts from expected prices or from beliefs through haggling.

between rival and counter offers. That is, a new bid is recognized as a bid by all, but buyers and sellers respond differently to it.[8] Robot traders with adaptive expectations, on the other hand, do not differentiate between bids and asks and pool all prices for a particular commodity. So-called Gjerstad-Dickhaut beliefs make use of quantities offered and traded; all other algorithms ignore quantity information.[9] The TU algorithm tries to achieve a target utility level; as a result, it perceives markets as interdependent. The other algorithms derive their expectations and beliefs with respect to a certain market from offers which have been observed in that particular market alone. Some algorithms generate admissible prices, while others allow inadmissible expectations that may oblige traders to wait and see.[10]

Each algorithm implements reservation prices, that are used to determine whether floor offers are acceptable or not. Method (a) of choice of action / reservation price puts $p^r = (1 + \mu)\, p^e$, where p^r and p^e are the reservation and the expected price respectively and where μ is a markup (non-positive for sellers and non-negative for buyers), c.f. section 3.4. ZI and ZIP do not use the markup (i.e. $\mu = 0$), because they maintain shouts rather than expected prices. A markup in case of eRnd has no added value and has been omitted for that reason. The belief-based algorithms could use a markup, but here we want to see the effect of utility maximization. Other values for μ (choice of action / markup) have been determined as part of the pre-calibration (see appendix B). If there is no markup and no haggling (as for ZI and ZIP) then the reservation prices are equal to planned shouts. Method (b) refers to specific rules that are implemented by the TU algorithm, c.f. appendix B. Methods (c) and (d) both implement monopolistic competition; that is, traders set prices by optimizing expected utility against their beliefs with respect to proposed prices being accepted. Method (c) conditions on floor offers, so as to yield admissible proposed prices. Method (d) may lead to preferred prices that are not admissible, in that case the robot trader waits.

All algorithms first determine the feasible actions that are available to them and they do so in the same way, method (a) of choice of action / perceived options (for details, see appendix B). When it comes to selecting the best action, the ZI-algorithm pools all available feasible actions and selects one at random (choice of action / best action, method (a)). Other algorithms apply preferences to the available opportunities (method (b), see appendix B).

Proposing an offer involves selecting a price and a quantity. In case of acceptances, the price is determined by the offer which is being accepted. For new proposals traders can use their reservation price, which is possibly equal to their expected

---

[8] In the belief-based algorithms, buyers and sellers use different beliefs. That is to say, they also react differently to observed offers.

[9] In the context of the experiments of Anderson et al. (2004) this gives an advantage to the GDA, GDW and eGD algorithms, because they are not misled by the practice of n times accepting one unit, instead of accepting n units once, c.f. section 3.3.

[10] Algorithms that are supposed to deliver an admissible offer may, however, fail to do so if (i) a trader owns a floor offer, (ii) his only option would be to improve upon that offer and (iii) either there is no counter offer or (iv) it is still unclear that the owner is the marginal trader. Before such a trader is willing to propose a better price, he first wants to see a counter offer, and if that is available, he wants to wait and see if someone else submits another proposal (possibly in another market). In scripted runs, if the human counterpart of this robot trader does propose an offer, this shows up as the algorithm considering his alter ego "too eager", c.f. table 4.7.

price (haggling = no), or they can propose a better than expected price by haggling (haggling = yes). A trader commits all or part of his endowments to a transaction and uses the offered price of a commodity to calculate what should be supplied in exchange for the quantity offered. This is quantity setting / method (a), which is used by all algorithms, including the belief-based algorithms.

4.1.3 Global parameters

The global parameters in FACTS can be fixed by aligning the program with the experiments of Anderson et al. (2004), c.f. table 4.4. These settings apply to all simulations.

4.2 Calibration methodology

This section proposes criteria for assessing how well algorithms capture human trading, and it elaborates on how to measure that performance. Measurements will be based on two types of simulations: scripted, or conditional simulations and semi-scripted, or unconditional simulations.

In conditional simulations each robot trader is assigned to a human trader. Based on the same preferences, endowments and on the same publicly available information robot traders calculate a set of feasible actions, select a best action and (if they are quickest) submit an offer. In scripted runs, robot offers are generated for the purpose of comparing them with human offers. These kind of simulations are best suited for testing behavior at the individual level. In semi-scripted runs, on the other hand, human offers are only used to trigger robot offers; the latter determine the evolution of price formation. Semi-scripted runs, therefore, are fit for comparing aggregate

Table 4.4 – Global parameter settings for replicating the Scarf examples

parameter	value	remarks
number of goods	3	see section 3.2.2 on page 25
number of traders	15	replication, five times
preferences	Leontief	see section 3.2.2 on page 25
endowments	$\mathbf{W}^A, \mathbf{W}^B, \mathbf{W}^C$	see section 3.2.2 on page 25
market protocol	MUDA	see section 3.2.3 on page 26
length of periods	variable	depends on experimental data
nr of periods	1+10, 1+15	stable, cw / ccw
nr of runs	1,000	

The Scarf examples already define the number of goods and traders and their preferences and endowments. The experiments of Anderson et al. fix other parameters like the market protocol and the length of simulations.

features of price formation, such as convergence and orbiting. These emergent phenomena at the macro level can also be used for testing behavioral hypotheses at the micro level. For instance, if robot traders are gullible and likely to accept any price as an equilibrium price, then even in the unstable Scarf economies this could lead to convergence. Such an outcome, however, is at odds with the findings of Anderson et al. (2004).

Section 4.2.1 presents a list of requirements that algorithms ideally should satisfy and discusses their relative importance. Subsequent sections flesh out the details of how to apply the criteria to the data. Section 4.2.5 addresses the dependency issue that is due to the fact that there are multiple aspects of robot trading that need to be calibrated, e.g. expectation formation and the selection of a best option from the set of perceived opportunities.

4.2.1 Criteria for assessing algorithms

We want robot trading to capture the following salient features of human trading in the experimental Scarf economies:

- Convergence:
  - Convergence to the Walrasian equilibrium prices in the stable Scarf economy.
  - Non-convergence in Scarf economies if tâtonnement is not stable.
  - Convergence in approximately the same number of transactions as required by human traders.
- Consistency:
  - Actual actions of human traders are recognized as feasible opportunities.
  - Correct predictions of the actions of human traders, conditional on the observed history of trading.
  - Algorithmic offers are close to the prices that were proposed by human traders.
- Orbiting (in the unstable economies):
  - Robot prices should orbit around the equilibrium.
  - Orbiting must be in the right direction (clockwise or counter-clockwise, depending on the initial allocation).
- End-of-period allocations:
  - By the end of a period, robot traders should be as close to the Walrasian equilibrium allocation as the allocations that were achieved by humans.

At the aggregate level, the most important requirement is that algorithms achieve convergence in the stable Scarf economy while avoiding convergence in the unstable examples. At the individual level the correct prediction of human moves is key. These phenomena reflect the way that traders process information, which is our main concern. Consistency at the individual level is direct evidence, easy to measure and less sensitive to the dependencies that are discussed in section 4.2.5. Convergence or lack thereof, on the other hand, is indirect and more difficult to establish. Therefore, for the purpose of calibrating parameters, we put more weight on data at the individual level. However, for the selection of methods of expectation formation achieving convergence is more important. The predictions of human moves are one-step-ahead predictions. Algorithms that perform well in this respect but which fail to achieve convergence ignore an essential part of human behavior. Just as good descriptive models do not need to predict out-of-sample data adequately, good one-step-ahead predictions are no guarantee for achieving convergence. The point is that good descriptions and good one-step-ahead predictions can be obtained through flexibility, while good predictions of out-of-sample data and convergence require that an essential continuity is captured.

Orbiting in experimental trading seems to be the most important finding of Anderson et al. (2004), because a priori it seems less likely to occur than convergence in the stable Scarf economy. Indeed, it is this outcome that ties the experimental results of Anderson et al. closely to tâtonnement theory. Should the detection of orbiting, then, not take precedence over observing convergence? Since orbiting is conditional on not having convergence and because orbiting does not apply to the stable Scarf economy it is more natural to treat this as a subsidiary requirement. Furthermore, since we do not know what causes orbiting it is difficult to assert its importance for expectation formation or trading behavior. Anderson et al. found a low-frequency fluctuation that spans multiple periods. We propose that the observed frequency of orbiting should have been much higher for it to be the expression of the mechanism that drives tâtonnement. We also have to keep in mind that orbiting in the experimental clockwise economy was not that pronounced (c.f. chapter 3). No matter how it should be interpreted, orbiting still remains a valid criterion for discriminating between competing hypotheses.[11]

We refrain from offering a complete ordering of the different requirements. There are far too many unknowns for assuming that a particular ranking of criteria automatically will guide us to the correct model. Assessing the relative performance of various algorithms in capturing human trading at this time cannot be anything else than a tentative qualitative judgment. Chapter 6 provides another perspective on ranking by considering the ecological rationality of different algorithms.

4.2.2 Consistency

If an algorithm successfully captures human trading, then the actual moves of human traders should be in the set of perceived opportunities. Different algorithms can be compared based on the percentage of actual moves that are recognized as a feasible move. This is particularly relevant for the calibration of price expectations, because ceteris paribus the perception of opportunities depends on price expectations. High recognition rates, however, are no guarantee for correct predictions of human moves. Since the latter matter most, we prefer to calibrate based on prediction rates.

[11] Also see footnote 18 on page 63.

In scripted runs, we may also consider the difference between prices proposed by human and algorithmic traders. Here, we cannot assume that the identity of the proposing agent is the same. A robot trader may prefer to wait whereas his human counterpart decides to submit an offer. As a result, there is also no reason why offers of the quickest human and robot traders will be similar. One may propose to sell commodity 2, while the other wants to buy commodity 3. Yet, the offers can be compared by looking at the synchronized time series of proposed prices (c.f. section 3.3). For each algorithm we can calculate the average distance between the actual (human) and virtual (robot) time series as:

$$\delta = \frac{1}{4N} \sum_{t=1}^{N} \sqrt{\left(p_{bid,2}^{human}(t) - p_{bid,2}^{robot}(t)\right)^2 + ... + \left(p_{ask,3}^{human}(t) - p_{ask,3}^{robot}(t)\right)^2}$$

A lower value of δ means that a particular algorithm generates prices that are on average closer to human prices. In scripted runs algorithms must improve upon pending offers submitted by human traders. This provides them with substantial guidance. Therefore we consider δ to be of lesser importance than the simulated confidence intervals that are described in section 4.2.3.

For getting a better understanding of why some algorithms perform better than others, we also use additional consistency measures:

- Consider a human buyer who at some point in time expects the fair price of commodity 2 to be 35. It seems safe to assume that this buyer will make an offer below, or at or perhaps even slightly above 35 (if he applies a markup). Offers to buy that substantially overshoot a reservation price are not, however, compatible with the algorithm.

- Similarly, certain algorithms sometimes stipulate that it is better for a trader to wait rather than to propose an offer. Observed discrepancies with respect to waiting also signal that a particular algorithm does not explain human offers.

- What is the probability that an actual trader and his simulated counter-part make an offer at a particular iteration? And, conditional on the identity of proposing traders, what is the probability that they select the same commodity for making an offer? And if they select the same commodity, do they select the same action (buying or selling)? And if they also select the same action, is this action similar in the sense of being either an acceptance or a rejection? By calculating these probabilities we can also gain some insight into why one algorithm is better in capturing human trading behavior than another.

4.2.3 Convergence

The difference between price formation in the stable and the counter clockwise unstable laboratory economies is both obvious and remarkable at the same time. Visual inspection of figure 3.1 suggests that price formation in session 414 was convergent, while it wasn't in session 511. However, it is not immediately clear in which sense price formation does or does not converge. There are two issues:

- Exhaustion of the set of feasible Pareto improvements makes prices formation unstable (c.f. appendix A).

- Furthermore, human traders at odd times may try to move the market. For instance, session 414 saw unexpected asks of 6000 or more when it was already clear that acceptable prices were approximately 40 and 20.

From these considerations, we may conclude that convergence manifests itself not as a limit, but rather as a degree of concentration of observed prices. For gauging convergence in the stable Scarf economy, we primarily rely on visual inspection of the graphs of the first twenty runs of unconditional price formation. The graphs shown in section 4.3 and in appendix B are the most representative; if there is substantial diversity across runs then this is mentioned in the captions. Additionally, we compare price processes in terms of (i) the distance between prices (see section 4.2.2) and (ii) the percentage of observations that are close to the average.

For both human and simulated price formation, we can measure the percentage of observed prices that are close to their average, say within a margin of $\pm 10\%$.[12] If an algorithm captures human trading, then it should have a comparable score on this concentration statistic. We determine this by simulating confidence intervals for the concentration statistic and by verifying whether these confidence intervals cover the observed scores for human trading. We calculate symmetric two-sided critical values at the 10% level, using 1,000 runs per algorithm. That is, the critical values are such that 50 runs will have values below the critical lower bound and 50 runs will exceed the critical upper bound.

In the stable Scarf economy, the observed concentration statistics for human trading are 74% and 78% for commodities 2 and 3 respectively, while in the unstable counter clockwise example both percentages are equal to 17%. These values are based on trading prices; therefore, we can also calculate the percentages for the clockwise unstable economy: 17% and 26% for commodities 2 and 3 respectively.[13]

Results of these tests will not necessarily be clear-cut. It is unlikely that good hypotheses have small confidence intervals that cover all observed values, while bad hypotheses cover none. Hence, we should look at the number of human concentration statistics covered. However, with robot trading it is far more difficult to achieve convergence than not. Covering the human concentration statistics in the stable treatment is therefore more important than covering the statistics of the unstable treatments. Furthermore, if confidence intervals are very wide then that may indicate that the number of simulation runs is not sufficient and / or that there is no robust

[12] Anderson et al. (2004) already defines "close" as $(p_2, p_3) \in [36.5, 43.5] \times [16.5, 23.5]$, but it does not define when observations are "sufficiently often" near to the equilibrium. This bandwidth is described as prices being at most four "cents" removed from their equilibrium values. This sounds comforting, but four "cents" amounts to a deviation of 20% of the equilibrium price in case of commodity 3. Actually, Anderson et al. (2004) argues that their bandwidth is conservative: usual conventions would justify a range nearly twice the size. We find that difficult to accept. We prefer a bandwidth of 10% for each commodity, while admitting that this value is hard to rationalize.

[13] The file AllTrades.csv from Anderson et al. contains trades from all sessions, but it excludes the trades from training periods. This would give 83% and 75% (stable), 9% and 10% (ccw) and 9% and 15% (cw) as concentration statistics. The percentages we use for the convergent and the counter clockwise treatment come from the raw data, which *include* the training period. Comparison shows that after the training period there is less volatility in the stable treatment and more instability in the counter clockwise treatment. Therefore, it makes sense to correct the observed clockwise concentration statistic for including a training period: $9\% \times \frac{17}{9} = 17\%$ for commodity 2 and $15\% \times \frac{17}{10} = 26\%$.

Table 4.5 – Sensitivity of confidence intervals to the number of simulations

runs	stable		ccw		cw	
	good 2	good 3	good 2	good 3	good 2	good 3
100	87-93	96-98	3-7	4-83	3-83	3-76
200	87-93	95-98	3-7	10-83	2-84	3-76
300	87-93	95-98	3-7	14-83	3-84	3-76
400	87-93	95-98	3-7	11-84	3-84	2-76
500	87-93	95-98	3-7	13-83	3-84	3-76
600	87-93	96-98	3-7	16-83	3-84	3-76
700	87-93	95-98	3-7	13-83	3-84	3-76
800	87-93	95-98	3-7	14-84	3-84	3-76
900	87-93	95-98	3-7	13-83	3-84	3-76
1,000	87-93	95-98	3-7	14-83	3-84	3-76

Simulated confidence intervals (symmetric, at confidence level = 0.1) for concentration statistics, per treatment and commodity, for varying number of runs. Intervals bounds are percentages. The table has been generated with the eGD-algorithm. It is clear that the size of the confidence intervals is quite stable. Whether a confidence interval is wide or not therefore cannot be attributed to an insufficient number of simulation runs.

convergence. Table 4.5 demonstrates that a wide confidence interval is a reliable indicator of a lack of robust convergence.

Based on the analyses of chapter 3, we expect the convergence of human price formation in the stable treatment to be quite robust. Wide confidence intervals therefore can be taken as a contra-indication. If confidence intervals are small and situated at low levels of concentration then that would be a strong indication of a lack of convergence, while small intervals at high levels of concentration are compatible with robust convergence.

The concentration statistic will not enable us to reliably detect convergence, but it will help in rejecting simulation results, that are not similar to human price formation. For instance, the statistic will be too low if simulated prices fluctuate around a stable equilibrium with large deviations, or if there are small deviations around a pronounced upward or downward trend. And, it will be too high if the economy quickly settles on a equilibrium, or if time series of prices exhibit minor fluctuations around a very mild trend. A comparable score on the statistic is not sufficient for concluding that simulated price formation is like human price formation. This is why visual inspection of representative simulated time series is also important.

We want to establish whether robot trading can achieve convergence similar to human trading. For a comparison to be fair, robot traders should get about the same "time" as human traders did. But how to align machine time and laboratory time? In conditional price formation, robot traders submit the same number of proposals as did human traders. Setting the number of offers equal is the closest one can get to synchronizing the length of the experiments in terms of the number of decisions taken.[14] By simulating in semi-scripted mode the same can be accomplished for unconditional price formation: every time a human trader proposes an offer, the auctioneer informs robot traders about the current spread, which makes them to submit a response time and possibly a proposal. Unfortunately, we have no human offers for the clockwise unstable Scarf economy. Instead, we use the data of the counter-clockwise treatment to also trigger robot offers in the clockwise treatment. Since session 511 (ccw) consisted of 15 periods (after the training period), and both clockwise treatments consisted of 10 sessions, in this case we truncate the data set after 1+10 periods.[15]

4.2.4 Orbiting

Although near-exhaustion of feasible Pareto improvements may contribute to unstable prices, this does not explain the lack of convergence in the counterclockwise unstable Scarf economy. Session 511 is characterized by long term fluctuations in trading prices across periods (c.f. figure 3.1). These fluctuations, which are evidence of orbiting, have been confirmed by Goeree and Lindsay (2016).

Anderson et al. (2004) has introduced the clock hand model as a way of discriminating between clockwise and counter clockwise orbiting.[16] Suppose that a virtual clock hand, that is fixed at the Walrasian equilibrium prices, points at observed (synchronized) prices. For each observation, one can calculate the angle that the clock hand makes relative to 12:00 hrs. Between two successive observations the clock hand either moves forward or backward. For instance, consider two successive observations which lie at 2 and 3 o'clock, relative to the Walrasian equilibrium prices. That is, the angles (as measured in hours) relative to 12:00 hrs are 2 (old) and 3 (new) respectively. Anderson et al. (2004), following mathematical convention, calculates the difference in angles as old - new = 2 - 3 = -1. Hence, clockwise movement generates negative increments and counter clockwise movement produces positive differences. We, however, prefer to calculate progress as new - old, leading to clockwise movements being classified as positive and counter clockwise movements as negative.

If there is orbiting then movements of the clock hand are likely to be highly correlated, successive increments of the angles will not cancel out and their sum will indicate the direction of orbiting. With FACTS we can simulate the distribution of the final values of accumulated differences between successive angles (or cumulative angles for short), to express both the propensity of orbiting and its direction.[17] If

[14] Most likely, the number of decisions was less than the number of offers in the experiments of Anderson et al., see footnote 6 on page 30.

[15] The complexity of trading in the ccw and cw unstable economies is comparable, and certainly less than trading in the stable Scarf economy (c.f. footnote 3 on page 133). In the clockwise treatment, therefore, the number of offers per period (before exhausting the set of feasible Pareto improvements), will be closer to the counter-clockwise unstable than to the stable treatment. Hence, it is better to use a truncated data set of session 511 instead of the full data set of session 414.

[16] Anderson et al. (2004) has also devised the so-called quadrant model, c.f. section 3.2.5. We prefer the clock hand model for its greater power (c.f. section 3.5.1 on page 44) and its ability to aggregate over an experiment (see below).

[17] This is inspired by Hirota et al. (2005), which plots the cumulative angles of a clockwise and

price formation in the clockwise (cw) and counter clockwise (ccw) examples is in accordance with tâtonnement theory, then we expect the probability density in the cw-case to be shifted to the right and in the ccw-case shifted to the left.

To illustrate these concepts, figure 4.3 shows the probability densities of cumulative angles of both Zero Intelligence (ZI-) and ZI-Plus (ZIP-) trading. The symmetry of the probability densities around zero in case of ZI-trading shows that FACTS is not biased in either direction; ZIP-trading, on the other hand, conforms to the predictions of tâtonnement theory.[18] This is interesting because ZIP-traders manage their competitiveness in both markets independently: prices, or price changes, in one market do not affect prices in the other markets, and yet there is systematic orbiting across different runs. Furthermore, densities can have multiple (local) maxima; this phenomenon also manifests itself with other algorithms.[19]

4.2.5 Dependencies

Capturing human trading behavior by means of robots requires the calibration of multiple fundamental parameters, such as expectation formation and the selection of a best option from the set of perceived opportunities. Ideally, one should evaluate all possible combinations of these parameters before identifying which configuration is best. Instead, we calibrate expectation formation conditional on rules of thumb for prioritizing feasible actions (c.f. B).

The reason for this is that nine out of thirteen algorithms use point expectations rather than distributions. Choice based on utility maximization, however, requires beliefs, i.e. expectations as distributions instead of point expectations. Put differently, only four out of thirteen algorithms can potentially be improved by considering other methods for selecting a best alternative from the sets of feasible actions. Furthermore, an algorithm's ability to recognize human actions as opportunities precedes the selection of a preferred alternative. To the extent that we rank algorithms accordingly our analysis is independent of the method of selecting a best alternative. Although limited in scope, section 4.3.4, addresses the dependency between expecta-

---

a counter clockwise experiment to demonstrate the difference in price dynamics between the two treatments. Goeree and Lindsay (2016) uses the same technique for demonstrating orbiting in the replication of the ccw-case.

[18]In our simulations, orbiting seems to be due to one price increasing beyond bounds: in ZIP-trading in the counter clockwise economy the price of commodity 2 increases over the course of the experiment, while in the clockwise treatment it is the price of commodity 3 that rises steadily. In session 511 (the counter-clockwise treatment) of Anderson et al. the price of commodity 2 also increases, but here at some point the long-term trend is reversed, c.f. figure 3.1. That probably means that human traders eventually start to doubt whether a further increase is credible. The point, however, is that human trading initially experienced the same unchecked upward pressure on the price of commodity 2 as in our simulated price formation. The sustained increase of a trading price seems to reflect the *absence* of a negative feedback mechanism, while orbiting in tâtonnement theory, of course, is the *expression* of negative feedback. That raises some doubts about the proper interpretation of experimental orbits.

[19]It is not exactly clear what is the reason for the secondary maxima. Possibly, this is a reflection of the probability that a single price starts increasing beyond bound at a certain time; if this occurs sooner, the cumulative angle is likely to be greater in absolute value. Alternatively, successive angles may offset each other; if there is a pattern in this then some final values can become more likely than others.

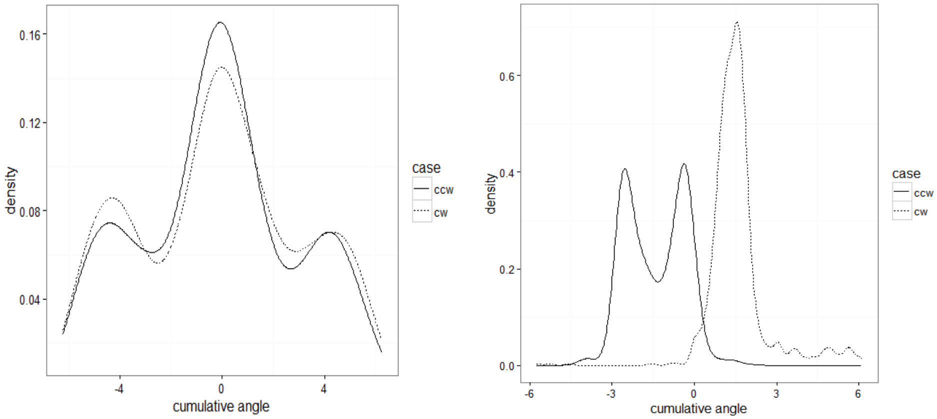

Figure 4.3 – *Simulated probability densities of cumulative angles in the unstable Scarf economies. The ZI-densities (left) are symmetric around zero, showing that (i) there is no inherent bias in the system and (ii) orbiting is due to expectation formation. ZIP-trading (right) exhibits systematic orbiting because the densities have shifted away from zero; furthermore, they each have shifted in the right direction.*

tion formation and choice among opportunities in the form of a sensitivity analysis.

There is another, deeper dependency: the rules that identify opportunities do not preclude arbitrage, but they substantially limit it by letting agents accept a pending offer provided that they do not intend to reverse it at a later stage. Arbitrage is considered in chapter 5. That is, robot traders buy what they need and sell what they can spare. This restricts the number of opportunities that robot traders observe. Table 3.1 shows that approximately 90% of all actions in session 414 (i.e. in the stable treatment) can be characterized as "regular". Regular actions are consistent with non-speculative behavior. When we determine whether actual moves are recognized as feasible opportunities, or whether human moves are correctly predicted we restrict ourselves to this set of regular actions. This largely takes care of the limitation. To the extent that regular opportunities are not correctly identified, all algorithms suffer (but not necessarily to the same degree). Based on the fact that recognition rates appear reasonable and consistent across the stable and counter clockwise treatment, we believe that the adverse effects are limited and that differences in recognizing human moves as feasible opportunities can be attributed to variations in the quality of price expectations, c.f. table B.1 on page 153. The aggregate analysis, however, can be biased because it uses all the data and because human traders were engaged in arbitrage.

4.2.6 End-of-period allocations

Although our focus is on price formation we should not lose sight of the purpose of trading, which is to improve the well-being of agents. This is achieved by the market through successive changes in the distribution of commodities. Although any core allocation could be used as a benchmark, we will use the distance of simulated allocations to the Walrasian equilibrium allocation as a way to assess the similarity of robot and human trading. That way we can be sure that our criteria with respect to allocations and prices are consistent. Furthermore, the stylized facts of chapter 3 also suggest that end-of-period allocations of human trading are gravitating towards the Walrasian equilibrium, c.f. figures 3.6 and 3.7.

Instead of measuring the Euclidean distance, it is more common to express this distance by means of efficiency. This is defined as the ratio of actual average utility divided by potential average utility (as realized in the Walrasian equilibrium).[20] However, since similar efficiency scores less than one can be realized in an infinite number of ways, we prefer to visualize the end-of-period allocations as in figure 4.4. In each 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. Initially, in the stable Scarf economy, traders of type III own all the money (i.e. commodity 1); during the course of trading, money should "traverse the triangle" and end up equally in the hands of traders of types I and II. The other commodities should also cross the triangle. The equilibrium shares (which are the same for each of the three Scarf economies) are marked "WE".

4.3 Calibration of expectation formation

4.3.1 Consistency

How well do different algorithms recognize human moves as feasible actions and how well do they predict them? Table 4.6 provides theanswers to these questions. The

[20] Interestingly, in the Scarf economies efficiency can also be calculated by using excess quantities. The maximum amount of a commodity that can be taken away from a trader without reducing his level of utility is called an excess quantity. By using excess quantities for calculating efficiency one does not have to assume the existence of a Walrasian equilibrium. Let $\xi(\mathbf{W})$ be the efficiency of allocation $\mathbf{W}$ and let $v_j(\mathbf{W})$ be the excess quantity of commodity j, i.e. the sum of excess holdings of commodity j. Then efficiency can be alternatively defined as

$$\xi(\mathbf{W}) = 1 - \frac{\frac{v_1(\mathbf{W})}{400} + \frac{v_2(\mathbf{W})}{10} + \frac{v_3(\mathbf{W})}{20}}{3}$$

with $0 \leq \xi(\mathbf{W}) \leq 1$. One easily derives:

$$\begin{aligned}
\mathbf{v}(\mathbf{W}) &= \begin{pmatrix} 400 - 400u_1 - 400u_2 \\ 10 - 10u_2 - 10u_3 \\ 20 - 20u_1 - 20u_3 \end{pmatrix} \Rightarrow \\
\xi(\mathbf{W}) &= 1 - \frac{3 - 2u_1 - 2u_2 - 2u_3}{3} \\
&= \frac{(u_1 + u_2 + u_3)/3}{1/2}
\end{aligned}$$

with 1/2 being the average of potential utility levels, as obtained in the Walrasian equilibrium.

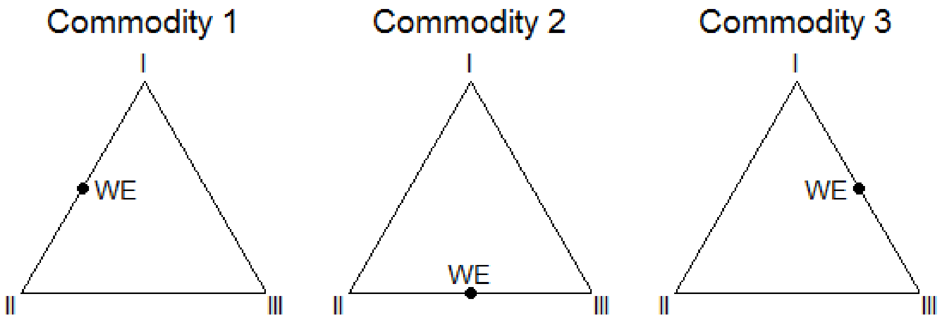

Figure 4.4 – *Shares per trader type in the available commodities in the Walrasian equilibrium. This equilibrium is the same for each of the three Scarf economies.*

MEW-algorithm has the highest recognition rates, while eBAS provides the best predictions of human moves.

GDW and MEW recognize acceptances better than other algorithms. One would expect that this would also lead to better predictions, because acceptances are generally preferred to other opportunities. GDW- and MEW-traders, however, set prices by maximizing expected utility against their own subjective beliefs. That is, these traders implement the hypothesis of monopolistic competition. Both types of traders are adverse to rejections, because their expected utility is very sensitive to the probability that an offer is accepted. As a result, they tend to propose prices that are favorable to their prospective trading partners. This raises their recognition rates, because (i) they can improve upon and (ii) even accept almost every pending offer. At the same time, their prediction rates are low because these algorithms accept far too many proposals: every offer appears to be a bargain. Remarkably, these algorithms are even outperformed by the Zero Intelligence algorithm. The underlying beliefs of MEW- and GDW-traders are not at fault; indeed, the eGD- and eME-algorithms (which derive expected prices from GD- and ME-beliefs respectively) are among the best algorithms in terms of recognition. The high recognition and low prediction rates of GDW and MEW are both due to selecting prices based on optimal expected utility.

The good recognition rates of the eBAS-algorithm are not surprising. Since expected prices are based on bid / ask spreads, possibilities to improve upon floor offers are always detected. The outperformance in terms of predictions presumably is due to these expectations being more elastic than others: they are not constrained in any way by a notion of plausible price levels. This may also explain the high propensity to disqualify human offers as overshooting the reservation price (c.f. table 4.7). Instead of "seeing through" haggling, eBAS takes such offers seriously. Haggling (especially if one or both floor prices are missing) continuously disturbs eBAS expectations. Convergence, so it seems, requires that expectations are somehow anchored.

The algorithms other than GDW and MEW are comparable in that they all try to learn what the equilibrium price is, rather then determine which price would be most advantageous. It makes sense, therefore, to rank them based on the recognition rate alone. Then the TU-algorithm stands out in the stable Scarf economy. Its

Table 4.6 – Recognition and prediction of human moves

Description	stable economy		ccw unstable economy	
	recognized (%)	predicted (%)	recognized (%)	predicted (%)
ZI	37.7	53.7	41.8	57.8
eRnd	73.3	52.7	66.4	55.1
eEMA	74.8	68.7	74.2	73.3
eBAS	74.1	69.0	74.0	72.5
eGD	75.5	68.0	73.2	62.5
eME	76.2	63.2	71.8	61.3
ZIP	76.5	67.1	76.9	69.8
AA	74.5	65.0	75.2	67.8
TU	77.5	65.4	73.7	62.4
GDW	77.0	51.8	68.9	51.7
MEW	83.7	48.7	85.5	48.5

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 those that appear to be motivated by arbitrage. The prediction percentages are conditional on the human action being recognized as a feasible option. Prediction percentages are calculated assuming that all algorithms (except ZI) use rules of thumb to select a best option from the set of perceived alternative actions. Recognition and prediction rates are fairly stable across the stable and counter clockwise treatments. The eBAS algorithm is best; even though its recognition rate is slightly lower than those of eEMA and eGD it outperforms in terms of predictions. The high recognition and the low prediction rates of MEW (and GDW) are both due to price setting by optimizing against subjective beliefs. Belief algorithms GDA and MEA have been omitted as a result of the pre-calibration, c.f. appendix B.

distinguishing feature is that it perceives opportunities in different markets as interdependent. Observe that in the stable treatment all traders operate in two markets, while in the unstable treatments only one in three traders is active in both markets. This explains why the recognition rate of TU-traders is higher in the stable Scarf economy. Apparently, interdependency also plays a role in the perception of opportunities by human traders. Interestingly, ZIP-traders perform well in the counter clockwise economy; they do better than AA-traders even though the latter can adjust to changes in volatility.

Table 4.7 provides some insight as to why some algorithms perform better than others. If we allow robot traders to approve or disapprove of offers of their human alter egos than they find that human offers often overshoot their reservation prices. To a large extent, such discrepancies are acceptable. If an algorithm correctly describes expectation formation and if there is a reasonable spread in expected prices, then approximately half the traders should be expected to overshoot reservation prices,

Table 4.7 – Qualitative comparison of human and robot offers

(a) The stable Scarf economy

probability	ZI	eRnd	ZIP	eEMA	eGD	eME	eBAS	AA	TU	GDW	MEW
overshooting	0.41	0.41	0.33	0.33	0.32	0.35	0.47	0.41	0.43	0.38	0.27
too eager	0.28	0.07	0.01	0.01	0.01	0.02	0.01	0.02	0.01	0.06	0.20
= agent	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.08	0.08
= good (2/3)	0.57	0.58	0.67	0.68	0.68	0.65	0.67	0.67	0.67	0.59	0.62
= action (buy/sell)	0.79	0.88	0.92	0.92	0.91	0.90	0.92	0.91	0.90	0.89	0.90
= accept / reject	0.20	0.46	0.30	0.26	0.26	0.33	0.25	0.26	0.31	0.49	0.41

(b) The unstable ccw economy

probability	ZI	eRnd	ZIP	eEMA	eGD	eME	eBAS	AA	TU	GDW	MEW
overshooting	0.47	0.47	0.33	0.35	0.40	0.41	0.38	0.36	0.47	0.40	0.19
too eager	0.22	0.23	0.06	0.06	0.10	0.11	0.06	0.07	0.10	0.24	0.30
= agent	0.08	0.07	0.07	0.07	0.08	0.08	0.08	0.08	0.08	0.08	0.08
= good (2/3)	0.65	0.75	0.78	0.78	0.75	0.74	0.77	0.77	0.76	0.71	0.78
= action (buy/sell)	0.83	0.89	0.90	0.90	0.88	0.88	0.92	0.88	0.87	0.89	0.91
= accept / reject	0.16	0.32	0.20	0.17	0.19	0.20	0.17	0.19	0.22	0.31	0.23

"Overshooting" and "too eager" give an assessment of human trading behavior from the perspective of an algorithm. "Overshooting" means that offers from human traders exceed the reservation price. "Too eager" means that a trader should have waited instead of submitting an offer. The line "= agent" gives the simulated probability that the trader proposing an offer in session 414 or 511 is the same as his corresponding robot (base rate: 1/15 = 0.07). The other simulated probabilities are conditional on the previous probabilities. For instance, "= accept / reject" refers to a robot trader accepting or rejecting an offer, conditional on having the same trader taking the same action in the same market. Best values have been underscored. MEW is best on overshooting because it has the lowest probability of disapproving of human offers. Predicting the market conditional on the identity of the trader is easier in the counter clockwise unstable economy because most traders are active in one market only. Interestingly, utility optimization guides GDW- and MEW-traders to the "wrong" market more often in the stable economy. This table also shows the high propensity of eBAS to predict human actions (buy / sell) correctly.

because of the random matching between traders and expectations. What is interesting, however, is which algorithms disapprove less of human offers than others. At first sight, it seems that MaxEnt beliefs are the most compatible with human offers. But this is a direct consequence of belief-based algorithms trying to avoid rejections: it is impossible to overshoot a bid of 200 or an ask of 1. If, for this reason, we ignore the belief-based algorithms then eGD yields reservation prices that are most similar to human offers in the stable treatment. This algorithm also is best in the sense that it disqualifies human traders less often as being too eager.

Avoiding rejections also leads to belief-based algorithms having good probabilities of accepting an offer. The high propensity of the eRnd algorithm to predict acceptance correctly is due to conditioning on agent, good and action. This, for instance, leaves buyers with random price expectations between the floor bid (if any) and the maximum price. Excluding these particular cases, we find that eME-reservation prices are relatively good in explaining acceptances.

Table 4.7 contains some other interesting elements, such as the better than random chance for the MEW-algorithm to predict the identity of the human trader submitting an offer correctly. This suggests that waiting is recognized by human traders as a distinct opportunity.

All algorithms use the same rules for selecting the market in which to submit an offer. As described in appendix B, this choice depends on comparing the best opportunities for each market. One reason why ZIP-traders may do relatively well in predicting the market and the action (buy or sell), conditional on the identity of the trader, is that ZIP-traders manage their competitiveness. In the absence of an opportunity to accept, and given the fact that human traders seem to prefer regular over strategic offers, the choice of market and of an action depends on the availability of a regular offer. That is, on being competitive.

Overall, table 4.7 shows that eGD-trading is a strong algorithm in the stable treatment: ignoring GDW and MEW for erratic behavior then eGD is the best or very close to the best, except for acceptance / rejection. On similar grounds, ZIP is strong in the ccw treatment.

4.3.2 Convergence

4.3.2.1 Confidence intervals and distance between prices

In this section we compare aggregate features of human and robot price formation. For this, we simulate (i) confidence intervals for the concentration statistic and (ii) the average distance between robot and human price formation, c.f. table 4.8.

Judging by the lengths of their confidence intervals, there are only two algorithms in the stable Scarf economy that have small confidence intervals situated at high levels of concentration, eME and eGD. Generally speaking, their convergence is too good because they are likely to have more concentration in trading prices compared to human trading. This is partly due to tabulating these beliefs based on all observations, which makes the expected prices relatively inelastic. The eME-algorithm does cover one observed concentration statistic in the stable Scarf economy. It also has the closest average distance to human prices. Neither eGD- nor eME-trading leads to convergence in the unstable economies, but here as well they are likely to exhibit too much concentration. The eGD-algorithm does cover three observed concentration statistics, but its confidence intervals are quite long. Observe that the eME- and eGD-algorithm indirectly profit from the fact that belief-based algorithms pool the initial expectations of all traders. From the start they are better informed than, for instance, the eEMA-algorithm.

The analysis illustrates that it is important to have tests at both the individual and the aggregate level, because this allows us to see through the high percentage of correct predictions of the eBAS-algorithm. These expectations work well at the individual level, because they never miss opportunities to improve upon pending offers. Nevertheless, because of a lack of anchoring, they are easily upset by haggling and therefore they do not lead to convergence at the aggregate level.

TU reservation prices, on the other hand, are anchored by traders pursuing an objective. TU-traders perceive markets as interrelated; this, however, also introduces instability to a degree that is absent in human trading. This could be due to our implementation, or to human traders focusing on one market at a time. Even though this algorithm is based on eGD-expectations, it does not deliver robust convergence in the stable Scarf economy.

Table 4.8 – 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
eRnd	9.1	9.3	9.2	11-16	12-17	10-15	12-17	12-19	10-16
eEMA	5.9	6.4	6.2	5-29	4-34	0-3	3-25	4-27	1-11
eBAS	6.1	6.6	6.4	12-66	8-28	6-16	6-43	26-65	10-23
eGD	5.8	6.4	6.2	87-93	95-98	3-7	14-83	3-84	3-76
eME	5.8	6.4	6.1	75-82	84-88	35-47	28-51	24-63	45-58
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
TU	6.1	6.7	6.5	23-74	18-62	19-44	14-61	9-62	11-33
GDW	8.2	8.8	8.5	8-38	11-39	22-45	2-44	3-43	16-46
MEW	6.4	6.8	6.7	6-69	5-16	13-23	19-27	8-22	7-18

Statistics are averages, calculated over 1,000 runs. Here, eME emerges as the best algorithm because (i) in the stable Scarf economy it has small confidence intervals with one covering the observed concentration statistic, while the other one is close; (ii) it also has the smallest average distance to human prices; (iii) no convergence in the unstable economies, although here we do have too much concentration of trading prices compared to human price formation; (iv) it is better than its closest rival, eGD. The ZIP-algorithm covers most human concentration statistics, but its intervals are generally quite wide indicating that price formation is not very robust.

4.3.2.2 Visual inspection of price formation in the stable treatment

Figures 4.5 and 4.6 show selected graphs of time series of trading prices in the stable Scarf economy. The examples of eME- and eGD-trading shown in figure 4.5 are representative. The graphs of TU-trading to a certain extent also look similar. For other algorithms, however, there is more variation in the graphs of the first twenty simulations.

Although it may seem that eGD-trading generates fewer transactions than TU-trading, this is actually not the case: TU-trading exhibits more variation in trading prices. The Adaptive-Aggressive (AA) algorithm automatically adjusts to observed volatility. If competitive pressure is high then traders engage in aggressive buying and selling. This could have been a comparative advantage; however, figure 4.6 suggests that AA-trading can create bubbles.[21]

[21] The possibility of a trading process endogenously generating bubbles raises some doubts about Fisher's assumption of the absence of favorable surprises, which is crucial for the convergence of the

It is clear that the resemblance between time series from robot and human trading requires improvement. Nevertheless, we assess eGD-expectations to be relatively good. Convergence is consistent (albeit too strong) and its average trading prices tend to approximate the Walrasian equilibrium values slightly closer than the prices of eME, c.f. figure 4.7.

4.3.2.3 A Marshallian path?

According to Plott et al., trading in a Continuous Double Auction follows a Marshallian path.

> "The Marshallian model incorporates a principle of self-organizing, coordination that mysteriously determines the sequence in which specific pairs of agents trade in an environment in which market identities and agent preferences are not public. Disequilibrium trades along the Marshallian path of trades do not change the theoretical equilibrium. The substance of this paper is to demonstrate that the Marshallian principle captures a natural tendency of the adjustment in single, continuous, double auction markets and to suggest how it takes place.", Plott et al. (2013, p. 190).[22]

The "mysterious" mechanism matches buyers and sellers with the greatest consumer and producer surplus. The paper of Plott et al. assumes exogenous reservation prices, as in the earliest experiments with trading at all prices. Absent such fixed reservation prices, it is unclear whether exchange occurs along a Marshallian path in case of trading at all prices in the Scarf economies. Easley and Ledyard (1993) is quite skeptical about the relevance of a Marshallian path in CDAs.

The underlying idea - that trading occurs between agents who stand to gain the most - may, however, be valid. If traders use the rules of thumb, then they will start to look for possibilities to cancel or accept a pending offer. If they find such an opportunity then they need to look no further. Put differently, traders who are able to

---

Fisher process, c.f. section 2.3.2.

[22]Many articles that discuss the Marshallian path do not give a reference. Plott et al. (2013), on the other hand, is very specific: "refer to appendix H of Marshall, 1961 eighth edition, p. 806". However, in appendix H we could not find any explicit mention of matching buyers and sellers according to the size of their consumer and producer surplus (these are explained in appendix H, footnote 86 on p. 811).

Book V, chapter V, section 4 suggests that suppliers with the greatest producer surplus will be the ones that most easily increase production after the "market price" increases due to excess demand. This mechanism, that brings (short term) "market prices" back to (long term) "normal prices" works in the medium to long run. In the experiments of Plott et al. supply is fixed; in effect they study the dynamics of "market prices".

Book V, chapter II presents an example of the equilibration of a market for corn, that seems to be more to the point. Here, Marshall comments: "In this illustration there is a latent assumption which is in accordance with the actual conditions of most markets; but which ought to be distinctly recognized in order to prevent its creeping into those cases in which it is not justifiable. We tacitly assumed that the sum which purchasers were willing to pay, and which sellers were willing to take, for the seven hundredth quarter would not be affected by the question whether the earlier bargains had been made at a high or a low rate.", Marshall (1961, p. 334) The seven hundredth quarter is the marginal unit. This behavioral assumption, which is *not* describing a Marshallian path, gives rise to what Milton Friedman describes as the "usual demand" curve, c.f. Friedman (2007, p. 15) and section 4.4.2.1.

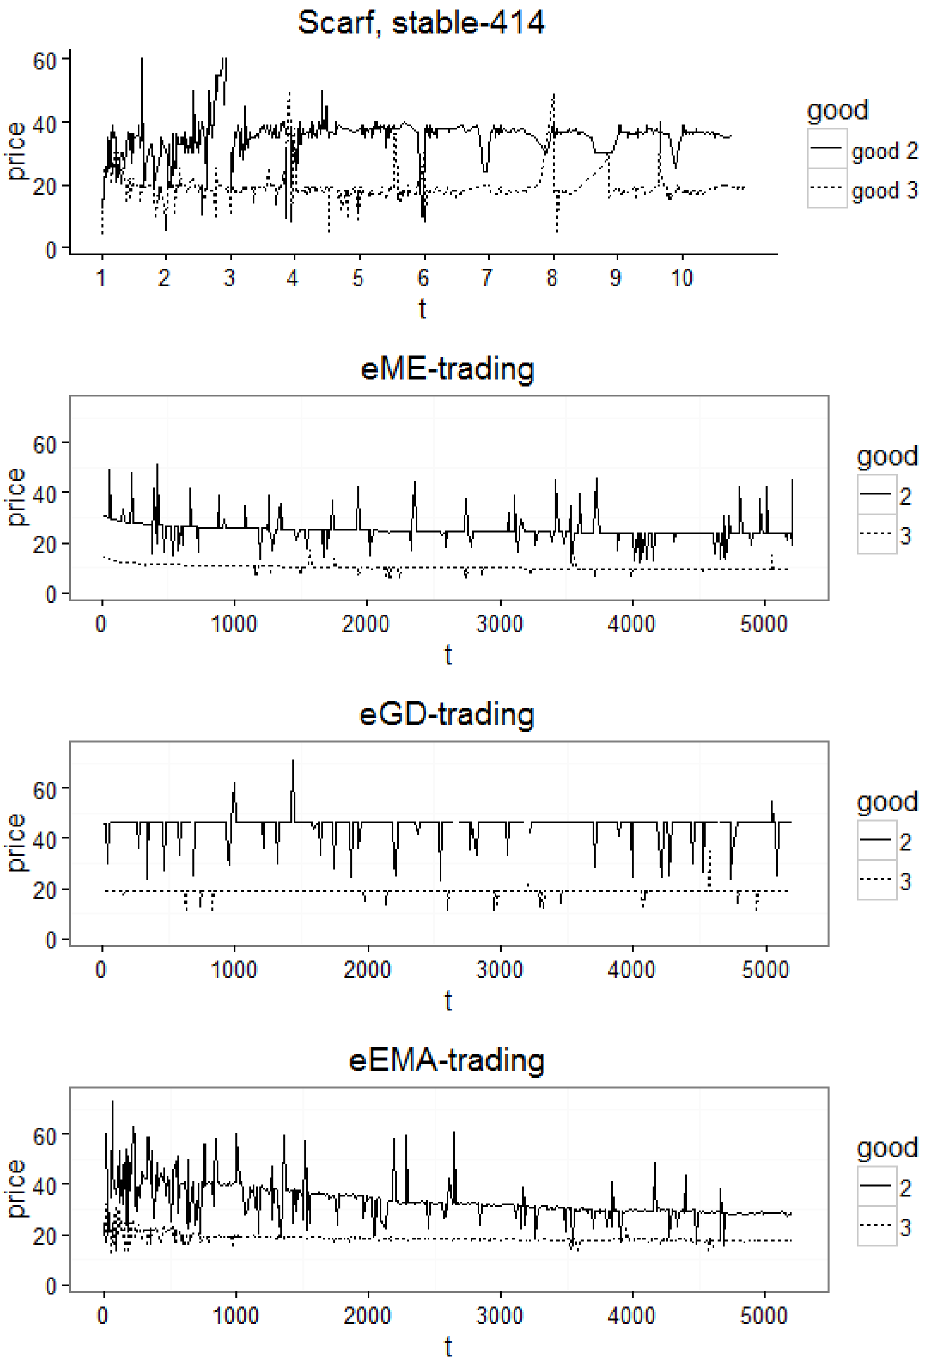

Figure 4.5 – *Visual inspection of time series of trading prices. Graphs for eGD- and eME-trading are representative; for eEMA, however, there is a fair amount of variation between the graphs of different runs. The eGD- and eME-algorithms exhibit too much concentration.*

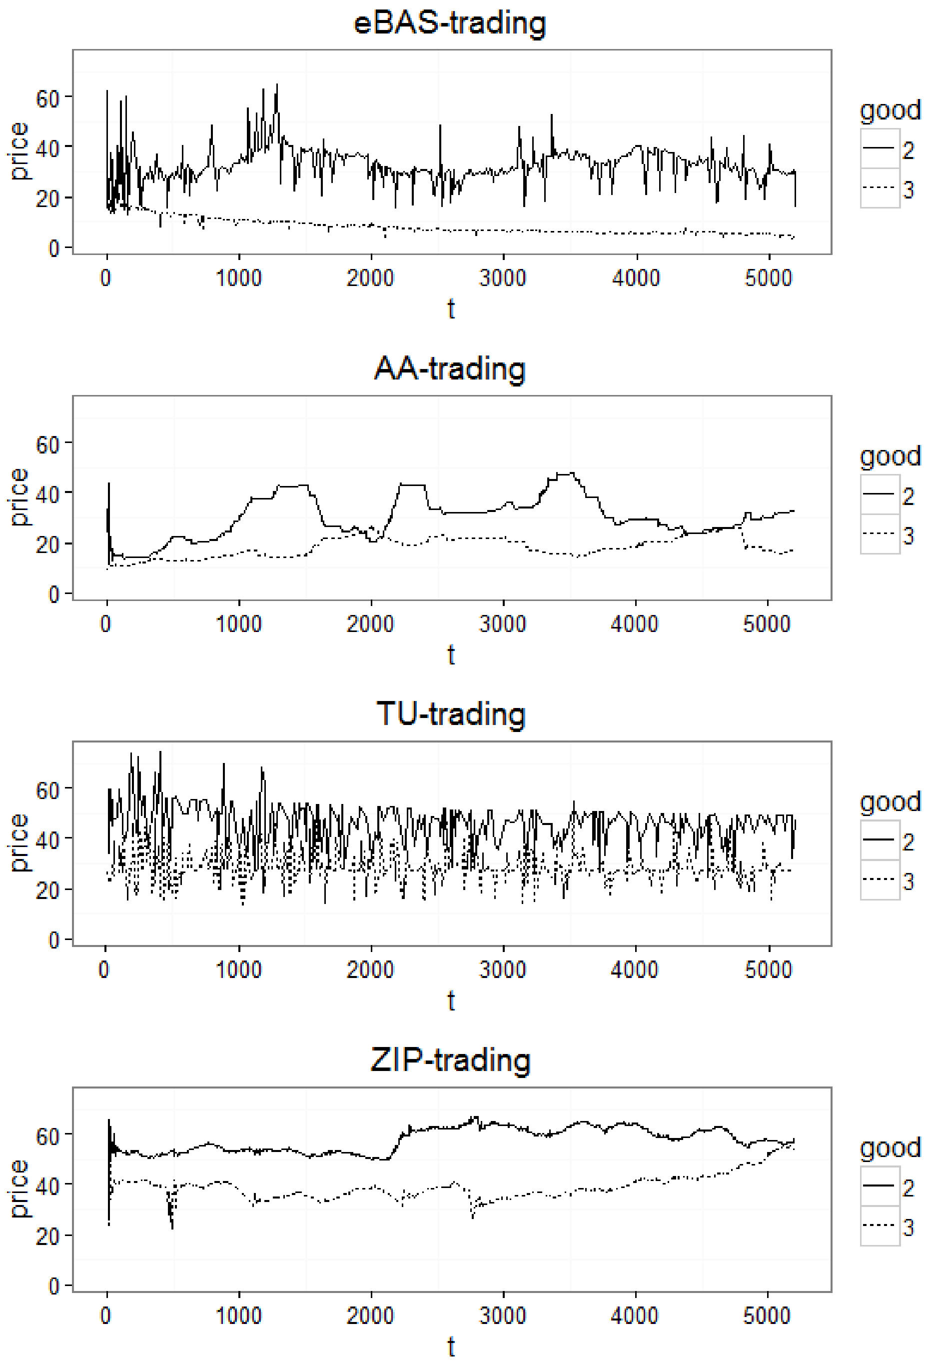

Figure 4.6 – *Visual inspection of time series of trading prices. The graph for TU-trading is representative; for other algorithms there is a fair amount of variation between graphs of different runs.*

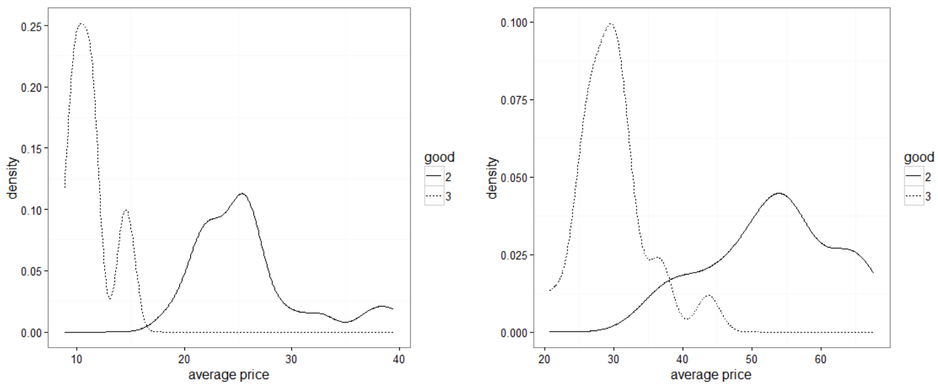

Figure 4.7 – *Simulated probability densities of average trading prices of eME-trading (left) and eGD-trading (right) in the stable treatment, based on 1,000 runs. Prices of eGD-trading are typically higher than prices of eME-trading. Both algorithms assign a higher price to commodity 2; for modal prices we have approximately* $p_2 \approx 2p_3$.

accept a floor offer may indeed respond quicker than others.[23] Whether this helps or hurts convergence depends on expectation formation. Under the protocol of Anderson et al., acceptances also clear pending rival offers. To the extent that algorithms rely on the presence of floor offers and / or can be easily upset by haggling, expectations may become unhinged and that will adversely affect convergence.[24]

Currently, acceptances are treated as "eager" options. This means that acceptances compete with other proposals for being the first to be submitted. If traders with acceptances would respond ahead of others with new proposals, i.e. if acceptances were "urgent", then there would be more transactions.[25] It is, therefore, an interesting experiment to vary the priority of acceptances. Unlike other algorithms (e.g. eBAS), eGD-trading does not suffer from having more acceptances and fewer floor offers. The recognition and prediction rates of eGD-trading are not sensitive to changing acceptances from eager into urgent options (e.g. average prediction in the stable treatment improves from 68.0% to 68.3% (the biggest change)). In the stable economy there are 8% more transactions, while there are 4% more in the counter clockwise unstable treatment. Furthermore, urgent acceptances slightly decrease concentration in the stable Scarf economy, while maintaining the small and high confidence intervals, c.f. table 4.9. Concentration in the counter clockwise example, on the other hand, has increased. Giving more priority to acceptances does not adversely affect convergence; the effect seems positive instead.

[23]Gjerstad and Dickhaut (1998) also makes an assumption to that effect.

[24]Initial simulations with a variant of the TU-algorithm confirmed the increased instability. The implementation of Gjerstad-Dickhaut beliefs in FACTS, on the other hand, is rather insensitive because it tabulates all offers.

[25]See section B.1.1 for an explanation of the priority rules in FACTS.

Table 4.9 – Sensitivity to the status of acceptances

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
eager	5.8	6.4	6.2	87-93	95-98	3-7	14-83	3-84	3-76
urgent	6.0	6.9	6.6	80-87	89-95	2-8	60-84	45-62	8-79

Statistics for eGD-trading in the stable treatment with a varying priority of acceptances. Concentration in the stable economy remains consistently high. Concentration has increased for commodity 3 in the counter clockwise and for commodity 2 in the clockwise treatment, suggesting that convergence benefits from acceptance getting a higher priority. The distance between prices slightly increases, especially in the unstable treatment.

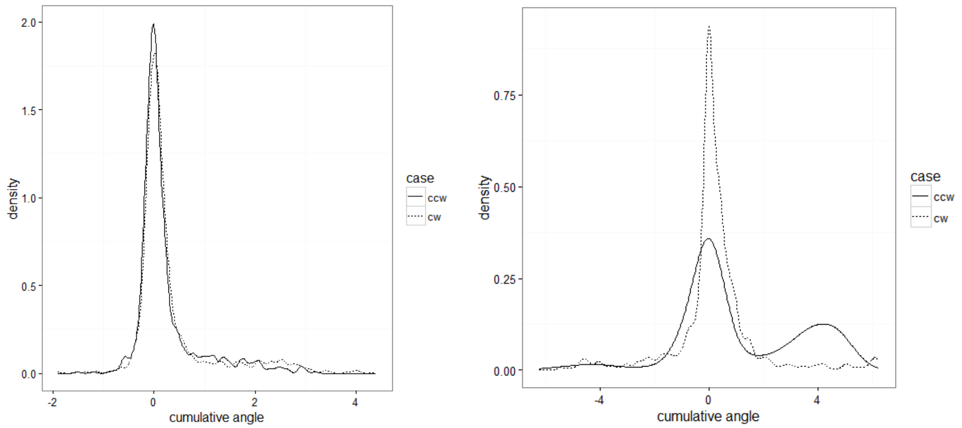

Figure 4.8 – *Simulated probability densities of cumulative angles in the unstable Scarf economies, as generated by eME-traders (left) and eGD-traders (right). The eME-traders generate no systematic orbiting. There is a better chance of observing orbiting in eGD-trading, albeit in the wrong direction. In the experiments of Anderson et al. there was correct orbiting in the counter clockwise example and little orbiting in the clockwise economy.*

4.3.3 Orbiting

Figure 4.3 illustrates that the ZIP-algorithm leads to systematic orbiting in the direction that is predicted by tâtonnement theory. Other algorithms produce less pronounced orbiting and often in the wrong direction.

In the light of the results of table 4.8 we focus our analysis of orbiting on the eGD- and eME-algorithm. The eGD-traders are more likely to generate prices that orbit, but chances are less than even that orbiting will occur in the correct direction, c.f. figure 4.8. Commodity prices in eME-trading, on the other hand, exhibit no systematic orbiting. Orbiting does not offer much help in discriminating between eME- and eGD-trading.

Table 4.10 – Sensitivity to rules of thumb

configuration	prediction rates (%)	
	stable	ccw
eME - EU	48.1	49.1
eME - RoTh	63.2	61.3
eGD - EU	48.7	49.3
eGD - RoTh	68.0	62.5

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 those that appear to be motivated by arbitrage. The prediction rates are conditional on the human action being recognized as a feasible option. Row "eME - EU" gives results for the combination of eME-expectations and selection of a best alternative based on expected utility maximization; row "eME - RoTh" gives analogous results for selection based on the rules of thumb. The latter clearly perform better; this justifies the use of the rules of thumb for the calibration of expectations. Furthermore, eGD consistently performs better than eME.

4.3.4 Dependencies

The calibration of expectation formation depends on how a best alternative is selected from a set of feasible alternatives. Algorithms that maintain point expectations need rules, like the ones that are derived in section B.1.3. Algorithms that depend on beliefs (whether directly or indirectly) can apply different methods, such as utility maximization (c.f. section B.1.2).

We have seen that price setting based on maximizing utility against subjective beliefs, i.e. monopolistic competition, leads to behavior that bears little resemblance to how human traders behaved in the experiments of Anderson et al.. Deriving expected prices as "no arbitrage" prices, on the other hand, leads to good recognition and prediction rates, to robust convergence in the stable economy and to lack of convergence in the unstable economies. A ranking of eGD- and eME-trading may depend on how best actions are selected from the set of perceived feasible actions; this is taken up in chapter 5.

The sensitivity analysis of table 4.10 vindicates the use of rules of thumb in the calibration of expectations. The combination of eME- or eGD-trading with the maximization of expected utility would lead to prediction rates that are worse than the results of other algorithms such as eEMA, eBAS, ZIP, AA and TU. In chapter 5, we consider a wider range of selection rules.

4.3.5 End of period allocations

How does expectation formation affect equilibrium selection? To find out, we plot the shares of trader types in the available commodities, c.f. figures 4.9 and 4.10. The blobs of similar algorithms, like eGD and eME, look similar, but different from the ones of ZIP and AA (these algorithms are also related, their blobs also resemble each other). Different spreads demonstrate that equilibrium selection is sensitive to expectation formation (which is to be expected).

The most significant feature of simulated allocations, however, is their spread and distance to the Walrasian equilibrium allocations. This implies that the efficiency of robot trading is substantially less than the efficiency of human trading. This is partly due to robot traders generating an insufficient number of transactions. Since traders buy what they need and sell what they can spare, more transactions generally mean greater concentration and more efficiency. The spread may also depend on the robustness of price convergence.

4.4 Discussion

4.4.1 Capturing human trading behavior

Different algorithms for expectation formation lead to substantially different price dynamics. So far, none of the algorithms generates prices that closely resemble human trading in each of the three different examples of Scarf. This does not seem due to lack of variation: we have considered (i) point expectations and beliefs (either ignoring or taking waiting and / or quantities into account); (ii) reservation prices based on observed prices (backward looking), on floor offers (forward looking), on utility maximization and on targets (with different attitudes toward target setting, c.f. appendix B); (iii) reservation prices with and without markup and even automatic adjustment to observed volatility in the form of adaptive-aggressive trading.

The calibration suggests that both the eME- and eGD-algorithm are better in capturing human expectations than the others: (i) they are likely to converge in the stable Scarf economy; (ii) they have confidence intervals for the concentration statistics that are small and that either cover or are close to the observed values of human trading; (iii) while eGD does better than eME, both have relatively good prediction rates; (iv) their average distance to human prices is also relatively small. However, there are important weak points as well: (i) even though there is no convergence in the unstable treatments there is currently too much concentration; (ii) if there is orbiting, then more likely than not it will be in the wrong direction. Of these two algorithms, we prefer the eGD-algorithm to the eME-algorithm because of its better prediction rate.

As a result of tabulating all observations, the GD- and ME-beliefs, and indirectly also the eGD- and eME-expectations, eventually become insensitive to new information and relatively inelastic with respect to time (c.f. figure 4.5 and table 4.8).[26] This contributes to convergence and also to concentration of trading prices in the unstable treatments. The belief-based algorithms can be improved by making beliefs dependent on fewer observations. This will also benefit the eGD- and eME-algorithms. Calibrating the number of observations that determine beliefs, however, seems premature because convergence also indirectly depends on the number of transactions in the simulations and that number currently is too low (c.f. table 4.11).[27]

If a trader perceives multiple opportunities, the rules of thumb stipulate that

[26] The motivation for tabulating all observations was that the subjects of Anderson et al. often submitted $n \times 1$ acceptances instead of $1 \times n$. Using a small number of observations for calculating beliefs would introduce artificial effects that do not seem warranted, c.f. section B.2.6.

[27] The significance of the low number of transactions became only gradually apparent.

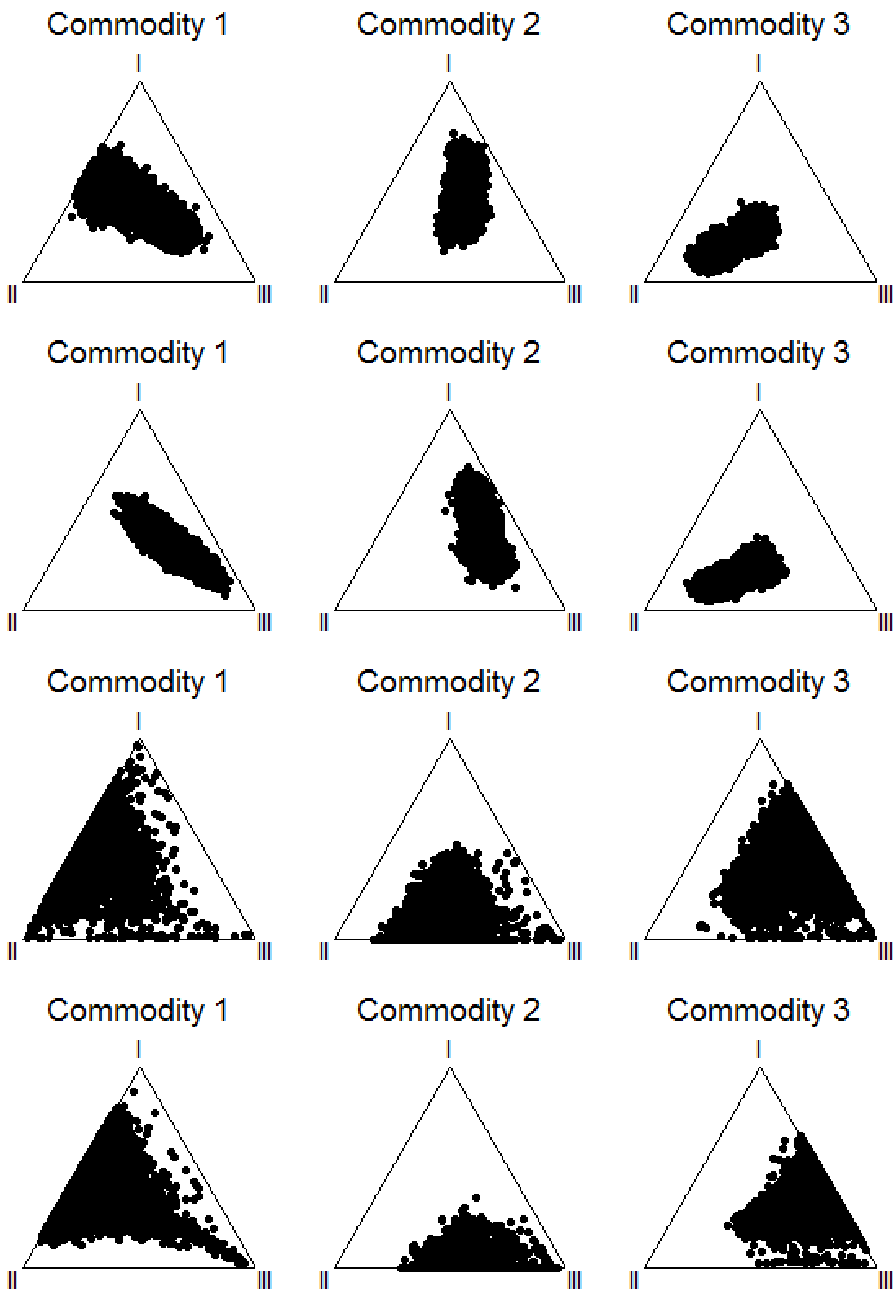

Figure 4.9 – *Shares per trader type in the available commodities in the stable treatment: eGD (row 1), eME (row 2), ZIP (row 3) and AA (row 4). Each dot represents an allocation at the end of a period. The results of 10+1 periods from 1,000 runs have been plotted. Equilibrium allocation selection is sensitive to expectation formation; related algorithms (eGD / eME and ZIP / AA) lead to similar results. The spread between allocations is considerable, even if price convergence is robust (eGD and eME). As a result, efficiency is substantially less than in human trading.*

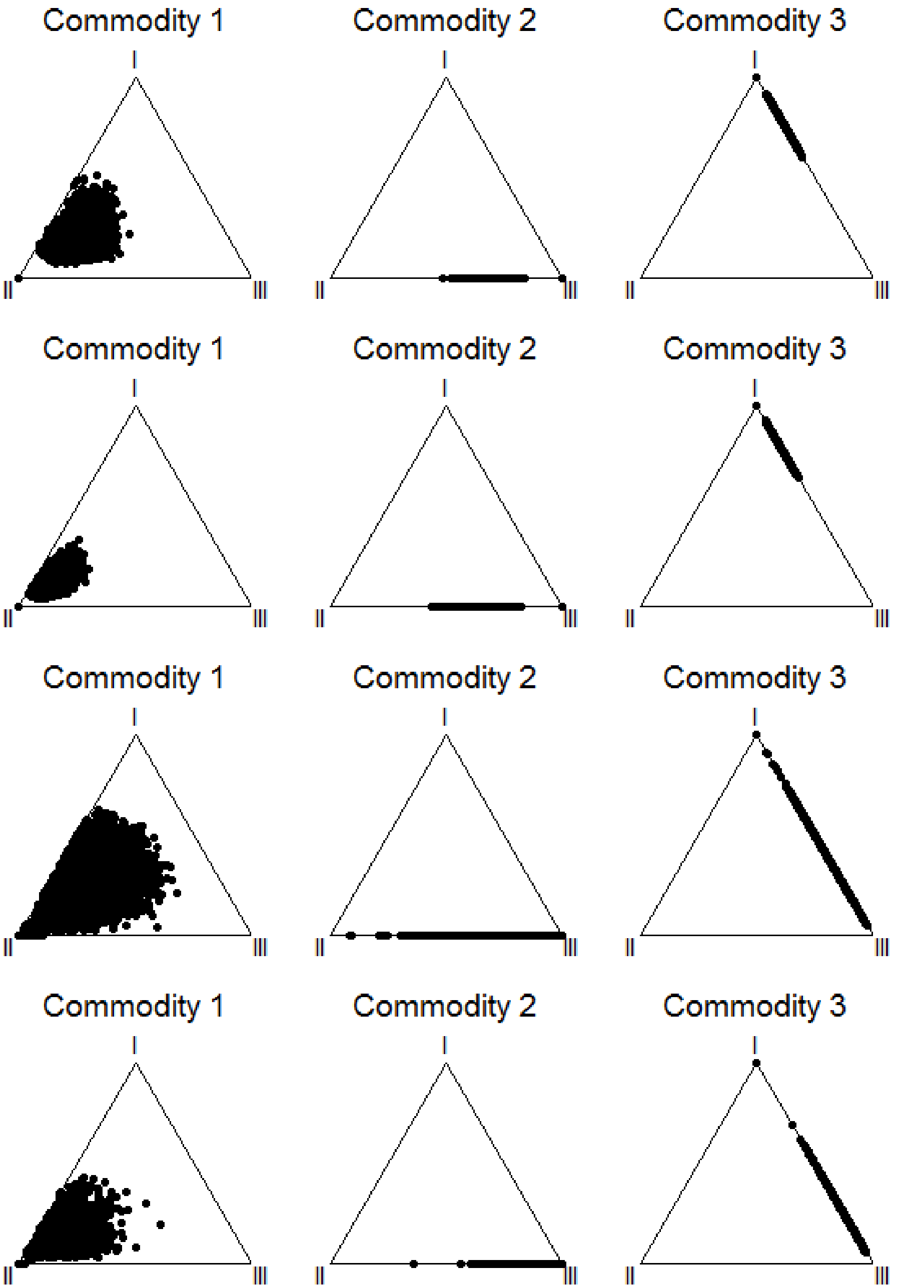

Figure 4.10 – *Shares per trader type in the available commodities in the counter clockwise treatment: eGD (row 1), eME (row 2), ZIP (row 3) and AA (row 4). Each dot represents an allocation at the end of a period. The results of 16+1 periods from 1,000 runs have been plotted. Equilibrium allocation selection is sensitive to expectation formation; related algorithms (eGD / eME and ZIP / AA) lead to similar results. The allocations of commodities 2 and 3 are on a side of triangle because initially they are owned by traders who derive utility from them.*

Table 4.11 – Number of transactions

description	stable			ccw		
	good 2	good 3	total	good 2	good 3	total
human traders	724	1272	1996	492	805	1297
eGD-traders	389	585	974	505	469	974
eME-traders	435	568	1003	472	219	691

Robot trading does not generate enough transactions compared to human trading. Since there is much more identifiable arbitrage in the practice periods, it would have been interesting to correct for that. Due to the unequal length of periods, however, we refrain from calculating average number of transactions for the practice and normal periods.

cancellations and acceptances have priority. That is already conducive to having many transactions. Having traders, who can accept a pending offer, respond quicker than others would lead to 4% - 8% more transactions. A higher value of the markup could contribute to more trading. The eGD-traders, for instance, could have had 10% more transactions if the markup would have been set to $\mu = 0.1$. That would not, however, generate a sufficient increase in the number of transactions. Furthermore, consistency with trader behavior suggests that the markup should be low (c.f. appendix B). It is also unlikely that arbitrage explains the discrepancy between robot and human trading.[28] The low number of transactions is symptomatic of a misspecification in the current model of trading behavior.

Most likely, the misspecification is due to quantity setting. We have opted to let traders propose single units of commodities 2 and 3, because initial simulations, with quantities derived from expected utility optimization, systematically ran into market failure (often after three trades, c.f. appendix A). Furthermore, the subjects of Anderson et al. also often decided to accept only one unit, e.g. in the stable Scarf economy this accounts for more 71%-87% of all accepted bids and asks in both markets, see table 3.6. Finally, trading one unit at a time simplifies weaving human and robot offers and has the added advantage that the speed of convergence depends on the quality of price expectations only. However, we do know that human traders were inclined to accept n times one unit instead of once accepting n units. If a robot trader submits an acceptance then chances are slim that he will submit two successive acceptances (because the robot has to take two decisions and compete for priority twice, while a human trader takes one decision leading to multiple transactions and hence also to more correlation between successive transaction prices).[29]

Improved modeling of quantity setting, and at the same time modeling the process of submitting one or more offers, creates several technical issues. For instance, traders with unprocessed offers should typically ignore triggers to submit new offers, to prevent a self-sustaining, exponential growth of submitted offers. On the other hand, they should resume responding to triggers when they believe that their pend-

[28] During the experiments, that is in periods > 0, identifiable arbitrage accounted for approximately 6% of the actions. The gap in table 4.11 exceeds that.

[29] As explained in section 3.3, clustering the observations to make them more reflective of individual decisions would distort price formation, i.e. publicly available information.

ing offers will not be processed any more because they have become unfeasible (e.g. if another trader was quicker to accept). "Machine time" and "laboratory time" will have to be synchronized more intimately for properly matching robot and human offers. Robot traders should be triggered by decisions of human subjects instead of by their (unclustered) offers, for giving robot and human traders an equal chance of achieving equilibrium. Solving these issues is left for future research.

4.4.2 Disequilibrium theory

4.4.2.1 Why would prices converge to the Walrasian equilibrium?

Price expectations can prevent convergence to the Walrasian equilibrium, but do they force convergence? If so, why? Where is the economics that can explain this phenomenon?

Algorithms can be distinguished according to whether they simply extract information from observed prices or whether they also infuse some information about a trader's own situation into his reservation prices. To underscore this point, in the unstable Scarf economies, sophisticated traders can deduce the Walrasian equilibrium prices based on introspection alone, without having to trade (c.f. chapter 3). Rather than looking at past prices a trader may also look at what he wants and/or at his opportunities. On the other hand, if traders use rules of thumb for prioritizing feasible actions then it seems less likely that they condition their reservation prices on what they have learned from contemplating their opportunities.[30]

Monopolistic competition and algorithms that manage reservation prices based on a target with respect to utility both "look inwards" and therefore should be better at anchoring price formation. Unfortunately, the performance of these algorithms is not (yet) good enough. In case of monopolistic competition this is due to traders being averse to rejections; the algorithms with a utility target may have suffered from the fact that beliefs over time become inflexible. Furthermore, if robot traders would generate more transactions then their utility levels are likely to increase. Perhaps that would benefit algorithms with a utility target more than other algorithms.

Milton Friedman has suggested that traders may correct for unfavorable transactions that they have committed to in the past. According to him, buyers (say traders of type III) determine their optimal expenditures on commodities based on expected prices; from these expenditures they deduct what already has been spent (possibly at unfavorable prices), and from this amount they deduce reservation prices by considering how much money is available for acquiring a certain amount of a commodity.[31] It is straightforward to generalize this kind of behavior to other types of traders who also have to sell. Here, there are similarities with algorithms that manage reservation

[30]It may be possible to test to what extent introspection or an analysis of opportunities determines expectation formation. If human traders with CES-preferences would have more difficulty in achieving convergence to the Walrasian equilibrium than traders with Leontief preferences, then that would be an indication that people form expectations by also looking at what they want. With CES-preferences the optimal ratio between quantities consumed is no longer fixed and is dependent on expected prices.

[31]This kind of behavior generates what Friedman calls "usual demand", as distinct from normal demand that derives from equating the rate of substitution of commodities with their relative prices, c.f. Friedman (2007, p. 15).

prices based on a utility target. While the latter may try to compensate adverse conditions in one market by adjusting reservation prices in another market, Friedman's traders typically feed corrections back into the same market in which previous "mistakes" were made.[32]

In section 3.5.2, we have suggested that experimental prices approximating the Walrasian equilibrium can be due to the demand symmetry, that is implicit in the preferences in the Scarf examples. A trade at a false price in one commodity does not distort demand for the other commodity. The eGD-algorithm can be useful in testing this hypothesis. We can measure the Euclidean distance between average commodity prices and their Walrasian equilibrium values when the demand symmetry becomes increasingly distorted. The latter can be achieved by changing the parameters of the Leontief preferences so that traders no longer divide their budget equally between the two commodities they like.

Consider seven exchange economies, that are characterized by their Walrasian equilibria (commodities j in rows, agents i in columns):

$$(\mathbf{p}^*, \mathbf{X}^*) = \left(\begin{pmatrix} 1 & 40 & 20 \end{pmatrix}, \begin{pmatrix} 50k & 400-50k & 0 \\ 0 & \frac{5}{4}k & 10-\frac{5}{4}k \\ 20-\frac{5}{2}k & 0 & \frac{5}{2}k \end{pmatrix} \right)$$

with $k = 1, 2, .., 7$. Let preferences be Leontief, i.e.

$$u_i(\mathbf{x}_i) = \min_{\alpha_{ji} \neq 0} (\alpha_{ji} x_{ji})$$

and

$$\alpha_{ji} = \begin{cases} & (i=1) \wedge (j=2) \\ 0 & (i=2) \wedge (j=3) \\ & (i=3) \wedge (j=1) \\ \frac{1}{x^*_{ji}} & \text{otherwise} \end{cases}.$$

Each trader achieves utility level $u_i(\mathbf{x}^*_i) = 1$ in each of the Walrasian equilibria. Endowments are the same as in the stable Scarf economy; therefore $k = 4$ corresponds to the stable Scarf example. For other values of k we have varying degrees of distortion of the symmetry in demand.[33] Let initial price expectations be random, with $p_{2i} \sim$ uniform $(20, 60)$ and $p_{3i} \sim$ uniform $(10, 30)$. Based on our hypothesis, we expect that for $k = 4$ average trading prices will be closest to the Walrasian equilibrium prices and that the distance increases if k moves to either 1 or 7. Similar to the experiments of Anderson et al. these economies are replicated five times.

The pattern in figure 4.11 suggests that the simulated distances are not random. Hence, in some economies it is easier to approximate the Walrasian equilibrium prices than in others. Interestingly, the demand symmetries in the stable Scarf economy $(k = 4)$ do not yield the smallest distance. This seems to falsify our hypothesis. It is not yet clear why the distance in case $k = 7$ is minimal. Knowing how difficult it is to

---

[32] When modeling this type of behavior, of course it would be possible (and interesting!) to treat gains and losses asymmetrically.

[33] An alternative approach is to replace the Leontief-preferences by CES-preferences, because the introduction of decreasing marginal utility also breaks demand symmetry.

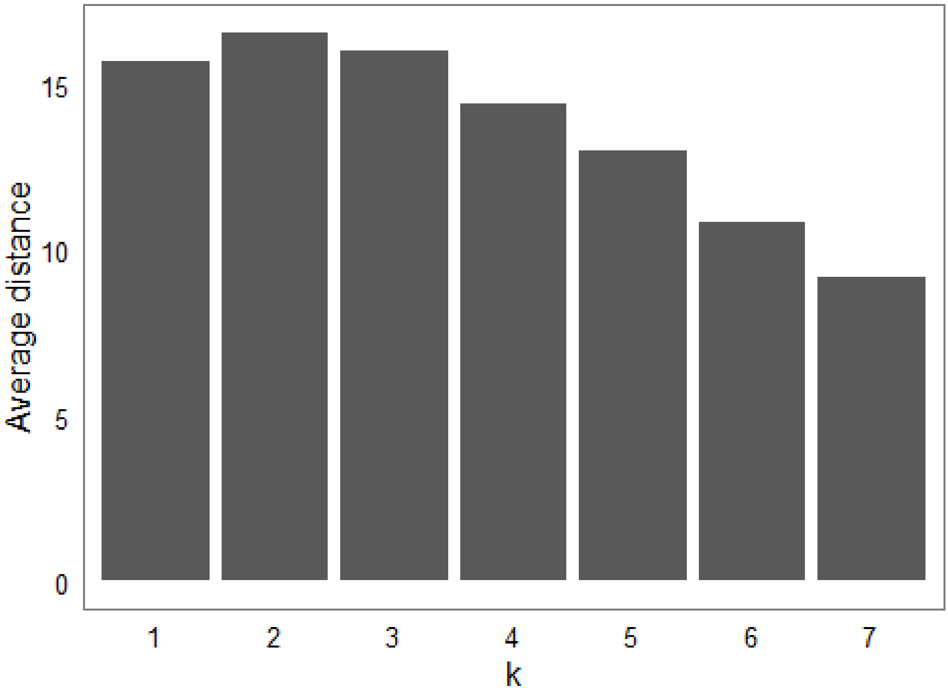

Figure 4.11 – *Euclidean distance between simulated commodity prices and Walrasian equilibrium prices. Distances are averages over 1,000 runs, with eGD-traders. The graph shows that in some cases it is easier to approximate the Walrasian equilibrium prices than in others. Surprisingly, the stable Scarf economy, which corresponds to $k = 4$, does not yield the best result even though here we have demand symmetries which neutralize transfers of wealth due to trading at false prices.*

approximate Walrasian equilibrium prices is useful for assessing human price formation. The eGD-traders are more sensitive to initial expectations and less responsive to observed prices than human traders are. Therefore, we expect that humans trading at all prices in these examples would exhibit a greater sensitivity to variations of k.

4.4.2.2 Monopolistic competition

The main lesson from the calibration of expectations is that monopolistic competition is not a good explanation of price formation.[34] Robot traders that set prices based on

[34]This is an illustration of how effective the Scarf examples can be as a critical test, as advocated by Karl Popper (e.g. Popper (1975, 1983)). Whether we have a falsification is a matter of critical, rational discussion and not of blindly applying an epistemological rule. One has to assess whether the alleged falsification is robust, and where to put the blame. This illustrates the importance of understanding, as advocated by Hahn, as opposed to Friedman's faith in following successful predictions (c.f. Hahn (1984); Friedman (1953)). Although science cannot be equated with an ongoing critical, rational discussion this may well be the best ideal that can be had. Rhetorics of science, which considers how scientists try to persuade, also focuses on this discussion, but it does not commit to the realist position (c.f. McCloskey (1998)).

expected utility maximization against beliefs that a proposed price will be accepted do not perform well, c.f. tables 4.6 and 4.8. The hypothesis is appealing because of its plausibility: it seems only rational for people to act in accordance with their expectations, especially if good expectations can lead to gains. This argument ignores the cost of decision-making, but our calibration points to another issue.

In the Scarf economies monopolistic competition leads to behavior that is observationally different from human trading: buyers propose high and sellers propose low prices, due to their aversion to rejections. In conditional simulations, too many offers are accepted because with Leontief preferences almost all floor prices appear to be bargains, even though by normal standards these prices would be considered as unfavorable. This finding illustrates that the Scarf examples (and the stable economy in particular) are harsh environments, and very well suited as testbeds for behavioral hypotheses. While utility maximization should not be invoked to explain proposed prices, it may still play a role in selecting a preferred opportunity from a set of feasible options. This is why we include it in the calibration of choice among opportunities, c.f. chapter 5.

4.4.2.3 Orbiting

In our simulations, orbiting reflects consistent, global features of price formation, such as one price steadily increasing while the other remains stable or is decreasing. The low frequency of orbits corresponds to a trend that is present over the duration of the experiment (or run). This raises the question whether orbiting in experimental markets is similar to orbiting in a tâtonnement process: in the experimental markets, the long term trend mirrors the absence of (sufficient) negative feedback, while orbiting in tâtonnement theory is the expression of a negative feedback loop.

4.4.2.4 Quantity signals

Do human traders take quantity signals into account, and if so, how and why? On the one hand, the market process is typically understood as communicating price signals only. On the other, rationality demands that agents use all available information.

Quantity information can indeed improve price expectations. By tabulating all observations, Gjerstad-Dickhaut beliefs are capable of correctly adjusting for the practice of accepting n units once instead of n times accepting one unit. This indirectly benefits eGD-expectations. One may wonder what would have happened if human subjects in the experiments of Anderson et al. were forced to accept n units once. Then there would have been fewer price signals and possibly more volatility. Would traders weigh observed prices by quantities traded?

Van der Hoog (2005) argues that agents pick up and use quantity signals that communicate rationing; for instance they can observe unemployment, overcapacity or underutilization of resources. It makes the point that it doesn't require a sophisticated model of the economy to make sense of such signals. This is valid if making sense amounts to becoming aware of a disequilibrium. Deriving further consequences from quantity signals, however, may quickly lead to adopting strong assumptions. Markets in Van der Hoog (2005) are characterized by complete and credible signaling. At any given price, all traders reveal their planned supply and demand, which makes excess demand observable for the auctioneer. After he communicates it to the traders in the form of rationing constraints, the latter can anticipate these quantity constraints.[35]

In a CDA, on the other hand, quantity signaling is incomplete and not necessarily credible. Here, agents become aware of rationing (if at all) if trading comes to a halt before endowments are reset (i.e. due to market failure, c.f. appendix A). Then they may learn rationing constraints that are coarse and not completely credible. Most of all, the quantity signals are also difficult to interpret. A trader of type I may have been unable to buy enough of commodity 3, leaving him with an excess amount of money. Should he conclude that the price he has paid for commodity 3 on average was too low? Not necessarily: he actually may have paid too much for commodity 3, resulting in a utility level that is substantially lower than it would have been in the Walrasian equilibrium. In the absence of knowledge of the Walrasian equilibrium it may be impossible to reach a correct conclusion.[36] In a simple model of exchange with trading at all prices, it seems unjustified to assume that traders will anticipate to be rationed in the next period. Therefore we will ignore this category of quantity signals.

4.4.2.5 Institutions versus intelligence

Gode and Sunder (1993) boldly claims that human intelligence is not necessary for achieving a competitive equilibrium, because a CDA would suffice. Its argument rests on Zero-Intelligence traders obtaining equilibrium in a specific example of a single, simple financial market. As pointed out by Cliff and Bruten (1997), this result breaks down if demand and supply are changed and are no longer symmetric around the equilibrium price. Ladley (2006) is aware of this critique; surprisingly, however, it is also quite generous to the view that institutions matter, whereas intelligence does not.

Although it is ironic that Zero-Intelligence performs better than monopolistic competition in predicting human moves, our results make it clear that ZI-traders do not achieve convergence in a more complicated environment with multiple markets: the simulated confidence intervals for the concentration statistic range from 6% to 10%, c.f. table 4.8. If Gode and Sunder were right, then it also should be less of a challenge to develop an algorithm that achieves convergence similar to human trading in the stable Scarf economy.

Nevertheless, if intelligence falls short of perfect foresight, institutions can certainly offer relief. The high probability of market failure (c.f. appendix A) can provide an economic motive for intermediaries. If long term players can acquire knowledge about the market, then they can offer brokerage services at a cost. This can be a Pareto improvement in so far as brokers succeed in eliminating (part of) the inefficiencies that remain after trading comes to a halt.

---

[35] The auctioneer can be eliminated from the story by assuming that traders have correct expectations with respect to rationing constraints.

[36] Communication between traders through channels other than the market, e.g. comparing endowments or price expectations, may help in interpreting the quantity constraints.

4.5 Conclusions

We have described FACTS, our platform for simulating trading behavior in an exchange economy; FACTS has been developed as part of this thesis. This description is followed by an overview of the algorithms that are used for investigating how people propose prices. Some of these algorithms are taken from the literature, others are variations or newly created. We also have made some improvements to existing algorithms. For details, refer to appendix B.

The inquiry into how people propose prices can be interpreted as part of the calibration of FACTS. We have proposed a methodology that derives criteria for assessing algorithms from the results of Anderson et al. (2004), and that details how these criteria should be applied.

The clearest result of the calibration is that the subjects of Anderson et al. did not behave in accordance with the hypothesis of monopolistic competition. Instead of setting prices by optimizing expected utility against beliefs that a proposed will be accepted, people use reservation prices that are anchored by expected prices. This is reminiscent of price taking, albeit in a more active form because traders face incomplete and false signals. This result agrees with the stylized facts that have been proposed in chapter 3.

Results with respect to how people form price expectations are currently inconclusive. Our methodology suggests multiple criteria and different algorithms perform best along the different dimensions. We believe that there are several explanations for this outcome. Expectation formation of our robot traders is relatively rigid; trading behavior possibly requires more sophistication and the capability to switch between different approaches.

However, before refining the algorithms a more mundane problem must be solved: our robot traders do not generate enough transactions and that adversely affects convergence of prices in the stable Scarf economy, and most likely also the convergence in allocation in each of the three Scarf economies. Our robot traders accept at most one unit; therefore it is less likely that the simulations will generate patterns in which a single robot trader accepts multiple units in successive moves (as human traders sometimes do). The motivation for our design is threefold: (i) initial simulations with quantity setting based on expected utility maximization quickly ran into market failure (c.f. appendix A); (ii) the subjects of Anderson et al. also often decide to accept only one unit; and (iii) technically, it is less complex if robot traders propose and accept at most one unit. With hindsight, however, it is clear that a proper calibration of expectation formation requires that robot traders generate an amount of trades that is comparable to human traders.

We consider the eGD-algorithm currently to be the best and we will apply this algorithm in chapter 5 for generating sets of perceived opportunities, from which robot traders have to select a preferred alternative. This analysis will be restricted to cases in which the actual choice of a human trader is recognized as a feasible choice. This allows us to compare the performance of different selection rules in predicting human moves.