



UvA-DARE (Digital Academic Repository)

Price discovery with fallible choice

Ruiter, A.G.J.M.

[Link to publication](#)

Citation for published version (APA):

Ruiter, A. G. J. M. (2018). Price discovery with fallible choice

General rights

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

Disclaimer/Complaints regulations

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

Appendix B

Algorithms

Trading behavior consists of selecting a preferred alternative from a set of perceived opportunities that largely depend on price expectations. This appendix prepares for the calibration of expectation formation (c.f. chapter 4) and of choice from a set of feasible actions (c.f. chapter 5). It discusses the identification of feasible actions (section B.1.1), the representation of opportunities as prospects (section B.1.2) and simple rules of thumb for selecting a best option (section B.1.3). Section B.2 describes individual algorithms and the pre-calibration of specific parameters (if any). Finally, section B.3 concludes.

B.1 Common elements

B.1.1 Perception of opportunities

Traders can submit proposals and accept or cancel pending offers. Their behavior can be described as follows:

- Submitting proposals:
 - There is no short-selling: traders must own the commodity they propose to sell; buying requires money.
 - Traders can submit an offer if there is no pending rival offer or else if they improve upon the floor price.¹
- Accepting offers:
 - Floor offers will be accepted if their price is considered favorable relative to expected prices.
 - An acceptance of an ask implies that the trader means to keep the commodity: conditional on expected prices there should be no need to sell

¹Traders can submit strategic offers that would create arbitrage opportunities if accepted. The purpose of a strategic offer is either to create such an arbitrage opportunity or else to depress ask prices / inflate bid prices.

(part of) the acquisition at a later time. Similarly, acceptance of a bid draws on excess stock: conditional on expected prices there is no need to buy back sold goods at a later time.²

- Canceling offers:
 - Pending offers will be canceled if (i) the trader has become inactive, (ii) their offered price is deemed unfavorable relative to current expectations, or else if (iii) the price is favorable but acceptance would reduce the level of utility (because of large quantities).^{3,4}

Opportunities are structured into different categories: urgent, eager and patient. Urgent responses always precede eager or patient responses, and patient responses are always slower than eager or urgent responses. Cancellations by inactive traders are treated as urgent (inactive traders do not need to digest the latest information from the auctioneer for arriving at their decision). Owners of pending offers, who are still active, are assumed to adopt a wait-and-see attitude with respect to their floor offer; they will not improve it unless (i) there is a counter-offer (not being their own strategic offer) and (ii) they themselves are the marginal trader (this will be the case if a patient response is the first to reach the auctioneer). Other actions are classified as eager.⁵ Traders select an action from a (non-empty) category with the highest urgency: urgent actions are preferred to eager actions, and the latter dominate patient options. If all categories are empty, then the trader waits.

B.1.2 Opportunities as lotteries

Each feasible action gives a decision-maker access to an uncertain future. We can model this by letting traders choose between different lotteries.⁶ The explicit valuation of individual opportunities requires agents to have beliefs with respect to the probability that an offer will be accepted. For instance, if a trader owns a pending offer with an outdated price, then he must know the probability that this offer will be accepted and that he will incur costs as a result of it. Or to take another example, the value of a strategic offer depends on whether it will be accepted or rejected (see below). Although both cases depend on the probability that an offered price will be accepted, there is a notable difference between these two examples. If a floor offer is not canceled and not accepted in the next iteration, then a new opportunity for

²Arbitrage is considered in chapter 5.

³By definition, inactive traders are at an individual optimum. Giving up one commodity in exchange for another would lower their level of utility; doing so implies arbitrage.

⁴Pending strategic offers can be canceled if they preclude regular offers (by assumption, traders with a pending strategic offer hold on to it until others improve upon the offered price).

⁵Gjerstad and Dickhaut (1998) asserts that agents with large potential gains tend to trade sooner rather than later. This hypothesis may affect the robustness of convergence in the stable Scarf economy; it is discussed separately in section 4.3.2.3. Elsewhere this queuing rule is not in effect. If one would classify acceptances as urgent actions then that would (i) increase the number of acceptances, but also (ii) increase the number of times that there are no floor offers, because the MUDA-rules stipulate that an acceptance erases a rival pending offer. This could make price formation less stable.

⁶For the definition and notation of lotteries or prospects, see section 5.2 on page 90.

canceling arises. On the other hand, if a strategic offer is not accepted immediately then it may be accepted later. For assessing the costs of keeping a pending offer with an unfavorable price we need the probability that it will be accepted in the next iteration; and for valuing a strategic offer we need the probability that it will be accepted before it is rejected.

As a matter of notation, let \bar{p}_j^b and \bar{p}_j^a be the floor bid and ask price (if available; $\bar{p}_j^b < \bar{p}_j^a$), and let p_j^b and p_j^a be proposed bid and ask prices. A strategically offered price is sometimes referred to as p_j^s and an expected price as p_j^e .

Suppose that trader i owns the floor bid, \bar{p}_j^b , and that he believes that offers by others are generated as follows: (i) with probability $\frac{1}{2}$ draw either a random bid or a random ask from the belief distribution; (ii) ignore offers with prices that do not improve upon the floor offers, i.e. ignore $p^b \leq \bar{p}_j^b$ and $p^a \geq \bar{p}_j^a$ (if present); let the process run until either $\bar{p}_j^b < p^b$ (a rejection) or else $p^a \leq \bar{p}_j^b$ (an acceptance) is observed. Then, the probability of bid \bar{p}_j^b being accepted is:

$$\mathcal{B}[\bar{p}_j^b] = \frac{P_A[p^a \leq \bar{p}_j^b]}{P_A[p^a \leq \bar{p}_j^b] + P_B[p^b > \bar{p}_j^b]}. \quad (\text{B.1.1})$$

If $P_A[p^a \leq \bar{p}_j^b] = P_B[p^b > \bar{p}_j^b]$ then we call \bar{p}_j^b a "no arbitrage" price, because there is no advantage in terms of the probability of acceptance. Expected prices of the eGD-algorithm are determined this way. For valuing the option of keeping a bid with an unfavorable price \bar{p}_j^b , we can use $P_A[p^a \leq \bar{p}_j^b]$ and for valuing a strategic offer at price p_j^s we can use either $\mathcal{B}[p_j^s]$ or $\mathcal{A}[p_j^s]$ as the probability that the offer (a bid or an ask respectively) will be accepted before it is rejected.

Suppose that a trader submits a bid that is subsequently rejected. How does he respond to this event? The trader could try and re-submit the same offer and expect to have practically the same chance of it being accepted. In a Continuous Double Auction rejections happen and do not necessarily imply that a trader should offer a higher bid price, or a lower ask price. If his beliefs are reasonably accurate, then repeated trials quickly will lead to a success, i.e. an acceptance.⁷ This suggests that traders can simplify and assume that proposals will be accepted with probability one, simply because they can re-submit proposals indefinitely. However, traders will generally be uncertain about whether their price expectations are correct. After observing a number of rejections, they will begin to believe that their estimate is either too optimistic or else that trading is coming to a halt.

We can model this by assuming that (i) traders are willing to submit the same proposal for a limited number of times; and (ii) that they are prepared to trade only small quantities for containing adverse effects of incorrect expectations. An outcome, in which some planned trades cannot be completed at currently expected prices, then becomes a distinct possibility. Put differently, uncertainty with respect to price expectations makes traders aware of and sensitive to the economy being out of equilibrium.

Let k be the maximum number of trials to get a proposal accepted, and let the probability of acceptance at the proposed price be π (dropping subscripts for traders

⁷There is no real need for discounting time, given that this trading model is valid in a very short run of, say, 15 to 30 minutes.

and commodities). Since the trader is prepared to propose the same price at most k times, the probability π is fixed during these trials. The probability of no acceptance in k trials is $(1 - \pi)^k$; hence, the probability of having at least one acceptance in k trials is $(1 - (1 - \pi)^k)$. However, after the first success, the trader resets the counter for trying to get a proposal accepted. The probability of having one acceptance then is equal to (i) not failing in the first set of at most k trials *and* having no success in the second set of k trials, i.e.

$$(1 - (1 - \pi)^k)(1 - \pi)^k.$$

More generally, the structure of this waiting problem is one of a geometric distribution. The probability distribution of the number of transactions the trader will be able to complete in a particular market given his price expectations is

$$P[m \text{ acceptances}] = (1 - (1 - \pi)^k)^m (1 - \pi)^k$$

with $m = 0, 1, 2, \dots$. If a trader proposes to trade at "no arbitrage" prices then $\pi = \frac{1}{2}$. This probability distribution applies to both markets and it allows traders to value actions based on the implied expected utility.

In valuing a proposal \mathbf{x} , traders assume that it will be accepted and then they evaluate the consequences in terms of the expected utility they can achieve from that position. Accepting \mathbf{x} yields utility $u_{00} = u(\mathbf{w} + \mathbf{x})$. Conditional on acceptance and on selling at most $\mathbf{q}^2 = (p_2^e, -1, 0)'$ in market 2 and buying $\mathbf{q}^3 = (-p_3^e, 0, 1)'$ in market 3, a trader of type I knows how many transactions he is removed from his utility target, $u^* = v(\mathbf{p}^e, \mathbf{w})$ with $v(\cdot)$ the indirect utility function. Suppose there are r transactions in market 2 and c transactions in market 3 required for closing the gap between u_{00} and u^* . Feasible outcomes can then be represented by a table, for instance

	0	1	..	c
0	u_{00}	u_{01}	..	u_{0c}
1	u_{10}	u_{11}	..	u_{1c}
..
r	u_{r0}	u_{r1}	..	u_{rc}

with $u_{rc} = u^*$ and $u_{fg} = u(\mathbf{w} + \mathbf{x} + f\mathbf{q} + \hat{g}(f)\mathbf{q}^3)$. The entries in the table give the best values for utility, given the number of transactions the trader is able / willing to complete in each market.^{8,9} There are feasibility constraints that must be observed: a trader may have to sell before being able to propose another bid. If this is the case,

⁸To see why the number of transactions in market 3 is a function of the number of transactions in market 2 for a trader of type I , consider the following. The best outcome will be obtained if the trader first plans to sell before buying. That way, he can plan buying conditional on the information of how much he is able to sell. The number of bids the market can absorb is an upperbound for the planned number of bids and not the actual number. The column index, g , of u_{fg} depends on being in row f . Then, the expected value of receiving extra money will always be positive. Otherwise the expected value of receiving a windfall gain could be negative, which does not make sense.

⁹Best values also depend on trading the maximum amount (here, by assumption one unit) whenever this is possible. The last transactions, r and c may involve amounts smaller than 1. If this is the case, then the formula $\mathbf{w} + \mathbf{x} + f\mathbf{q}^2 + \hat{g}(f)\mathbf{q}^3$ needs a minor modification.

then utility does not increase even though the market can absorb more bids within the tolerance of the trader.

This approach allows a trader to estimate his expected utility, conditional on the acceptance of \mathbf{x} and given the above probabilities of being able to complete any number of transactions. As a matter of fact, he can assign a value

$$s(\mathbf{x}) = \sum_{f=0}^r \sum_{g=0}^c u_{fg} P[f \text{ acceptances}] P[g \text{ acceptances}] - u(\mathbf{w})$$

to submitting \mathbf{x} . Given the valuation $s(\mathbf{x})$, and the probability that \mathbf{x} will be accepted, say η , we can form the lottery $\mathcal{L}_0 = (s(\mathbf{x}); \eta)$. Then we still have to consider the probability that the trader will have access to \mathcal{L}_0 ; this depends on who is first to propose.

Assume that traders perceive the market as depicted in figure B.1. This model consists of a compound lottery that also requires the specification of $\mathcal{L}_1.. \mathcal{L}_4$.¹⁰ If a trader has no pending offer with an unfavorable price then the value of a particular action is $(\mathcal{L}_0; \frac{1}{15})$. That is, it doesn't matter to the current decision what happens if another trader is quicker to respond: $\mathcal{L}_k = (0; 1)$, $k = 1..4$. If there is a pending offer with an unfavorable price, however, it does matter what happens. In that case, any action other than canceling also implies keeping the pending offer and running the risk that it will be accepted. The value of keeping an offer with an unfavorable price can be determined by applying the approach outlined above; this value then corresponds with an \mathcal{L}_k , with $1 \leq k \leq 4$.

What remains to be done is the valuation of a strategic offer, $\tilde{\mathbf{x}}$. The arbitrage opportunity that results from its acceptance has value $s(\tilde{\mathbf{x}})$. The probability of an offer, say a bid, being accepted before it is rejected is equal to $\mathcal{B}(p^s)$. If the offer is rejected then it will still have helped to improve the next transaction price. For instance, if the strategic offer is a bid then this impact can be estimated as $\frac{1}{2}(p_j^s - p_j^b)$, with p_j^s the proposed strategic price and p_j^b the pending bid price: this is the difference between the old and new average of the pending bid and ask prices. If we define $\mathbf{y} = (\frac{1}{2}(p_j^s - p_j^b), 0, 0)^T$ then we can assign a value of $s(\mathbf{y})$ to this monetary gain. Whether the trader, or some other seller, will appropriate the advantage is a matter of chance. Suppose there are n_j sellers of commodity j , then the probability of reaping the profit is $1/n_j$.¹¹ If another seller benefits from the higher transaction

¹⁰We may ask whether and how the resulting prospects can be simplified? Entropy-sensitive preferences require prizes to be different, c.f. section 5.3. Folding back (or reducing) compound lotteries, however, is delicate if preferences are not linear in probabilities, c.f. Machina (1989); Sarin and Wakker (1994); Laped and Toquebeuf (2007); Dillenberger (2008). Even though the branches of the compound lottery do not represent successive decisions, there is the issue of consistency. Decisions can be sensitive to whether traders perceive opportunities as simple or as compound lotteries. Once again, it matters how people understand their decision problem. FACTS implements both representations, so in principle we will be able to compare choice conditional on a simple and a compound specification. Unfortunately, however, compound prospects take long to evaluate (runs with compound lotteries take about 12 times longer than runs using the rules of thumb of section B.1.3). Therefore we work with simple prospects.

¹¹For simplicity's sake, FACTS sets $n_j = 5$ or $n_j = 6$, ($n_j = 10$ or $n_j = 11$) depending on whether the trader is an "original" seller (buyer) of commodity j . If a trader of type I acquires too much of commodity 3 then he will become a seller of that commodity, after being an "original" buyer of commodity 3.

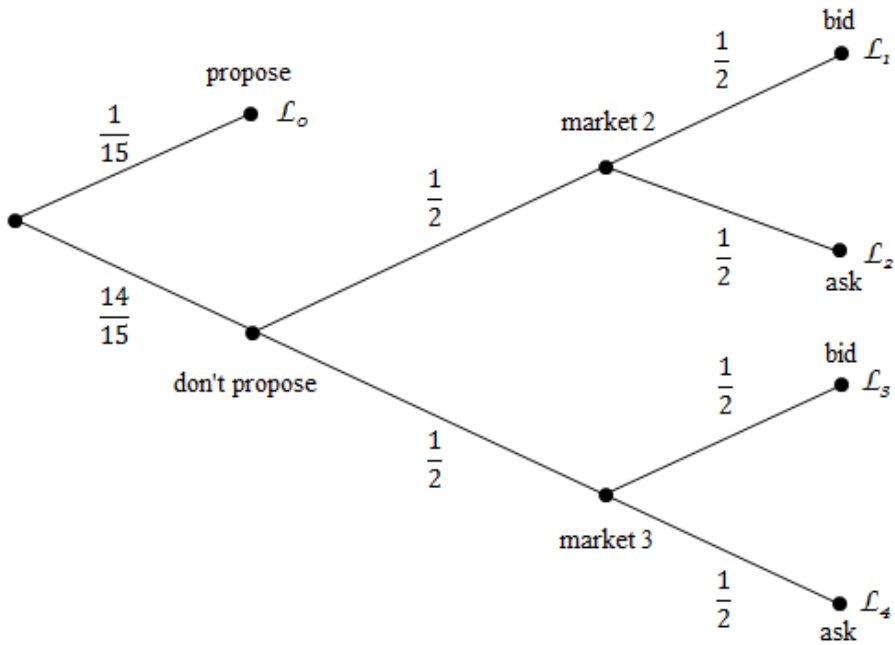


Figure B.1 – Model of offer generation. A trader obtains lottery \mathcal{L}_0 if he is the first to propose; otherwise \mathcal{L}_k occurs ($k = 1, \dots, 4$). Typically these alternatives don't change the status quo, i.e. they yield zero with certainty: $\mathcal{L}_k = (0; 1)$, $k = 1, \dots, 4$. However, if there is a pending offer with an unfavorable price, then one or two of the \mathcal{L}_k represent a loss that materializes if that proposal is accepted. Calibration of the eGD-algorithm suggests that, conditional on not proposing, a trader is agnostic about the market and the type of action (c.f. section B.2.6). This illustrates the fallibility of choice because in the stable Scarf economy there are twice as many buyers as there are sellers: people are either unaware of this fact or else they simply ignore it (assuming that eGD-trading captures their behavior).

price, then submitting a strategic offer still has a positive value because it signals to buyers that there is demand already at p_j^s . However, since it is difficult to quantify, we expect traders to ignore this consequence and set its value equal to 0.¹² With these definitions we can set \mathcal{L}_0 in figure 1 equal to:

$$\mathcal{L}_0 = \left(s(\tilde{\mathbf{x}}), s(\mathbf{y}); \mathcal{B}(p_j^s), \frac{1 - \mathcal{B}(p_j^s)}{n_j} \right).$$

To recap, the three possible outcomes of submitting a strategic offer are (i) the offer will be accepted, creating an arbitrage opportunity; (ii) the offer will be improved

¹²Should we be pressed to make a distinction between bounded rationality and fallible choice, then it would be at a point like this: bounded rationality can avoid a bias by substituting a lottery, even though it is unclear how to define that prospect. Fallible choice, on the other hand, simplifies to circumvent assessments that are too difficult, if necessary by adopting (behavioral) restrictions.

upon and the trader himself is able to reap the benefit; and (iii) the offer will be improved upon and some other trader gets the advantage. The value of a strategic ask can be determined similarly.

Observe that our approach to valuation is greedy in the sense that it gives priority to proposals that raise direct utility the most: a higher value of u_{00} leads to a higher expected value. The approach also favors acceptances over proposals, because acceptances eliminate transactions and hence uncertainty. It is better to act than to wait, because by assumption, waiting yields $(0; 1)$ (provided there is no pending offer with an unfavorable price).

The valuation ignores the possibility that expectations can be revised. Our traders assign a value u_{00} to the event that they are unable to complete any additional transactions at currently expected prices. In this case, however, they will revise their expectations and probably will do better than u_{00} . Instead of assuming that the valuation takes this contingency into account as well, we propose that conditioning on current price expectations is a natural way to simplify the decision problem.

B.1.3 Rules of thumb

If traders perceive multiple opportunities, then they have to choose which alternative is best. Selection of the best urgent or patient actions is random. Determining the preferred alternative from the set of eager options is more complicated. Each particular action can give rise to different futures that can be represented by lotteries (c.f. section B.1.2). Maximization of expected utility and cumulative prospect theory can then be used for ranking different lotteries and selecting a best one. Both methods, however, require substantial computational abilities. As a simpler alternative this section derives rules of thumb. Chapter 5 compares the performance of the different methods for selecting a preferred action by measuring how well each approach predicts the moves of human traders.

The rules of thumb are based on qualitative, heuristic arguments with respect to types of actions.¹³ To illustrate, consider a trader of type *III* with endowments $\mathbf{w} = (350, 0, 2)'$. First, suppose that this trader has the opportunity to either submit a bid for one unit of commodity 2 or to accept an ask which also delivers one unit of commodity 2. How to compare these opportunities? Both actions can deliver an increase of utility of $\frac{1}{10}$: accepting the ask does so with certainty and submitting the bid only if it is accepted (before it is rejected). A simple calculation of the expected value of these actions would lead to a preference for accepting the ask. However, such a calculation ignores the difference between the lower price that could be proposed and the higher floor price. Properly ranking the actions requires that the certain gain of $\frac{1}{10}$ utils be compared with $\frac{1}{10} \times \mathcal{B}(p_2^e)$ plus the expected value of owning an additional budget of $(\bar{p}_2^a - p_2^e)$.¹⁴

¹³There is no pre-ordained way of how decision-makers frame their decision problems. Indeed, defining and understanding what needs to be decided is a critical element of taking a decision. The proposition of Prospect Theory and of behavioral economics, that people analyze choices in terms of changes, is convincing because it is often easier to rank alternatives than it is to explicitly value them on a fair and consistent basis.

¹⁴The "additional budget" only exists in the comparison of proposing a bid and accepting an ask, provided the acceptance is taken to be the reference point. If the proposal would serve as the

The latter can be estimated as the indirect utility of this budget at expected prices:

$$\frac{\alpha_2 \alpha_3 (\bar{p}_2^a - p_2^e)}{\alpha_2 p_3^e + \alpha_3 p_2^e} \times \mathcal{B}(p_2^e).$$

For the sake of argument, suppose that traders analyze their options assuming "no arbitrage" prices, then $\mathcal{B}(p_2^e) = \frac{1}{2}$. Setting $\alpha_2 = 2\alpha_3 = \frac{1}{10}$ as in the Leontief utility function, we find in this case that proposing to buy at p_2^e is better than accepting at \bar{p}_2^a if $\bar{p}_2^a > 2(p_2^e + p_3^e)$. Hence, unless the ask price is extremely high compared to expected prices, it is better to accept than to propose at a better price. We will treat this as a rule of thumb: accept before proposing another offer. Behavioral economics maintains that losses hurt twice as much as equal-sized gains give pleasure, c.f. Tversky and Kahnemann (1974); Thaler (2009, 2015). We adopt this assumption to arrive at another rule: cancellations will be preferred to acceptances, because cancellations typically prevent losses.¹⁵

Next, consider the possibility that the trader can bid for one unit of commodity 2 or for two units of commodity 3. The latter action anticipates a utility level of $\frac{1}{5}$, which requires that two units of commodity 2 are bought at a later time. The bid for one unit of commodity 2 can be embedded in a scenario which also yields $\frac{1}{5}$ utils and which consists of the same proposed trades, albeit in a different order. Then, the expected values of the scenarios are the same. However, traders may well prefer the scenario which starts with the bid for commodity 2, because it raises utility sooner rather than later. This preference does not depend on specific beliefs with respect to acceptable prices. We'll assume that most people are greedy with respect to utility, which leads to another rule of thumb: propose offers which, if accepted, raise utility directly.

How to compare regular and strategic offers? The value of submitting a strategic offer is to create an arbitrage opportunity or to improve the floor price. There is an immediate impact on the next transaction price, but it also sends a signal which affects expectations: supply or demand already exists at the better (strategic) price. The long term effect most likely is more valuable, but it is difficult to quantify. The immediate value of the strategic offer depends on whether it is accepted or rejected.

reference point then the acceptance would entail a supplementary loss. The fact that the acceptance yields a certain increase in utility, however, intuitively makes it the more likely reference point. In pairwise comparisons, agents can choose a reference point at will and different comparisons may be based on alternative reference points (as in the text below). Comparing three or more alternatives at the same time quickly becomes more complicated due to consistency requirements.

¹⁵This is different from the valuation of prospects: suppose there is a pending offer with an unfavorable price; conditional on being the first to propose, a cancellation yields a prize of zero while submitting an offer will yield an expected gain. Both actions run the same risk of the floor offer being accepted. Therefore, proposing a new offer will dominate canceling the offer with an unfavorable price. What we have here is an incentive compatibility problem: in valuing a cancellation, the decision-maker uses a short horizon because the next iteration presents another opportunity to cancel the pending offer.

Upon reflection, however, a trader may conclude that the chance of incurring a loss actually is much higher if he always prefers to postpone canceling an offer with an unfavorable price. That could be another rationale for the rule of thumb (other than loss aversion). Observe that this type of more sophisticated reasoning does not fit into the standard framework of maximizing expected utility. In a context of sequential choice, the latter can even be characterized as an example of bounded rationality.

If rejected, it will have improved the next transaction price. It is uncertain whether the trader can appropriate this gain, or whether someone else seizes the advantage, but the expected value in this case is positive. If accepted, then this leads to a monetary gain of $(p_j^s - p_j^e) x_j$ or $(p_j^e - p_j^s) x_j$, depending on whether it is an ask or a bid; either way, this gain is also positive. It seems that there is no downside to submitting a strategic offer. However, if that were the case, then we should have seen more strategic bids in the experiments of Anderson et al. (2004). Apparently, the value of a strategic bid, if accepted, is not just $(p_j^e - p_j^s) x_j$; there are costs attached to giving up money.¹⁶ Suppose that traders only submit strategic bids when they are really sure that they can sell the commodity at a better price later, then the expected value of strategic offers is always positive and then they should generally be preferred to regular offers, because any regular offer is dominated by a scenario in which the same regular offer is preceded by a strategic offer.¹⁷

Interestingly, however, human traders seem to prefer regular to strategic offers: by reversing the preference between regular and strategic offers, the overall percentage of correctly predicted offers goes up from 38% to 66% in the stable Scarf economy and from 41% to 68% in the counter clockwise unstable economy, c.f. table B.1. One possible explanation could be that traders narrowly focus on available alternatives. Floor bids will often be at least half the expected price. A strategic bid price will lie between this floor and the expected price, leading to an expected monetary gain which is less than half the expected price. Selling one unit, on the other hand, will yield the full expected price if it is accepted (and following the earlier considerations, traders could believe that this happens with probability $\frac{1}{2}$). Hence, traders could conclude that it is better to submit a regular ask than a strategic bid. This argument, however, ignores the fact that strategic offers do not preclude regular offers, while the latter do affect the opportunities for submitting a particular strategic offer later on.

The rules of thumb can be summarized as follows:

- within each single market, the best action is selected according to: cancellations \succ acceptances \succ regular offers \succ strategic offers (\succ expresses a strict preference);
- if both markets admit a cancellation, then one is selected at random;
- if both markets have an acceptance as their best action, then traders prefer the one which yields the highest utility level; ties are broken depending on trader type (for instance, traders of type *I* (*II*) prefer to sell their stock of commodity 2 (3) before buying more of commodity 3 (2); traders of type *III* randomize between accepting asks for commodities 2 and 3);

¹⁶About 75% of the strategic offers would lead to a lower level of utility if accepted. If traders would expect these offers to be rejected, then there would be no reason for having less strategic bids than asks. If, on the other hand, traders expect that some strategic offers will be accepted, then apparently they worry more about being able to sell at better prices at a later stage than about being able to buy at better prices. This asymmetry could be due to a preference for money, that goes beyond what is specified in the utility function.

¹⁷A scenario, in which the regular offer is *followed* by the strategic offer, would also dominate the regular offer, provided that the opportunity to submit this strategic offer still exists. The latter, however, is less likely, which would make this scenario also dominated by the one in which the strategic offer is submitted first.

- an acceptance in one market dominates a regular proposal in another market;
- if both markets have a regular offer as their best action, then traders prefer the proposal that increases utility most;
- a regular offer in one market dominates a strategic proposal in another market;
- if both markets have a strategic offer as their best action:
 - and these are either both bids or both asks then traders select the market in which the relevant floor price is furthest removed from the expected price (ties are broken at random);
 - otherwise, the strategic ask is preferred to the strategic bid.

Does behavior like this recognize the actual offers of human traders as opportunities? To find out, we must make preliminary assumptions about how expectations are formed. A simple approach is to have traders use an expected moving average with individual random weights. The latter prevents that all expectations are perfectly correlated (each trader observes the same history).

For the calibration of price expectations, all robot traders use the same methods for perceiving opportunities, selecting the best alternative, haggling and for quantity setting. With eEMA-expectations, for instance, about 75% of human actions is correctly recognized as an opportunity; and, conditional on recognition, about 70% of the actions is also correctly predicted (c.f. table B.1). These rates are averages over 1,000 runs and apply to both the stable and the counter clockwise treatment.

The recognition rates of bids and asks are an average of regular and strategic offers. Opportunities for regular offers are more or less correctly identified: 96% (90%) in the stable (unstable) example. The recognition rate of strategic moves on average is 68% (62%) in the stable (counter clockwise unstable) economy, which is still reasonable. Opportunities for cancellations and acceptances are poorly detected. Probably, this is due to differences between human and algorithmic price expectations. The (unfavorable) difference between expected and accepted prices frequently is substantial (well in excess of 10%) and this may prevent some traders from updating their initial price expectations.¹⁸

It is, of course, highly unlikely that a trader dismisses almost all observed prices as implausible. Therefore, one may argue that either the initialization of price expectations or the plausibility weights are not properly calibrated. We expect that human trading is quite robust; in particular we do not assume that convergence, as

¹⁸Since we try to understand human trading as self-interested behavior, we ignore the possibility that traders exchange simply for the reason of actively participating in the laboratory experiment. Most likely, the poor recognition rate is due to initial expectations. In a single simulation run with the stable treatment there were 112 acceptances of a bid for commodity 2 by human traders that were not detected because the bid price was well below the expected price of the robot traders. Of these, 80 were due to a single trader who initially expected the price of commodity 2 to be around 83. As a result of applying plausibility weights, this robot trader did not lower his expected price, while his alter ego most likely did. The average recognition and matching rates have been calculated over 1,000 runs. Although that does reduce the impact of this particular example, there will always be outliers. If their human counter parts are active traders, then their behavior won't be properly understood. Furthermore, the practice of accepting n times one unit rather than one time n units amplifies the discrepancies.

Table B.1 – Perception and selection of opportunities

action	stable economy			ccw unstable economy		
	freq	detected (%)	predicted (%)	freq	detected (%)	predicted (%)
propose bid 2	732	95	60	696	89	80
propose bid 3	554	96	66	681	96	66
propose ask 2	760	94	58	734	87	41
propose ask 3	964	93	67	1043	90	87
accept bid 2	189	24	99	116	27	88
accept bid 3	80	31	80	202	21	100
accept ask 2	406	30	94	251	34	100
accept ask 3	1006	49	88	474	30	93
cancel bid 2	26	81	100	55	52	100
cancel bid 3	7	57	100	16	22	100
cancel ask 2	26	19	100	32	47	100
cancel ask 3	78	70	100	44	46	100
total	4828	75	69	4344	74	73

Figures are averages over 1,000 runs of eEMA with $\omega = 0.3$ and $\mu = 0.00$ for both the stable and ccw treatment. Actions, which are thought to refer to arbitrage, have been excluded. The percentages, which express the extent to which algorithmic moves correctly predict human moves, are conditional on the human actions being recognized as feasible options. Conditional on recognition, acceptances and cancellations receive appropriate priority. The selection of bids and asks is less good. This may be due to incompletely recognizing arbitrage and to the relatively crude character of the rules of thumb. The performance of the rules of thumb is consistent across the economies.

observed in session 414, is due to initial price expectations being very concentrated. Indeed, initial trading prices show a lot of volatility. For a discussion of the flexibility of price expectations, see section 4.4.1.

B.1.4 Other elements of robot trading

B.1.4.1 Acceptable prices

All prices p_j are assumed to lie between bounds m_j and M_j . In the Scarf economies, for both commodities the values are set to $m_j = 1$ and $M_j = 200$. Having such bounds is convenient, e.g. for implementing Zero-Intelligence traders and for tabulating beliefs. The bounds also prevent unstable behavior when amounts are divided by prices. The bounds, which are treated as a given by all algorithms, are not subject to calibration. Prices are not integer-valued (for making it somewhat more challenging to reach convergence).

B.1.4.2 Haggling

In FACTS, traders can haggle.¹⁹ If they exist, let α_j be the floor ask and let β_j be the floor bid for commodity j , and suppose that the spread is sufficiently large, i.e. $\beta_j < 0.95\alpha_j$. If the proposed price covers more than half of the current spread, then a haggling trader will adjust his shout (offer). For instance, for a robot trader contemplating to propose selling at price \tilde{p}_j with $\alpha_j > \tilde{p}_j > \beta_j > 0$:

$$p_j = \begin{cases} \theta\beta_j + (1 - \theta)\alpha_j, & \tilde{p}_j \leq (\beta_j + \alpha_j)/2 \\ \tilde{p}_j, & \text{otherwise.} \end{cases}$$

Here, $\theta \sim \text{uniform}(0.2, 0.5)$. The maximal weight of 0.5 is due to the fact that the seller is not prepared to "cross" more than half of the current spread if his reservation price is closer to the floor bid. The lower bound of 0.2 makes sure that the haggling process does not take too long. Buyers behave similarly. If the floor bid and ask do not exist simultaneously, then the agent checks the presence of a rival floor offer: if it is absent then buyers will propose $\theta\tilde{p}_j$ and sellers \tilde{p}_j/θ , with $\theta \sim \text{uniform}(0.5, 0.8)$. If a rival offer is available (while the opposite floor offer is not), then traders improve upon the floor offer by, say, at most 5%, e.g. a seller puts

$$p_j = \begin{cases} (1 - \theta)\alpha_j, & \tilde{p}_j < 0.95\alpha_j; \\ \tilde{p}_j, & 0.95\alpha_j \leq \tilde{p}_j < \alpha_j; \end{cases}$$

with $\theta \sim \text{uniform}(0, 0.05)$. Afterwards, shouts are constrained to lie between m_j and M_j . These rules seem reasonable, and since they are common to all algorithms which use haggling the various distributions will not be optimized.

B.1.4.3 Quantity setting

In the experiments of Anderson et al. (2004) human traders often proposed to exchange small amounts and quantities traded typically were one unit for both commodities 2 and 3. There are two possible explanations for this: (i) trading in small amounts is prudent given the uncertainty over prices; (ii) the software which was used by Anderson et al. may have made it easier for traders to accept n times one unit rather than one time n units, c.f. section 3.3. The latter consideration would justify that robot traders propose quantities bigger than one. However, we have all robot traders proposing and accepting quantities of at most one unit. In so far as the speed of convergence depends on quantities, all algorithms face the same constraint. Now, if one algorithm converges faster than another, it must be due to how prices are processed and not to any differences in quantity setting.

The drawback of this design choice is that algorithmic price formation will be different from human trading in the sense that there will be no series of n acceptances of one unit of a commodity coming from the same trader. First because the quantity on offer is at most one and second because after each acceptance the auctioneer

¹⁹Haggling only occurs when traders propose a regular offer; if they submit a strategic offer these rules do not apply.

informs all traders, and there is only a slim chance that the same trader repeatedly will submit the next offer.²⁰

Quantities can be smaller than one, for instance if traders do not have enough money to buy a full unit, or if they do not need a whole unit to obtain an optimum, or if a floor offer has been partly accepted previously.

B.2 Pre-calibration

Several algorithms that are considered here are taken from the literature (ZI, ZIP, eEMA, AA, GDA). Minor changes are sometimes necessary, because they have been developed for a simple financial market, with one asset and money and with exogenous reservation prices. We propose improvements to the so-called Adaptive-Aggressive and Gjerstad-Dickhaut algorithms. In addition, we add some variations and new algorithms.

The goal of the pre-calibration is to fix parameters, so that the various algorithms describe the experimental data as good as possible. In the calibration of expectation formation, they have to compete against each other (c.f. chapter 4). Parameters include the source of price expectations, haggling and the markup (c.f. section 3.4.2) and also parameters which are specific to an algorithm (if any). The objective is to select values that are good for all three examples of Scarf. There is no compelling argument why expectations would depend on offered prices in the unstable economies and on trading prices in the stable Scarf economy. A similar objection applies to the question whether traders haggle or not. On the other hand, differences in volatility may justify differences in the markup.²¹ Therefore, we do allow different values of the markup in the stable and unstable treatments.

The pre-calibration relies on the same methods as used in chapters 4 and 5: (i) visual inspection and confidence intervals of concentration statistics for assessing convergence; (ii) shifts in the distribution of cumulative movements of the so-called clock hand model for detecting orbiting; (iii) the distance between the algorithmic and human (synchronized) prices for gauging similarities in prices; and (iv) the recognition of actual human moves as feasible options and the extent to which human moves are correctly predicted for discriminating between behavioral hypotheses. For a discussion of these methods, see section 4.2.

²⁰The latter may affect the average distance between human and robot prices, which is one of the measures which is used in the calibration. In the fully-scripted simulations, human offers are executed including their proposed quantities. If one trader offers to sell, say, five units of commodity 2, then this allows another trader to submit five subsequent acceptances of one unit each. Although possible, it is unlikely that the same robot trader will be selected for submitting an offer five times in a row. As a result, one should expect more discrepancies between prices as proposed by human and robot traders than would have been the case if robot traders behaved in the same way as humans in terms of quantity setting.

²¹Such differences may arise through competitive pressure, c.f. 3.4.2. This is also the basic idea underlying the Adaptive-Aggressive algorithm, which endogenously adjusts to differences in volatility, c.f. section B.2.5.

B.2.1 Zero Intelligence and random expectations

B.2.1.1 Zero Intelligence (ZI-)trading

Gode and Sunder (1993) has introduced Zero Intelligence (ZI-) traders. These are characterized by (i) treating all feasible opportunities as equally good and (ii) randomization of the determinants of feasibility (such as price expectations, c.f. equation (B.2.1)). In FACTS, ZI-traders pool all feasible urgent, eager and patient actions and select one alternative at random. If they have no feasible action, then ZI-traders wait. Given common behavior, as described above, feasibility mainly depends on price expectations. Let m_j and M_j be the minimum and maximum price for commodity j and let $\theta_{ji}^t \sim \text{uniform}(0, 1)$. Then the expected prices, p_{ji}^t , $j = 2, 3$, of ZI-trader i at iteration t are

$$p_{ji}^t = \theta_{ji}^t m_j + (1 - \theta_{ji}^t) M_j. \quad (\text{B.2.1})$$

ZI-traders have no memory and they do not learn. By assumption, they quote their expected prices, i.e. they do not haggle. Since m_j and M_j are treated as given for all algorithms, there are no other discretionary parameters to calibrate. ZI-trading is useful as a diagnostic: if other robot traders perform worse than ZI-traders, then that signals a bias, which most likely is unintended and due to an implementation error.

B.2.1.2 Random expectations (eRnd)

The eRnd-traders are a variation on ZI-traders. They have random price expectations as in equation (B.2.1), and they know how to haggle. Contrary to ZI-traders, eRnd-traders do not pool all available opportunities and they do not select a feasible action at random.²² Enabling haggling makes the prices proposed by eRnd-traders more similar to human traders: with haggling the average distance between robot and human offers is 9.2 units of money, which is slightly better than the 9.5 units that apply to no haggling. Haggling will be used subsequently by all other robot traders, except for the belief-based algorithms (these set prices based on expected utility maximization against their beliefs).

B.2.2 Adaptive expectations (eEMA)

An early application of adaptive expectations is Nerlove (1958). FACTS implements traders who update their expectations, $p_j^e(t)$, by taking an (exponential) moving average of plausibility weighted observed prices as in equation B.2.2

$$p_j^e(t+1) = \left(1 - \theta^t e^{-\left(\frac{p_j^{last} - p_j^e(t)}{2\sigma_j^t}\right)^2} \right) p_j^e(t) + \theta^t e^{-\left(\frac{p_j^{last} - p_j^e(t)}{2\sigma_j^t}\right)^2} p_j^{last} \quad (\text{B.2.2})$$

²²In principle these and later algorithms that generate price expectations can be combined with different rules for selecting a best alternative from a set of perceived possibilities. However, since expected utility maximization, cumulative prospect theory and entropy-sensitive preferences require beliefs, all algorithms that yield point expectations must be combined with the rules of thumb for prioritizing feasible actions.

Table B.2 – eEMA: source and weights

Description	trading prices		offered prices	
	stable	ccw	stable	ccw
$\omega = 0.1$	68.5	70.8	68.6	72.0
$\omega = 0.3$	<u>68.7</u>	<u>73.3</u>	68.4	70.8
$\omega = 0.5$	65.9	73.2	66.4	68.9
$\omega = 0.7$	64.9	73.2	65.6	67.9
$\omega = 0.9$	64.1	70.8	65.1	66.9

Percentages of correctly predicted human actions, excluding arbitrage, conditional on recognizing the move as a feasible opportunity. Percentages are an average over 1,000 runs and they are based on markup $\mu = 0.00$ (see text). At $\omega = 0.3$, trading prices dominate offered prices.

with $\theta^t \sim \text{uniform}(\max(0, \omega - 0.2), \omega)$. This latter formula limits the spread of weights between iterations (consistency of behavior over time), and also between different goods and traders (limited disagreement over expected prices). The use of plausibility weights is motivated by the presence of haggling: very low bid prices and / or very high ask prices can make expectations more volatile than what seems realistic.²³ The weights decrease the further an observed price is removed from its expected value. Here, we simplify by putting $\sigma_j^t = 0.25p_j^e(t)$. Even with this basic model, there are different possibilities to consider: does p_j^{last} represent offered or trading prices, and which value of ω to apply? To calibrate these parameters, we look at the percentage of correctly predicted human actions. Table B.2 shows that trading prices should be used and $\omega = 0.3$.

The next step consists of fixing the markup, μ (c.f. table B.3). If the markup becomes more pronounced robot traders will be inclined to accept more offers, mainly at the expense of submitting regular proposals. If human traders are more reluctant to accept then prediction rates will drop. The best values are $\mu = 0.00$ for both the stable and counter-clockwise treatments.

B.2.3 Expectations based on bid / ask spreads (eBAS)

The eBAS-algorithm takes a forward looking approach by deriving expected prices, p_j^e , from current floor offers. Prices are updated if a trader is no longer competitive, e.g. $p_j^e \leq \beta_j$ if the trader is a buyer, with β_j the floor bid for commodity j . Furthermore, traders, who can accept a counter offer but who deem the spread between the current floor bid and ask too wide, also update their expectation (fearing that their current expected price is too high (buyers) or too low (sellers)), c.f. table B.4. If the current spread is acceptable, $\frac{\beta_j}{\alpha_j} \geq \vartheta$, and if $p_j^e > \alpha_j$ then a buyer accepts the floor ask, α_j .

²³Even though the plausibility weights slightly improve both the recognition rate and the prediction of human moves as presented in table B.1, they can make expectations relatively inelastic if traders have extreme initial expectations. With hindsight, a better approach may be to update expected prices if a trader is surprised; for instance, after observing a disproportionate number of rejections or acceptances.

Table B.3 – eEMA: determining the markup

markup	stable (%)	ccw (%)	markup	stable (%)	ccw (%)
$\mu = 0.00$	<u>68.7</u>	<u>73.3</u>	$\mu = 0.06$	53.4	66.8
$\mu = 0.01$	66.7	72.2	$\mu = 0.07$	51.5	65.4
$\mu = 0.02$	63.8	71.1	$\mu = 0.08$	50.1	64.2
$\mu = 0.03$	61.1	70.2	$\mu = 0.09$	49.0	63.2
$\mu = 0.04$	58.4	69.3	$\mu = 0.10$	48.0	62.0
$\mu = 0.05$	55.7	68.1			

Percentages of correctly predicted human actions, excluding arbitrage, conditional on recognizing the move as a feasible opportunity. Percentages are an average over 1,000 runs. The best values for the markup are $\mu = 0.00$ for both the stable Scarf economy and the counter-clockwise unstable economy.

The threshold for the ratio of the floor bid and ask is randomized and constrained to lie between 0.75 and 0.85. These values seem reasonable. A buyer, who is no longer competitive, sets his new expected price equal to $p_j^e = \theta\beta_j + (1 - \theta)(\beta_j + \alpha_j)/2$, i.e. to a random mixture of the floor bid and the average of the floor bid and ask, with $\theta \sim \text{uniform}(0,0.1)$. The relatively low value of 0.1 speeds up learning. If there is only a floor bid (ask), then only buyers (sellers), who are not competitive, update their expected prices relative to the floor (by approximately 30%). If there are no offers, then every trader is competitive and holds on to his current reservation prices. After observing an acceptance, buyers (sellers) lower (raise) their reservation price relative to the accepted price.

From table B.5, we conclude that the markup is $\mu = 0.00$ for both the stable and the unstable treatments.

B.2.4 ZI-Plus

Cliff and Bruten (1997) refutes the claim of Gode and Sunder (1993) that a Continuous Double Auction is sufficient for letting ZI-traders achieve competitive equilibrium in a single financial market with money and one asset, in which traders have to execute limit orders. The ZI-Plus, or ZIP-algorithm, is proposed as having the minimum level of intelligence that is required for obtaining robust convergence to the competitive equilibrium.

Let the offered price, or shout price p_j , be the product of an (exogenous) reservation price and a profit margin. We have $p_j = p_j^r(1 + \varphi_j)$, with p_j the shout price of commodity j , p_j^r the limit or reservation price and φ_j the profit margin ($\varphi_j > 0$ for sellers and $\varphi_j < 0$ for buyers). The ZIP-algorithm specifies rules for updating profit margins. In general equilibrium models, there are no exogenous limit prices. However, by interpreting the ZIP-algorithm as providing rules for updating shouts rather than profit margins, it is possible to apply it in the context of the Scarf economies.

Table B.6 presents the updating rules for shouts.²⁴ For instance, a buyer who is

²⁴Note that an offer can have a status other than accepted or rejected. In that case traders leave

Table B.4 – Learning expected prices in the eBAS-algorithm

Situation	conditions	expectation
<i>acceptance observed</i>		
- buyer	acceptance	$p_j^e = (1 - \theta) p_j^{last}$
- seller	acceptance	$p_j^e = (1 + \theta) p_j^{last}$
<i>bid and ask exist</i>		
- buyer	$p_j^e \leq \beta_j$ or $(p_j^e > \alpha_j) \& \left(\frac{\beta_j}{\alpha_j} < \vartheta \right)$	$p_j^e = \theta \beta_j + (1 - \theta) (\beta_j + \alpha_j) / 2$
- seller	$p_j^e \geq \alpha_j$ or $(p_j^e < \beta_j) \& \left(\frac{\beta_j}{\alpha_j} < \vartheta \right)$	$p_j^e = \theta \alpha_j + (1 - \theta) (\beta_j + \alpha_j) / 2$
<i>only a bid exists</i>		
- buyer	$p_j^e \leq \beta_j$	$p_j^e = (1.3 - \theta 0.3) \beta_j$
<i>only an ask exists</i>		
- seller	$p_j^e \geq \alpha_j$	$p_j^e = (0.7 + \theta 0.3) \alpha_j$

After an acceptance, buyers (sellers) push for lower (higher) prices. Otherwise, eBAS-traders update their reservation prices if they are no longer competitive (e.g. $p_j^e \leq \beta_j$ for a buyer). Suppose that both the floor bid, β_j , and the floor ask, α_j , exist, and that $\beta_j < \vartheta \alpha_j$; then traders, who can accept a counter offer, also seek a better price. If there are no floor offers, then traders keep their expectations. With respect to the random variables we have $\theta \sim \text{uniform}(0,0.1)$ and $\vartheta \sim \text{uniform}(0.75,0.85)$.

Table B.5 – Determining the markup for eBAS

markup	stable (%)	ccw (%)	markup	stable (%)	ccw (%)
$\mu = 0.00$	<u>69.0</u>	<u>72.5</u>	$\mu = 0.06$	56.2	69.2
$\mu = 0.01$	67.4	72.2	$\mu = 0.07$	54.4	68.3
$\mu = 0.02$	65.7	71.8	$\mu = 0.08$	52.8	67.4
$\mu = 0.03$	63.2	71.3	$\mu = 0.09$	51.5	66.6
$\mu = 0.04$	60.7	70.7	$\mu = 0.10$	50.6	66.0
$\mu = 0.05$	58.2	69.9			

Percentages of correctly predicted human actions, excluding arbitrage, conditional on recognizing the move as a feasible opportunity. Percentages are an average over 1,000 runs. The best values are $\mu = 0.00$ for both the stable and the counter-clockwise treatment.

no longer competitive after observing the latest bid price p_j^{last} will increase p_{ji} if p_j^{last} is rejected; on the other hand, he will leave p_{ji} as is if p_j^{last} is accepted. It will not be immediately clear whether a new offer has been rejected or not. As a matter of fact, while acceptance is easy to recognize in the experiments of Anderson et al. (2004), rejection is not because there is no explicit acceptance / rejection step. An offer will

their shouts unchanged, which is why the ZIP-algorithm sometimes decides that a human trader should have waited, c.f. table 4.7.

Table B.6 – Learning shouts in the ZIP-algorithm

Market condition	buyers	sellers
Bid accepted	$p_j \geq p_j^{last}$: lower p_j otherwise: leave p_j as is	$p_j \leq p_j^{last}$: raise p_j active & $p_j \geq p_j^{last}$: lower p_j otherwise: leave p_j as is
Ask accepted	$p_j \geq p_j^{last}$: lower p_j active & $p_j \leq p_j^{last}$: raise p_j otherwise: leave p_j as is	$p_j \leq p_j^{last}$: raise p_j otherwise: leave p_j as is
Bid rejected	active & $p_j \leq p_j^{last}$: raise p_j otherwise: leave p_j as is	leave p_j as is
Ask rejected	leave p_j as is	active & $p_j \geq p_j^{last}$: lower p_j otherwise: leave p_j as is

Cliff and Bruten (1997a) presents the rules in a slightly different form, because it assumes that the market protocol includes an explicit acceptance / rejection step. Here, rejection is implicit.

be considered as being rejected after it has been replaced by a better one. That is, rejection will be understood as failure to accept in time.²⁵ With respect to updating a shout price, p_{ji} , traders have three options: leave it as is, raise or lower it. The adjustment of shouts is as follows:

$$\begin{aligned}
 p_{ji}^{t+1} &= p_{ji}^t + \Gamma_{ji}^t; \\
 \Gamma_{ji}^t &= \gamma_i \Gamma_{ji}^{t-1} + (1 - \gamma_i) \Delta_{ji}^t; \\
 \Delta_{ji}^t &= \kappa_{ji} (\tau_{ji}^t - p_{ji}^t); \\
 \tau_{ji}^t &= R_{ji}^t p_j^{last} + A_{ji}^t; \\
 R_{ji}^t &= \begin{cases} 1 + \lambda^t, & \text{raise } p_{ji}^t; \\ 1 - \lambda^t, & \text{lower } p_{ji}^t; \end{cases} \\
 A_{ji}^t &= \begin{cases} \delta^t & \text{raise } p_{ji}^t; \\ -\delta^t, & \text{lower } p_{ji}^t. \end{cases}
 \end{aligned}$$

The change of p_{ji}^t is determined by a momentum-coefficient, Γ_{ji}^t , which is a weighted average of its previous value and an increment, Δ_{ji}^t . The latter is a correction-factor depending on a learning rate, κ_{ji} , and the deviation between the shout and a target price, τ_{ji}^t . The target price is calculated when the shout price has to be updated, otherwise it is kept fixed. It is a linear transformation of the current offered price, p_j^{last} , with coefficients that are random and which depend on whether p_{ji}^t needs to be raised or lowered. This specification keeps buyers (sellers) pushing for a lower (higher) price. The distributions of γ_i , λ^t and δ^t are taken from Cliff (1997): $\gamma_i \sim \text{uniform}(0.2, 0.8)$ and $\lambda^t, \delta^t \sim \text{uniform}(0, 0.05)$. These distributions will not

²⁵C.f. footnote 11 on page 37. This definition of a rejection is the same as used by Gjerstad and Dickhaut (1998).

be calibrated for keeping our implementation comparable with others. Furthermore, it is clear that ZIP-traders respond to offered prices rather than trading prices (for observing rejections), and also that their rules for adjusting shouts already stipulate how ZIP-traders haggle. That is, with the distributions of γ_i, λ^t and δ^t fixed, there are no other discretionary parameters that need to be calibrated.

B.2.5 Adaptive-Aggressive expectations

Adaptive-Aggressive traders were introduced in Vytelingum (2006); they manage a trade-off between adaptive and aggressive bargaining. Aggressive buyers (sellers) offer prices above (below) the expected equilibrium prices, while passive buyers (sellers) offer prices below (above) the expected equilibrium. Lower prices will give a better price for the proposing buyer, at the expense of the probability that the offer will be accepted. Aggressive bargaining is supposed to occur when uncertainty with respect to prices is great (e.g. after an external shock); when prices are stable bargaining is assumed to be passive.

AA-traders were designed for the simple financial market with given limit prices.²⁶ An intra-marginal buyer maintains reservation prices as follows (dropping the subscripts i and j):

$$p^r = \begin{cases} p^e \left(1 - \frac{e^{-\rho\vartheta} - 1}{e^{\vartheta} - 1}\right), & -1 < \rho < 0 \\ p^e + (\ell - p^e) \frac{e^{\rho\vartheta} - 1}{e^{\vartheta} - 1}, & 0 < \rho < 1 \end{cases}$$

with p^r the reservation (or target) price, p^e the expected (estimated) equilibrium price, ℓ the buyer's exogenous limit price, and ρ the degree of aggressiveness. Traders can set reservation prices below, at or above the expected equilibrium price, depending on ρ . The parameter ϑ must be calculated so as to make the derivative of p^r continuous. The complexity of this formula is due to the presence the exogenous limit price and the desire to let the reservation price be a smooth function of its parameters. In FACTS, this function is implemented as

$$p^r = \begin{cases} p^e \frac{2e^{\vartheta} - e^{-\rho\vartheta} - 1}{2(e^{\vartheta} - 1)}, & -1 < \rho \leq 0 \\ p^e \frac{e^{\vartheta} + e^{\rho\vartheta} - 2}{e^{\vartheta} - 1}, & 0 \leq \rho < 1 \end{cases}$$

This simpler alternative is both continuous and smooth in $\rho = 0$, and it has $\frac{1}{2}p^e \leq p^r \leq 2p^e$. These bounds are expected to be sufficient: bids below $\frac{1}{2}p^e$ practically have a small chance of being accepted; the same holds for asks in excess of $2p^e$. Furthermore, traders will be allowed to haggle, that is if there is no floor bid then

²⁶According to De Luca and Cliff (2011), AA-traders achieve better results than human traders in the simple financial market with one asset and money. It is understandable that developers of trading algorithms want to benchmark the performance of their algorithms against the performance of people. While that may be difficult to accomplish, this thesis shows that it is equally challenging to capture, rather than outperform, human trading behavior.

Since they are designed to do better it makes sense to ask whether the alleged dominance of AA-traders extends to the more complicated environments of the Scarf economies? In order to assess this we compare the efficiency of both human and AA-trading. Efficiency is defined as actual average utility divided by potential average utility. While human traders achieve an efficiency of 0.92 in both the stable and the ccw treatment, AA-traders score 0.79 and 0.66 respectively. Trading in the Scarf economies, once again, proves to be something else.




Figure B.2 – Aggressiveness model. The degree of aggressiveness, ρ , determines the reservation price, p^r . If neutral ($\rho = 0$), the reservation price equals the expected equilibrium price. Aggressive buyers (sellers) have $\rho > 0$ ($\rho < 0$). Different values of θ lead to alternative relationships between ρ and p^r : e.g. if $\theta = 2$ then the reservation price is relatively inelastic in the degree of aggressiveness.

buyers can propose bids below $\frac{1}{2}p^e$ as they please. The model for sellers is the same as the one for buyers; it is shown in figure B.2, with $p^e = 3$.

The target price p^r depends on ρ and θ . The degree of aggressiveness, ρ , is updated through short term learning. Table B.7 gives the rules, with p^{last} representing the latest offered price. At this price, there can be a transaction, or just a bid or an ask. After observing a transaction, all (active) traders update ρ , by either increasing or decreasing it. If an offer is not an acceptance, then buyers update ρ if the offer is a bid, and sellers update if the latest offer is an ask. In Vytelingum (2006), the first step in updating ρ consists of determining ρ_{shout} , this is the degree of aggressiveness that would set the reservation price equal to p^{last} . The desired degree of aggressiveness at t , δ^t , is defined relative to ρ_{shout} : $\delta^t = (1 \pm \lambda)\rho_{shout}$. The purpose of applying the factor $(1 \pm \lambda)$, with λ a small random value, is to make sure that the desired degree of aggressiveness results in an admissible offer.

The implementation in FACTS is equivalent, but slightly different. If an agent has to modify his degree of aggressiveness, then he first determines an admissible

Table B.7 – Short term learning of the rate of aggressiveness

trader	transaction	bid	ask
buyer	$p^r \geq p^{last}$: decrease ρ	$p^r > p^{last}$: do nothing	
	$p^r < p^{last}$: increase ρ	$p^r \leq p^{last}$: increase ρ	
seller	$p^r > p^{last}$: increase ρ		$p^r \geq p^{last}$: increase ρ
	$p^r \leq p^{last}$: decrease ρ		$p^r < p^{last}$: do nothing

The behavior of AA-traders can be characterized as maintaining their competitiveness: if their reservation price is competitive, compared to the latest rival offer, then they do nothing otherwise they adjust; if a transaction price is favorable compared to their target price, then they relax the latter.

offer, p^{adm} , and then calculates the implied desired degree of aggressiveness directly as:

$$\delta^{t+1} = \begin{cases} -1, & p^{adm} < \frac{1}{2}p^e; \\ \frac{-\log\left(1+2(e^\theta-1)\left(1-\frac{p^{adm}}{p^e}\right)\right)}{\theta}, & \frac{1}{2}p^e \leq p^{adm} < p^e; \\ 0, & p^{adm} = p^e; \\ \frac{\log\left(1+(e^\theta-1)\left(\frac{p^{adm}}{p^e}-1\right)\right)}{\theta}, & p^e < p^{adm} \leq 2p^e; \\ 1, & 2p^e < p^{adm}. \end{cases}$$

Given a new value δ^{t+1} , we put $\rho^{t+1} = \rho^t + \kappa_1(\delta^{t+1} - \rho^t)$, with $0 < \kappa_1 < 1$.

The value of θ is the result of long term learning, i.e. θ is updated after each transaction. Let v^t be the volatility at t (i.e. after observing the latest transaction):

$$v^t = \frac{1}{p^e} \sqrt{\frac{\sum_{s=t-N+1}^t (p^s - p^e)^2}{N}}.$$

Volatility estimates differ across individual traders because volatility is calculated relative to different expected equilibrium prices. The updated value of v^t is used to determine a new target value of θ , with the help of the mapping²⁷

$$\theta^*(v^t) = (\theta_{max} - \theta_{min}) \left(1 - \frac{v^t - v_{min}^t}{v_{max}^t - v_{min}^t} e^{\gamma \left(\frac{v^t - v_{min}^t}{v_{max}^t - v_{min}^t} - 1 \right)} \right) + \theta_{min} \quad (\text{B.2.3})$$

with $v_{max}^t = \max_{\tau \leq t} v^\tau$ and $v_{min}^t = \min_{\tau \leq t} v^\tau$ (these variables are updated before the new target is calculated). Finally, the new value of θ is calculated as

$$\theta^{t+1} = \theta^t + \kappa_2(\theta^*(v^t) - \theta^t)$$

²⁷There is some confusion around this mapping: Vytelingum (2006) specifies $\theta^*(v) = (\theta_{max} - \theta_{min}) \left(1 - \frac{v - v_{min}}{v_{max} - v_{min}} \right) e^{\gamma \left(\frac{v - v_{min}}{v_{max} - v_{min}} - 1 \right)} + \theta_{min}$, but in this formula θ^* does not reach θ_{max} as one may expect and as suggested by a graph illustrating the mapping. Equation B.2.3 (which is suggested in <https://github.com/davecliff/BristolStockExchange/wiki/3.9-Robot-AA>) does match that graph.

with $0 < \kappa_2 < 1$. After updating the target price, p^r , it may be the case that p^r is not admissible, because it does not improve upon the floor offer. In that case the AA-trader waits. If, on the other hand, p^r is admissible then the agent may haggle toward it.

As a result of a calibration Vytelingum (2006) sets $\theta_{max} = 2$, $\theta_{min} = -8$ and $\gamma = 2$; we use the same values. The learning parameters, κ_1 and κ_2 , are specific for individual traders and are taken from the uniform(0,1) distribution. This leaves the expected equilibrium prices, \mathbf{p}^e , as the only parameter which still needs to be fixed. Expected prices are calculated as in the eEMA-algorithm (see section B.2.2 for the pre-calibration). In case of AA, the markup is already included in p^r .

B.2.6 Gjerstad-Dickhaut beliefs, with variations

Gjerstad and Dickhaut (1998) proposes belief functions that are based on accepted and rejected offers.²⁸ Basically, the probability that a bid at price p will be accepted is the number of accepted bids below price p divided by the sum of accepted bids below p and rejected bids above p . If sellers propose asks below p then this also signals a willingness to accept bids below p . Here we will use quantities accepted and rejected, instead of frequencies. This takes care of divisible commodities and also of exchanging of multiple units of a commodity.²⁹ For each commodity there are separate beliefs that a bid and an ask will be accepted:

$$P[\text{bid accepted at } p_{bid}] = \frac{\sum_{p \leq p_{bid}} q_{bid}^{acc}(p) + \sum_{p \leq p_{bid}} q_{ask}(p)}{\sum_{p \leq p_{bid}} q_{bid}^{acc}(p) + \sum_{p \leq p_{bid}} q_{ask}(p) + \sum_{p \geq p_{bid}} q_{bid}^{rej}(p)}$$

$$P[\text{ask accepted at } p_{ask}] = \frac{\sum_{p \geq p_{ask}} q_{ask}^{acc}(p) + \sum_{p \geq p_{ask}} q_{bid}(p)}{\sum_{p \geq p_{ask}} q_{ask}^{acc}(p) + \sum_{p \geq p_{ask}} q_{bid}(p) + \sum_{p \leq p_{ask}} q_{ask}^{rej}(p)}$$

Having beliefs with respect to prices allows the agents to optimize utility, in accordance with the theory of monopolistic competition. Upon further reflection, it becomes clear that the GD-belief that a bid will be accepted at price p_{bid} is in fact a probability that is conditional on observing an ask. What matters, however, is the unconditional probability that the bid will be accepted *before it is rejected*. That is, the better belief function is given by expression (B.1.1).³⁰

²⁸Their distinction between accepted and rejected offers is the same as ours (c.f. the description of the ZIP-algorithm).

²⁹Tesauro and Das (2001) proposes modifications that cover trading multiple units and different market rules, e.g. an order book with a depth greater than one.

³⁰Gjerstad and Dickhaut (1998) suggests that traders remember a short history of, say, five offers, to which a spline function is fitted in order to arrive at the beliefs. In FACTS, however, agents remember the complete past as tabulated distributions. Using only a few observations makes price expectations more sensitive to the recent past. Since the subjects of Anderson et al. often submitted $n \times 1$ acceptances instead of $1 \times n$ this could trick robot traders into believing that the spread in prices is zero (and that could cause numerical instability). Using all observations, on the other hand, consolidates mistakes but also gives greater stability (i.e. consistency over time). Not having to fit splines saves time, but a reasonably granular belief function can also take quite long to evaluate, especially since there are multiple markets.

It is clear that the number of observations that determine a price expectation can and should be calibrated after robot traders generate enough transactions.

Table B.8 – GD: Awareness of the relative number of buyers

Description	aware		unaware	
	stable	ccw	stable	ccw
GDA	45.7	51.2	45.1	51.2
GDW	<u>51.8</u>	<u>51.7</u>	50.6	51.7

Percentages of correctly predicted human actions, excluding arbitrage, conditional on recognizing the move as a feasible opportunity. In the stable Scarf economy, buyers outnumber sellers (in the unstable economies there are as many buyers as there are sellers). Columns 'aware' refer to the situation that traders are aware of this fact; 'unaware' means that traders assume that there are equally many buyers and sellers. Percentages are an average over 1,000 runs. Although the differences are small, GDA and GDW-traders appear to be aware of the relative number of buyers. GDW-traders outperform GDA-traders.

Optimizing utility, based on beliefs with respect to the acceptability of offers, yields both preferred prices as well as preferred quantities. Here, however, we only use the preferred prices, because we are mainly interested in expectation formation. Furthermore, human traders tend to exchange small quantities, presumably as a safeguard against adverse effects of incorrect expectations. Conditional on these preferred prices as expected prices, we let GD-traders behave like other robot traders. Since a markup on expected prices would distort the optimization over prices, we set $\mu = 0$.

We consider a few variations on the theme of GD-beliefs. GDA-traders always derive admissible offers from their beliefs. They do so by restricting the domain of the belief function, conditional on floor prices (if any). GDW-traders on the other hand, consider the trade-off between proposing and waiting for better circumstances. They optimize over the full domain of their beliefs, and if their preferred price is not admissible then they wait. Another variation refers to awareness of the relative number of buyers compared to sellers. In deriving the probability that a bid is accepted before it is rejected, expression (B.1.1) assumes that the propensities of observing a bid or an ask are the same. In the stable Scarf economy, however, buyers outnumber sellers by a margin of 2 to 1. If traders are aware of this, and if they take it into account, they would use

$$\mathcal{B}'[\beta] = \frac{P_A [p^a \leq \beta]}{P_A [p^a \leq \beta] + 2P_B [p^b > \beta]}. \quad (\text{B.2.4})$$

in the stable Scarf economy, rather than (B.1.1). To the extent that monopolistic competition based on GD-beliefs captures the behavior of human traders, it seems that the subjects of Anderson et al. took the relative difference into account, c.f. table B.8. In principle, it is unfair to compare GDA- and GDW-traders based on correct predictions, because we can only verify predictions in case the human trader decided not to wait. However, conditioning on recognition can compensate for the disadvantage.

Given beliefs with respect to the acceptability of bids and asks, it is also possible to derive an estimate of what may be called a "no arbitrage" price, i.e. a price at which

Table B.9 – Determining the markup for eGD

markup	stable (%)	ccw (%)	markup	stable (%)	ccw (%)
$\mu = 0.00$	69.1	63.1	$\mu = 0.06$	56.0	60.6
$\mu = 0.01$	68.0	62.5	$\mu = 0.07$	54.1	59.5
$\mu = 0.02$	65.3	62.4	$\mu = 0.08$	53.2	59.2
$\mu = 0.03$	62.4	62.1	$\mu = 0.09$	51.5	59.0
$\mu = 0.04$	59.8	61.9	$\mu = 0.10$	50.1	58.6
$\mu = 0.05$	57.3	61.6			

Percentages of correctly predicted human actions, excluding arbitrage, conditional on recognizing the move as a feasible opportunity. Percentages are an average over 1,000 runs. The best values for the markup are $\mu = 0.00$ in both the stable and the counter-clockwise economy. However, we use $\mu = 0.01$ for avoiding perfectly correlated expectations.

acceptance of bids and asks is equally likely.³¹ Lower prices would invite arbitrage on the part of buyers, because if a seller were to accept a bid below the "no arbitrage" price then buyers have a better chance of re-selling the commodity at a slightly higher price. The "no arbitrage" prices are a good candidate for estimating the equilibrium price, because in an equilibrium every trader should be able to buy or sell at the same price.³² Suppose that GD-traders reflect on offers, by comparing what would be acceptable from a buyer's but also from a seller's point of view, then it is a small step to using "no arbitrage" prices as reservation prices. This is what eGD-traders do. Based on the percentage of correctly predicted human moves, eGD-traders ignore the fact that buyers outnumber sellers in the stable Scarf economy: 69% versus 60%.

Table B.9 presents the calibration of the markup for the unaware eGD-algorithm. Higher values of μ lead to poorer prediction rates, and more so in the stable economy. Although $\mu = 0.00$ is best, we use $\mu = 0.01$ in order to avoid perfectly correlated expectations.³³

³¹With hindsight, traders can avoid a complicated calculation of "no arbitrage" prices by directly estimating prices that they think have an even chance of being accepted and rejected. This would also introduce more heterogeneity in expectations.

³²"No arbitrage" prices do not require beliefs. It suffices for traders to estimate at which price an offer is equally likely to be improved upon or accepted. To get the most out of observed data, a proposed ask (bid) price should be interpreted as a willingness to accept bids (asks) at better prices. This is similar to the analysis in Gjerstad and Dickhaut (1998).

³³In the unstable economies, there is a discontinuity at $\mu = 0.00$ that manifests itself in the confidence intervals for concentration statistics (c.f. section 4.2.3 for an explanation of the confidence intervals). For instance, for the counter clockwise economy, goods 2 / 3 we have intervals 80-94 / 76-87 in case of $\mu = 0.00$; that compares to 3-7 / 14-83 if $\mu = 0.01$. In the clockwise economy, there is a similar discontinuity. For $\mu \geq 0.01$, on the other hand, the confidence intervals are very robust in both unstable economies. The same phenomenon also occurs in eME-trading, albeit less pronounced, c.f. footnote 34.

Table B.10 – ME: Awareness of the relative number of buyers

Description	aware		unaware	
	stable	ccw	stable	ccw
MEA	43.3	45.7	42.9	45.7
MEW	<u>48.8</u>	<u>48.5</u>	45.1	48.5

Percentages of correctly predicted human actions, excluding arbitrage, conditional on recognizing the move as a feasible opportunity. Percentages are an average over 1,000 runs. MEA- and MEW-traders also appear to take into account that buyers outnumber sellers. In the unstable economies, the number of buyers and sellers is the same. MEW-traders outperform MEA-traders.

B.2.7 MaxEnt beliefs, with variations

Let traders believe $m_j \leq p_j \leq M_j$ and consider all possible distributions of p_j over the support $[m_j, M_j]$ that have an expected value that is equal to the observed mean of p_j . Within this class, the truncated exponential distribution has maximal entropy (c.f. Conrad (2004)).

$$f(p_j) = \begin{cases} \frac{\theta e^{\theta p_j}}{e^{\theta M_j} - e^{\theta m_j}}, & \theta \neq 0; \\ \frac{1}{M_j - m_j}, & \theta = 0. \end{cases}$$

This means that this distribution does not assume information that is not embodied in the data. To assume any other distribution for commodity prices implies that extra information is used. A justification for that assumption cannot be found in the observations. Distributions that maximize entropy relative to constraints are called MaxEnt distributions, c.f. Theil and Fiebig (1984). The parameter θ can be solved by setting

$$\frac{1}{n(T)} \sum_{t \leq T} p_j^t = \frac{M_j e^{\theta M_j} - m_j e^{\theta m_j}}{e^{\theta M_j} - e^{\theta m_j}} - \frac{1}{\theta}$$

with $n(T)$ the number of observations until time T . Traders can tally the means of observed prices per commodity / type of offer. In contrast to GD-beliefs, here they ignore quantities and they just focus on offered prices. This allows us to see if there is an advantage to observing quantities. A buyer contemplating a bid price p can estimate its acceptability by the probability of observing an ask price lower or equal to p before the bid is improved upon.

If available, these distributions can be used for different purposes, just like the GD-beliefs: for having traders select the best admissible price (MEA-traders), or for letting them consider the trade-off between waiting and proposing (MEW-traders), or for estimating the implied "no arbitrage" prices (eME-traders).

The calibration of awareness of the relative number of buyers shows that MEA and MEW-traders take the difference between the number of buyers and sellers into account (c.f. table B.10), eME-traders on the other hand do not: the prediction rate of an unaware traders is 64% and the rate of aware traders is 32%.

As is the case with GD-beliefs, for MEW traders a markup would distort the optimization over prices. For eME-traders, we do calibrate the markup: table B.11 shows that the best values for the markup are $\mu = 0.00$ for both the stable economy

Table B.11 – Determining the markup for eME

markup	stable(%)	ccw (%)	markup	stable(%)	ccw (%)
$\mu = 0.00$	<u>64.1</u>	<u>61.7</u>	$\mu = 0.06$	58.8	59.3
$\mu = 0.01$	63.2	61.3	$\mu = 0.07$	58.1	58.9
$\mu = 0.02$	62.3	60.9	$\mu = 0.08$	57.4	58.5
$\mu = 0.03$	61.4	60.4	$\mu = 0.09$	56.7	58.2
$\mu = 0.04$	60.5	60.0	$\mu = 0.10$	56.0	57.8
$\mu = 0.05$	57.9	59.6			

Percentages of correctly predicted human actions, excluding arbitrage, conditional on recognizing the move as a feasible opportunity. Percentages are an average over 1,000 runs. The best value for the markup are $\mu = 0.00$ for both the convergent and the counter-clockwise treatment. However, we use $\mu = 0.01$ for avoiding perfectly correlated expectations.

and for the counter-clockwise unstable economy.³⁴ We use $\mu = 0.01$, however, for preventing perfectly correlated price expectations. Otherwise, there are no specific parameters to be fixed.

B.2.8 Target utility

TU-traders try to achieve a target utility level, say \hat{u} . As a result of this, they perceive markets as interconnected. For instance, if a trader of type *III* has to pay more for acquiring commodity 2 then he wants to pay less for getting commodity 3. Similarly, if a trader of type *I* receives a better price when selling commodity 2, he can offer to pay more for commodity 3. We propose different flavors of TU-algorithms depending on how on the utility target is set and on how reservation prices are derived from the target.³⁵

TU-traders distinguish between plausible and implausible prices, with the intention to restrict trading to plausible prices. A maximum level of utility occurs if prices in both markets are plausible and optimal. If the target is fixed at such an optimum then the trader has to reject almost all plausible prices as being not good enough. On

³⁴In the unstable economies, there is a discontinuity at $\mu = 0.00$ that manifests itself in the confidence intervals for concentration statistics (c.f. section 4.2.3 for an explanation of the confidence intervals). For instance, for the counter clockwise economy, goods 2 / 3 we have intervals 39-52 / 41-86 in case of $\mu = 0.00$; that compares to 35-47 / 28-51 if $\mu = 0.01$. In the clockwise economy, there is also a discontinuity. For $\mu \geq 0.01$, on the other hand, in both unstable economies the confidence intervals are very robust. The same phenomenon also occurs in eGD-trading where it is more pronounced, c.f. footnote 33.

³⁵An interesting variation on utility targets is that traders use their individual demand curves (based on initial endowments) to determine how much they want to spent on each commodity, given expected prices. Traders can then calculate average prices for subsequent transactions that would bring planned total expenditure to the desired level. Friedman (2007) appeals to this kind of behavior to clarify the difference between a normal, theoretical demand curve and what is described as the "usual" demand curve. The latter corrects for deviations from currently expected prices in previous transactions. Friedman (2007) assumes that previous prices are less favorable; implicitly, this is tantamount to positing a Marshallian path, c.f. 4.3.2.3.

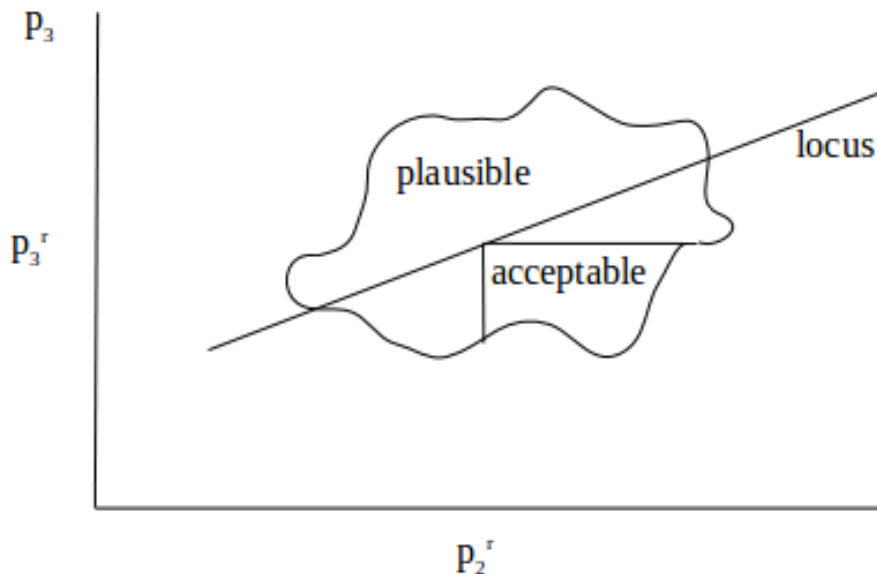


Figure B.3 – The locus consists of prices which all yield a particular target for a trader of type I; the locus cuts through the set of plausible prices. For a trader of type I, all prices below the locus would yield a utility level which exceeds the target. The prices, that are acceptable given unconditional reservation prices, p^r , lie below / to the right of p^r . Clearly, the set of plausible, acceptable prices, given p^r , is typically smaller than the set of plausible prices which lie below the locus.

the other hand, if a trader is satisfied with the minimum level of utility then all plausible prices in principle will be acceptable to him. This trade-off between the utility target and the number of prices being acceptable is a trade-off between the utility target and its feasibility (i.e. the probability that the target will be achieved), given that the agent wants to trade at plausible prices. We distinguish different attitudes towards target setting: traders can optimize expected utility or they can maximize the target conditional on a probability that the target will be achieved.

Suppose that a utility target has been fixed, the trader then knows the locus of prices that yield the same utility level (assuming that he can trade any quantity at these prices). Selecting a vector of reservation prices means choosing a point on the locus. This point defines a set of acceptable prices that are compatible with the reservation prices, c.f. figure B.3. If a trader applies the reservation prices unconditionally (i.e. simultaneously), then effectively he is seeking profits in all markets. The probability of achieving a utility target with specific reservation prices depends on the size of the set of acceptable prices relative to the set of plausible prices. The trader can select a point on the locus by maximizing the probability that the target will be achieved. Once we have a relation between the utility target and optimal reservation prices we can go further and vary the target so as to maximize the expected utility, $\hat{u} \times P[u \geq \hat{u}]$.

Another approach would be to select a plausible price in one market and use the locus to map it to a reservation price in another market. This way the trader can condition a decision on the situation at hand. Suppose a trader wants to act in the market for commodity 2. In this case, he derives a reservation price from the locus, conditional on the currently expected price of commodity 3. This expected price may depend on floor offers or, if these are not available, on the latest synchronized prices. These kind of reservation prices will be called conditional reservation prices; they permit traders to make use of available information.

To appreciate the difference in terms of behavior, consider the latitude a trader may allow himself during trading. A trader with unconditional reservation prices will reject a counter offer if it compares unfavorably with the reservation price. Similarly, he will choose to wait if a rival offer is better than his own unconditional reservation price. Traders with conditional reservation prices may still be able to accept or to propose an offer instead of having to wait, because their reservation prices are more flexible.

Let prices be plausible if they fall in a particular window, i.e. $p_j \in [p_j^d, p_j^u]$.³⁶ The position of the plausibility window (i.e. its center) depends on expected prices. The size of the window depends on the spread of recently observed prices relative to a trader's *own* expected price and on a parameter, δ_i , that optionally can be randomized.

$$p_j^u - p_j^d = \delta_i \sigma_{ji}.$$

Given a notion of plausibility, offers by other agents will sometimes be perceived as favorable surprises, e.g. if an ask is below the lowest price thought to be plausible. If a trader observes a favorable surprise then he updates his price expectations and shifts the location of the plausibility window so as to make the surprise (just) plausible. For instance, an unexpected high bid for commodity 2, say p_2 , may lead a trader of type I to shift his plausibility window such that $p_2^d = p_2$. That is, he now expects that other buyers may be willing to bid even more.

Let \mathbf{p}^r be the vector of reservation prices and let \hat{u} be the utility target. The target restricts reservation prices via the budget constraint. Consider the case of traders of type I ; with a target equal to \hat{u} , in an optimum they consume $\hat{x}_1 = 400\hat{u}$ and $\hat{x}_2 = 20\hat{u}$; then their budget constraint is

$$(400 + 20p_3^r) \hat{u} \leq w_1 + p_2^r w_2 + p_3^r w_3 \tag{B.2.5}$$

If the weak inequality is an equality then (B.2.5) defines a locus of reservation prices that are sufficient for achieving the target. The position of the locus depends on the utility target. If $20\hat{u} > w_3$ then the locus slopes upward; if $20\hat{u} < w_3$ then the slope of the locus is downward. If $w_2 = 0$ or $20\hat{u} = w_3$ then the locus is vertical or horizontal respectively. If traders sell or buy in both markets then reservation prices will be inversely related. If they buy in one market while selling in the other then reservation prices will be perceived as positively related. If the optimal quantity of one

³⁶We did some work on an algorithm that used prediction ellipses instead of a rectangular plausibility window. However, its computational requirements made it too slow. With hindsight, the distinction between plausible and implausible prices (and opportunities) is rather coarse and applying more complexity to making that distinction does not pay off. Probably it would be better to have simpler rules for assessing plausibility.

commodity has been obtained, or if traders do not have anything left of the commodity which they do not prefer then reservation prices will be treated as unrelated.

Table B.12 gives an example of the *ex ante* probabilities that traders of type *I* will achieve their utility target in case of conditional reservation prices and equi-probable plausible prices.

Adaptation to current market conditions can be taken one step further, if traders allow both reservation prices to adjust. At each time the trader has to take a decision, he selects (notional) reservation prices so as to be competitive, either in both markets or in a preferred market. This may result in meaningful behavior because traders can use the flexible reservation prices to determine whether their target utility level is consistent with rival offers. If both rival offers exist, and lie above the locus, then the target is currently not feasible. This can be a trigger for lowering it. Table B.13 gives a global description of how to maintain a "sticky" target.

In principle, TU-traders respond to rival offers which exceed their own reservation prices and to favorable counter offers. In case of a rival offer, reservation prices are updated so as to admit an improving offer. If the utility target allows a trader to be competitive in both markets then he will choose his reservation prices accordingly (for having more freedom of future action). Otherwise, traders will compete in a preferred market, forsaking opportunities to trade in the other market. This preference for markets depends on (in order of decreasing importance): (i) the availability of a plausible reservation price; (ii) a direct contribution to achieving a higher utility level, or (iii) increasing future freedom of action by obtaining more money, or (iv) an equi-probable choice between both markets. TU-traders wait if rival offers in both markets are deemed implausibly unfavorable.

The TU-simulations vary across three dimensions: (i) how the target is set and maintained (optimized, fixed feasibility or a sticky target); (ii) how the fixed feasibility is set (at random or at a probability of 0.5) and (iii) the type of reservation prices (unconditional, conditional or sticky). Table B.14 presents the results. As expected conditional reservation prices, which allow traders to take the current situation into account, perform better than unconditional reservation (with which traders seek profits in all markets).

We determine the markup for fixed feasibility with a target that has a probability of 0.5 to be achieved with conditional reservation prices, c.f. table B.15.

B.3 Conclusions

In this appendix we have introduced some technical elements of FACTS; in particular how opportunities are perceived and represented as lotteries. We have offered heuristic arguments for a set of rules of thumb for prioritizing feasible actions. In addition, this appendix describes and pre-calibrates algorithms for maintaining reservation prices.

Table B.12 – Conditional probabilities of achieving target \hat{u}

Bounds on \hat{u}	$P[u \geq \hat{u}]$
$\frac{w_1 + p_3^d w_2 + p_3^n w_3}{400 + 20p_3^n} \leq \hat{u} \leq \min \left(\frac{w_1 + p_2^d w_2 + p_3^d w_3}{400 + 20p_3^d}, \frac{w_1 + p_2^n w_2 + p_3^n w_3}{400 + 20p_3^n} \right)$	$1 - \frac{((400 + 20p_3^n)\hat{u} - (w_1 + p_2^d w_2 + p_3^n w_3))^2}{2w_2(20\hat{u} - w_3)(p_2^n - p_2^d)(p_3^n - p_3^d)}$
$\frac{w_1 + p_2^n w_2 + p_3^n w_3}{400 + 20p_3^n} \leq \hat{u} \leq \frac{w_1 + p_2^d w_2 + p_3^d w_3}{400 + 20p_3^d}$	$\frac{w_1 + \frac{1}{2}(p_2^d + p_2^n)w_2 + p_3^d w_3 - (400 + 20p_3^d)\hat{u}}{(20\hat{u} - w_3)(p_3^n - p_3^d)}$
$\frac{w_1 + p_2^d w_2 + p_3^d w_3}{400 + 20p_3^d} \leq \hat{u} \leq \frac{w_1 + p_2^n w_2 + p_3^n w_3}{400 + 20p_3^n}$	$\frac{w_1 + p_2^n w_2 + \frac{1}{2}(p_3^d + p_3^n)w_3 - (400 + 20(p_3^d + p_3^n))/2\hat{u}}{w_2(p_2^n - p_2^d)}$
$\max \left(\frac{w_1 + p_2^d w_2 + p_3^d w_3}{400 + 20p_3^d}, \frac{w_1 + p_2^n w_2 + p_3^n w_3}{400 + 20p_3^n} \right) \leq \hat{u} \leq \frac{w_1 + p_2^d w_2 + p_3^d w_3}{400 + 20p_3^d}$	$\frac{((400 + 20p_3^d)\hat{u} - (w_1 + p_2^n w_2 + p_3^n w_3))^2}{2w_2(20\hat{u} - w_3)(p_2^n - p_2^d)(p_3^n - p_3^d)}$

The probabilities apply to traders of type I, who sell commodity 2 and buy commodity 3 (these traders face an upward sloping locus). There are similar probabilities for the cases that the locus is downward sloping, vertical or horizontal, as well as for traders of types II and III.

Table B.13 – Maintaining a sticky target

the locus is elastic with respect to the observed offer and the offer is submitted by another trader	rival offer	no longer competitive (observed offer is better than own reservation price)	offer in preferred market	plausible offer	improve upon the observed offer by mixing it with the least best of the counter floor offer (if any) and the appropriate end point of the locus; if the observed offer is a (partial) acceptance set the reservation price equal to it
			offer not in preferred market	implausible offer	select a point in the locus which is closest to the observed rival offer
otherwise	counter offer	still competitive better than expected not better than expected	ignore rival offer	implausible offer	improve upon the observed offer by mixing it with the least best of the counter floor offer (if any), the constraint induced by the rival floor offer in the preferred market (if any) and the appropriate end point of the locus; if these constraints are incompatible then ignore the observed rival offer
			ignore rival offer	ignore rival offer	ignore rival offer
			set the reservation price equal to the least best of the observed counter offer and the appropriate endpoint of the locus	ignore counter offer	
	ignore the observed offer				

Rules for maintaining a sticky target, i.e. a target which is decreased only when rival floor offers are inconsistent with the target. For traders of type I this happens when the rival floor offers lie above / to the left of the expected prices.

Table B.14 – Target calibration

Description	stable (%)	ccw (%)
Optimized, unconditional	47.8	50.2
Optimized, conditional	61.1	61.4
Fixed feasibility: randomized target, unconditional	48.3	50.7
Fixed feasibility: randomized target, conditional	<u>65.4</u>	62.4
Fixed feasibility: 0.5, unconditional	48.3	50.8
Fixed feasibility: 0.5, conditional	65.3	<u>62.5</u>
Sticky target, notional	63.0	59.7

Percentages of correctly predicted human actions, excluding arbitrage, conditional on recognizing the move as a feasible opportunity. Percentages are averages over 1,000 runs. Traders appear to pursue a target that has fixed feasibility. Whether the probability of meeting the target is randomized or set to a probability of 0.5 is not clear. Henceforth, we'll use the randomized version because it is more generic. Conditional reservation prices outperform both unconditional and notional reservation prices. Interestingly, sticky targets do better than optimized targets.

Table B.15 – Determining the markup for TU-trading

markup	stable(%)	ccw (%)	markup	stable(%)	ccw (%)
$\mu = 0.00$	<u>65.4</u>	<u>62.4</u>	$\mu = 0.06$	60.2	61.2
$\mu = 0.01$	65.1	62.1	$\mu = 0.07$	59.4	60.2
$\mu = 0.02$	63.9	62.2	$\mu = 0.08$	59.1	60.1
$\mu = 0.03$	63.0	62.0	$\mu = 0.09$	58.3	60.0
$\mu = 0.04$	61.8	62.0	$\mu = 0.10$	57.6	59.7
$\mu = 0.05$	60.8	62.0			

Percentages of correctly predicted human actions, excluding arbitrage, conditional on recognizing the move as a feasible opportunity. Percentages are an average over 1,000 runs. The best values are $\mu = 0.00$ in the both the stable and the ccw treatment.