Price discovery with fallible choice

Ruiter, A.G.J.M.

Chapter 6

Robustness

This chapter investigates the robustness of our previous simulation results. As an out-of-sample test, eGD-expectations are applied to trading in the so-called unstable Gale economy (c.f. section 6.1). Other results are subjected to evolutionary competition: traders exchange information with respect to the success of strategies and select the one that suits them best (c.f. section 6.2). Section 6.3 considers the impact of heterogeneity on price formation and end of period allocations; section 6.4 concludes.

6.1 The unstable Gale economy

In the calibration, the data of Anderson et al. (2004) have been extensively used for learning about human trading behavior and for tuning the algorithms. We like to assess the performance of the eGD-algorithm in another, independent, context. Crockett et al. (2011) reports on price formation in an experiment in which human traders re-enact the (unstable) Gale economy.

In the Gale examples, there are two commodities, the first of which is money and it has a price equal to 1.[1] There are two types of traders, both having a utility function of the form

$$u^i = \min(\alpha_1^i x_1 + \beta_1^i; \alpha_2^i x_2^i + \beta_2^i).$$

Table 6.1 presents the parameters for the unstable Gale example. The parameters of the stable economy are the same, except for the endowments which are switched between the traders of type I and II. For the experiments of Crockett et al. they imply the following (approximate) equilibrium prices: $p_2^* = 158$ in the unstable and $p_2^* = 1714$ in the stable economy respectively.

Gale's examples demonstrate the possibility of tâtonnement being unstable almost everywhere. In particular, in the unstable example, if the initial price is above its equilibrium value then subsequent prices will increase beyond limits in the unstable example. On the other hand, if the initial price is below the equilibrium, then subsequent prices will decrease sharply, approaching zero. That is, depending on how price formation starts, either buyers or sellers will be prepared to practically give away their endowments.

[1] In Crockett et al. (2011) the second commodity, i.c. y, is money.

Table 6.1 – Parameters in the unstable Gale economy

Item	Gale	Crockett *et al*
$\boldsymbol{\alpha}^1$	$\left(1, \frac{1}{2}\right)$	$\left(1, \frac{1}{0.028167}\right)$
$\boldsymbol{\beta}^1$	$(0, 0)$	$\left(\frac{38}{0.028167}, 0\right)$
$\mathbf{w}^1$	$(0, 1)$	$(400, 15)$
$\boldsymbol{\alpha}^2$	$(1, 2)$	$\left(1, \frac{1}{0.00152}\right)$
$\boldsymbol{\beta}^2$	$(0, 0)$	$\left(\frac{6}{0.00152}, 0\right)$
$\mathbf{w}^2$	$(1, 0)$	$(5600, 5)$

Crockett et al. have changed the parameters of the original examples to make it more difficult for human traders to guess the Walrasian equilibrium prices.

In the unstable Gale economy, Crockett et al. (2011) finds that (i) human traders learn the correct price if initial prices are sufficiently close to the equilibrium value; furthermore (ii) prices increase (decrease) if the initial price is well above (below) 158. Neither result is found in our simulations.

Figure 6.1 shows eGD price formation in the unstable Gale economy. Typically prices exhibit a downward trend and even more so if the average of the first 10 observations is above 161; then 96% of the runs (instead of 92%) show a downward trend. The slow rate of decrease reflects the fact that new observations successively receive less weight and that the algorithm in the long run is likely to settle on a stable state. Therefore the Gale simulations confirm what was already apparent from the Scarf simulations: eGD-expectations over time become inflexible (this feature is robust).[2] Here it even leads to convergence. Human price expectations are less well and differently anchored.

6.2 Ecological rationality

A strategy will be called ecologically rational if it survives a competition with other strategies in an environment where traders can learn and can adopt the strategy that they deem most successful.[3] Ecological rationality depends on the environment, because there can be a symbiosis between different strategies. This was clearly demonstrated in the so-called Santa Fe tournament (c.f. Friedman and Rust (1993)). This tournament was won by the famous sniper algorithm of Kaplan. His algorithm waits in the background until other, active, traders are close to an agreement; then it tries to steal the deal. The sniper algorithm is a parasite that needs other, active algorithms; in an economy populated with snipers nothing happens.

[2]The degree of inflexibility or inelasticity can be reduced by having fewer observations determine the underlying beliefs. Before this variable is calibrated, however, robot traders should first generate more transactions. Just as in the Scarf examples, human traders generate more transactions in the Gale example as well.

[3]Smith refers to this notion as ecological rationality, c.f. Smith (2008).

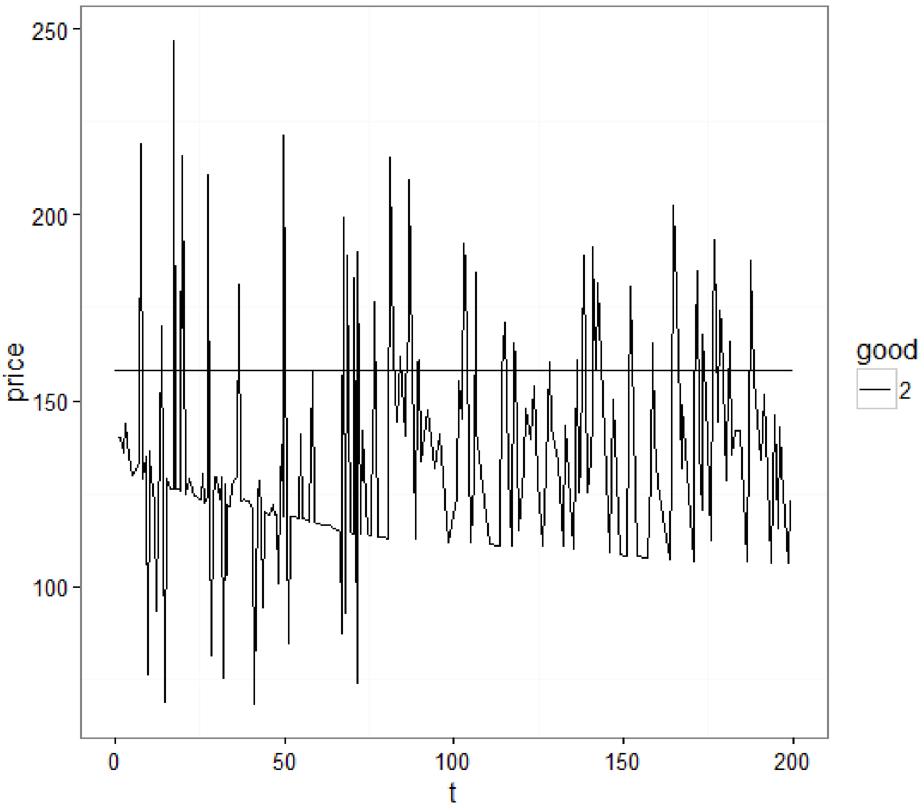

Figure 6.1 – *Representative price formation in the simulated unstable Gale economy. The horizontal line indicates the Walrasian equilibrium price. Even though trading prices can spike well above the equilibrium price, the trend typically is downwards. The rate of decrease is slow; in each simulation the trend remains above 60.*

Below, we detail how traders learn about strategies, which is a combination of replicator dynamics and reinforcement learning (section 6.2.1) and we revisit several topics of the earlier calibrations. Sections 6.2.2 and 6.2.3 discuss learning the markup and an attitude towards a utility target respectively. Other subjects are monopolistic competition (section 6.2.4), expectation formation (section 6.2.5), entropy-sensitive preferences (section 6.2.6), rules for selecting a best alternative (section 6.2.7) and arbitrage (section 6.2.9).

6.2.1 Learning

In FACTS evolutionary learning is implemented as follows. At the end of each run, traders receive a message that triggers them to share their experience with another, random trader.[4] The recipient returns the favor by revealing his experience. In

[4] These messages are loaded into the message queue before the run starts. By constraining the recipients one can implement a network with a pre-defined structure. Currently, the messages are

particular, traders communicate what they consider to be the best strategy, and they reveal the average level of utility they have obtained by applying that strategy.

Traders have memory; they remember the average utility (and its spread) they have experienced for each of the strategies they have applied so far. As a result of this, strategies can re-emerge, even if they are not recommended. Suppose that a trader has tried two alternative methods and that one has produced a better result on average. If he proceeds with the best alternative then it may become clear that the high score was exceptional and that its new average utility falls below the average of the other strategy.

Traders are willing to experiment: if they become aware of a new strategy they are willing to try it simply because it comes recommended. After having tried a particular strategy at least five times, traders apply additional constraints before adopting it again in case they prefer another strategy. Suppose that another trader recommends strategy r, while a trader himself prefers strategy p. Let μ_r and σ_r be his own estimates of expected utility and the standard deviation of experienced utility after applying strategy r. Then the trader verifies (i) if the revealed average utility, $\tilde{\mu}_r$, exceeds his own estimate and also (ii) if strategy r can improve upon his currently best average utility:

$$\begin{aligned} \mu_r &< \tilde{\mu}_r \\ \mu_p &< \mu_r + 2\sigma_r. \end{aligned}$$

That is, traders acknowledge that their own estimate of the average utility may be too low. If μ_p exceeds $\mu_r + 2\sigma_r$ then adopting strategy r is deemed not credible, even though other traders recommend it.[5] Since traders base their decisions on their own experiences, FACTS can accommodate heterogeneity between types of traders, because convergence of evolutionary competition does not require unanimity with respect to the value of alternative strategies.

Initial simulations demonstrated that convergence can be slow. Therefore, the simulation terminates if not one trader adopts a different strategy, or after 2,000 runs, whichever comes first. In order to avoid premature terminations, we require that each simulation completes at least 750 runs.[6]

---

dated per the time horizon, but it would also be possible to have information exchange between two traders at random times.

[5] Observe that traders do not assess the credibility of the message they receive; i.e. they do not verify $\mu_r < \tilde{\mu}_r < \mu_r + 2\sigma_r$. Instead they assume that other traders tell the truth, and they believe that strategy r is good for the messenger even though they themselves prefer p to r. If $\mu_r < \tilde{\mu}_r < \mu_p < \mu_r + 2\sigma_r$ then the trader is still willing to experiment, in search of clarity in the form of $\mu_r + 2\sigma_r < \mu_p$.

[6] In previous chapters, each run constituted a simulation experiment; here, all the runs constitute a single experiment. The time required for a sufficient number of replications of the learning experiments, unfortunately, is prohibitive. Simulations that quickly yield strong dominance, preferably in line with previous calibrations, have some claim to robustness, however, because the robot traders are not fooled by utility scores that qualify as outliers. They are willing to experiment unless their preferred strategy does better by at least two standard deviations.

Table 6.2 – Learning the markup

treatment	0.01	0.02	0.03	0.04	0.05	length
stable				2	13	868
ccw			1	1	13	1303
cw				1	14	1033

Final distributions of traders over different values of the markup and the length of the simulation by treatment. The initial distribution is the same for all three examples; care has been taken to vary behavior across trader types. In each of the three treatments, traders prefer high values of the markup. The simulation is based on eGD-expectations and rules of thumb for prioritizing feasible actions.

6.2.2 Learning the markup

Best values for the markup (if applicable) have been determined as part of the pre-calibration, c.f. appendix B. In principle, there are three approaches that can be applied, based on optimizing (i) the similarity between aggregate results; (ii) the number of transactions and (iii) the prediction rate. The markup is quite sensitive to the calibration method: method (iii) leads to low values, method (i) to intermediate values and method (ii) to (implausibly) high values. We have decided to use prediction rates for the calibration, because similarity at the aggregate level (small average distance between robot and human prices and comparable degree of concentration of trading prices) is indirect, less transparent and to a certain degree arbitrary (because there are multiple criteria that are difficult to rank and weigh). Based on prediction rates, the optimal value of μ is zero (c.f. table B.9); however, in order to prevent perfectly correlated eGD-expectations, the markup is set to $\mu = 0.01$.

Given the option to change between strategies, traders opt for higher values of the markup, c.f. table 6.2. This is as expected, because robot traders do not generate enough transactions. Since traders typically buy what they need and sell what they can spare a higher number of transactions tends to increase levels of utility and that means a comparative advantage of higher values of the markup. If one would allow values $\mu > 0.05$ then these would have been preferred. A proper calibration of μ therefore requires that robot traders generate a sufficient number of transactions.

6.2.3 Attitudes toward a utility target

In the TU-algorithm, traders try to achieve a target level with respect to utility. As a result of this, they perceive markets as interdependent. There are different flavors of TU-algorithms depending on how traders set a target and how they derive reservation prices from it (c.f. appendix B for a detailed explanation and for the pre-calibration).

In the pre-calibration, it was found that traders do not use unconditional reservation prices. That is, they do not apply their reservation prices simultaneously (that would amount to seeking profits in each individual market). Instead they condition their action in "one" market on (floor) prices in the "other" market. Table 6.3 contains results of learning if traders get to choose between sticky targets, or one of the other algorithms with conditional reservation prices. The results are more pronounced, treatment-specific and also different from the pre-calibration (i.e. table B.14).

Fixed feasibility with randomized targets does best in predicting the moves of human traders, but apparently it is not ecologically rational. Interestingly, learning leads to a symbiosis of sticky and optimized targets in the stable and counter clockwise treatments. This is also unexpected since the prediction rate of sticky targets in the counter clockwise example is less than those of the other algorithms.

While it is plausible that different environments may favor different strategies, it is not yet clear why the ccw and the cw treatment should lead to different results, both in terms of the dominant algorithm and the speed of convergence. Sticky targets consist of a number of behavioral rules, c.f. B.13. It would be an interesting experiment to subject this collection of rules to learning.

6.2.4 Monopolistic competition

The theory of monopolistic competition suggests that traders set prices by maximizing expected utility (or profit) against their own beliefs. In chapter 4 we did not find any evidence to support this claim. The Scarf examples induce behavior that is very sensitive to the belief that a proposal will be accepted. Assuming that trading behavior, and information processing in particular, fundamentally do not depend on the preferences of the Scarf examples, we propose that monopolistic competition is an unlikely explanation of human trading at all prices in a CDA. That does not, however, preclude that monopolistic competition can be ecologically rational.

The GDW-algorithm sets prices based on optimizing utility against unconditional Gjerstad-Dickhaut beliefs (i.e. GDW-traders are prepared to wait, c.f. the pre-calibration in section B.2.6). We consider two tests of monopolistic competition, one with rules of thumb and another with selection of a best opportunity based on expected utility maximization. The results of these two tests are similar; the only difference is that convergence in the clockwise economy is much slower in case of expected utility maximization, c.f. table 6.4. Price taking strongly dominates price setting in the stable and counter clockwise treatments. In the clockwise treatment there is a symbiosis between type III traders acting as price takers and the other traders acting as price setters. It is unclear why the results of the clockwise and counter clockwise are different. In both cases the type I and II operate in one market only, and in both cases type III traders have to sell one commodity and buy another with the proceeds, c.f. section 3.2.2.

Table 6.3 – Attitudes toward a utility target

treatment	stk	opt	fxf-rnd	fxf-avg	length
stable	5	10			936
ccw	11	4			1853
cw		14		1	753

Final distributions of traders over different attitudes towards target setting. The initial distribution is the same for all treatments; care has been taken to vary behavior across trader types. Sticky targets (column "stk") have notional reservation prices, while optimized (column "opt") and fixed feasibility (columns "fxf-") targets use conditional reservation prices (see section B.2.8 for an explanation of these concepts).*

Table 6.4 – Learning monopolistic competition

treatment	RoTh			EU		
	taking	setting	length	taking	setting	length
stable	15	0	750	15	0	750
ccw	15	0	750	15	0	750
cw	5	10	819	5	10	1689

Distributions of traders over different behaviors: (i) price-taking (columns "taking", based on the eGD-algorithm) and price-setting (columns "setting", based on the GDW algorithm) combined with a selection of a best alternative based on rules of thumb (columns "RoTh") and expected utility maximization (columns "EU"). The initial distribution is the same for all three treatments; care has been taken to vary behavior across trader types. The maximum number of runs is 2,000. In the cw-treatment there is a symbiosis between price setters (type I and II traders) and takers (type III traders). Here, convergence is much slower if the selection of a best alternative is based on EU.

6.2.5 Expectation formation

The calibration of expectation formation resulted in the eGD-algorithm being the best, especially in the stable Scarf economy (c.f. chapter 4). Therefore, we let it compete with six alternative methods that also performed relatively well. Table 6.5 largely vindicates the earlier conclusion that eGD is a strong algorithm for expectation formation; except for the fact that, in the stable treatment, it is strongly dominated by the ZIP-algorithm. This result is replicated in a three-way runoff between the eGD-, ZIP- and eME-algorithm, c.f. table 6.6.

ZIP traders manage their competitiveness: if their reservation price is no longer competitive they update it, otherwise they keep it unchanged. ZIP-traders therefore will be inclined to generate more transactions, making it ecologically rational in the stable Scarf economy. This comparative advantage will be smaller after robot traders generate a sufficient number of transactions.

The runoff between eGD, ZIP and eME produces a very low efficiency. This becomes clear from figure 6.2 and is confirmed in table 6.11.

6.2.6 Entropy-sensitive preferences

Adding an entropy term to a von Neumann-Morgenstern utility function introduces a trade-off between utility and uncertainty. The trade-off can express uncertainty / complexity aversion (negative sensitivity), but also a love of gambling (positive sensitivity). Although it is not clear whether people can easily switch between different sensitivities toward uncertainty, we proceed as if they can.

Trading in the unstable treatments apparently favors gambling over uncertainty aversion. Interestingly, in the clockwise treatment entropy neutrality is the best strategy. It is not clear why the unstable economies would be different in this respect. Table 6.7 shows that convergence is slow. After 2,000 runs traders still alternate between different "strategies", and figure 6.3 shows that the frequency of these changes on average is fairly stable. We conclude that neutrality (i.e. standard expected utility

Table 6.5 – Pairwise competition with eGD

expectation	final frequency			length		
	stable	ccw	cw	stable	ccw	cw
eBAS	0	5	4	750	864	1,000
eEMA	0	2	0	750	1,000	750
ZIP	15	5	9	750	753	1,000
AA	0	5	4	750	765	1,000
eME	0	10	10	750	792	752
TU	0	0	5	750	782	891

Final frequencies of traders adopting a method of expectation formation in a pairwise competition with eGD, plus the lengths of the simulations by method and treatment. Here, the maximum length is 1,000 runs. The initial distribution is the same for all three examples; care has been taken to vary behavior across trader types. At the end of the simulation in the stable treatment (after the minimum length of 750 runs), the eBAS-algorithm is adopted by none of the traders. This implies that all traders have adopted eGD, because of pairwise matching. The eGD-algorithm is selected as a benchmark, because it came out best in the calibration in chapter 4. The table demonstrates that the eGD-algorithm indeed is strong in the stable Scarf economy. The fact that here it is strongly dominated by the ZIP-algorithm may be due to robot traders generating not enough transactions. In the ccw-treatment there appears to be a symbiosis of the eME- and eGD-algorithm, with the former being more popular than the latter.

Table 6.6 – Learning expectation formation

treatment	eGD	ZIP	eME	length
stable	0	15	0	750
ccw	10	5	0	889
cw	7	8	0	2,000

Final distributions of traders over methods of expectation formation and the length of the simulations. The maximum length is 2,000 runs. The initial distribution is the same for all treatments; care has been taken to vary behavior across trader types. The results of the runoff agree with pairwise matching: ZIP strongly dominates eME, reducing the runoff to a confrontation of the eGD- and ZIP-algorithm.

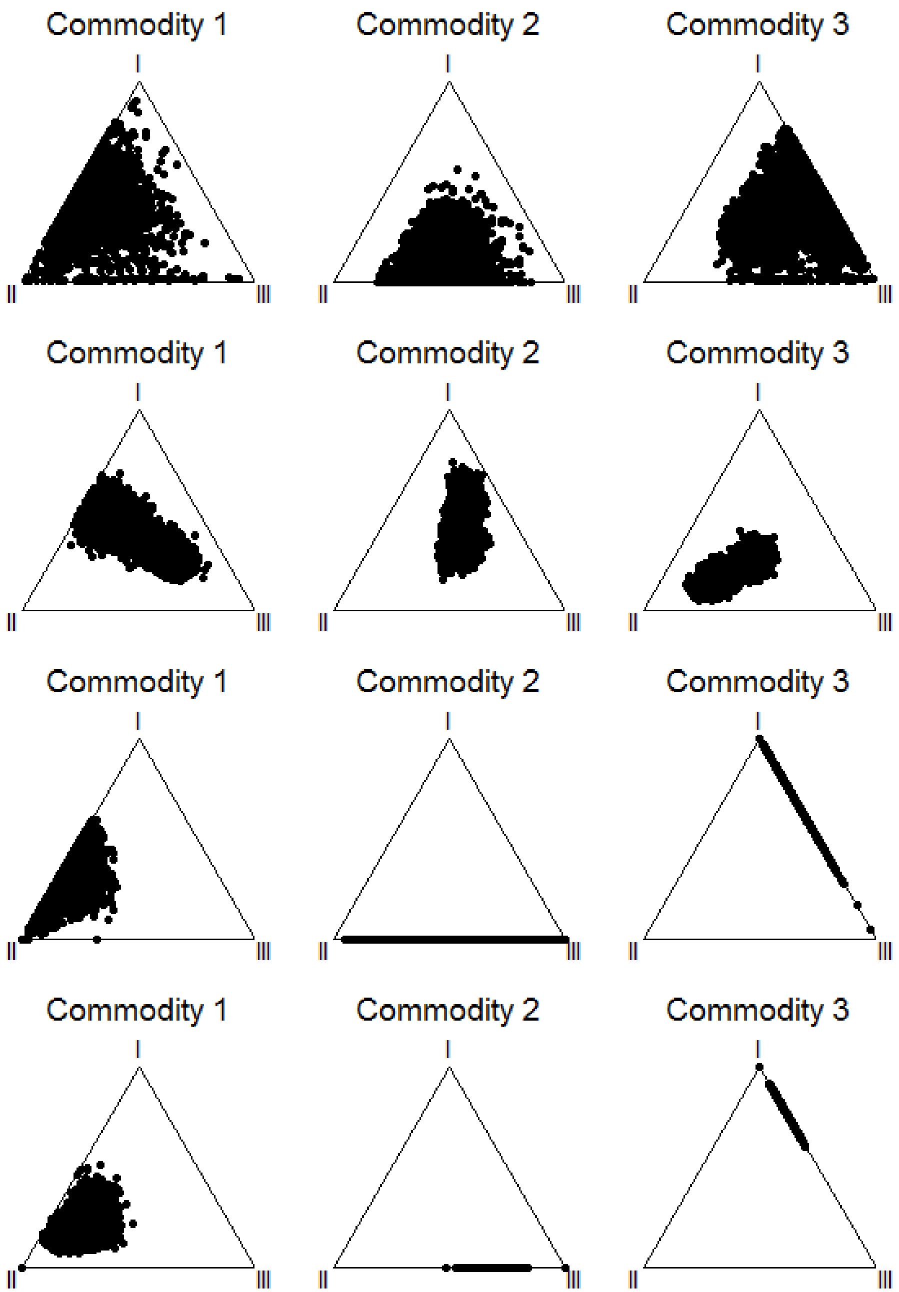

Figure 6.2 – *Shares per trader type in the available commodities in the runoff between eGD, ZIP and eME stable (row 1), eGD stable (row 2), eGD, ZIP and eME ccw (row 3) and eGD ccw (row 4). Each dot represents an allocation at the end of a period. The results of 10+1 (stable) and 17+1 (ccw) periods from 1,000 runs have been plotted. For keeping the comparison fair, the eGD graphs show 750 and 889 randomly selected dots. Rows 1 and 3 illustrate the low efficiency that is achieved in the runoff between eGD, ZIP and eME.*

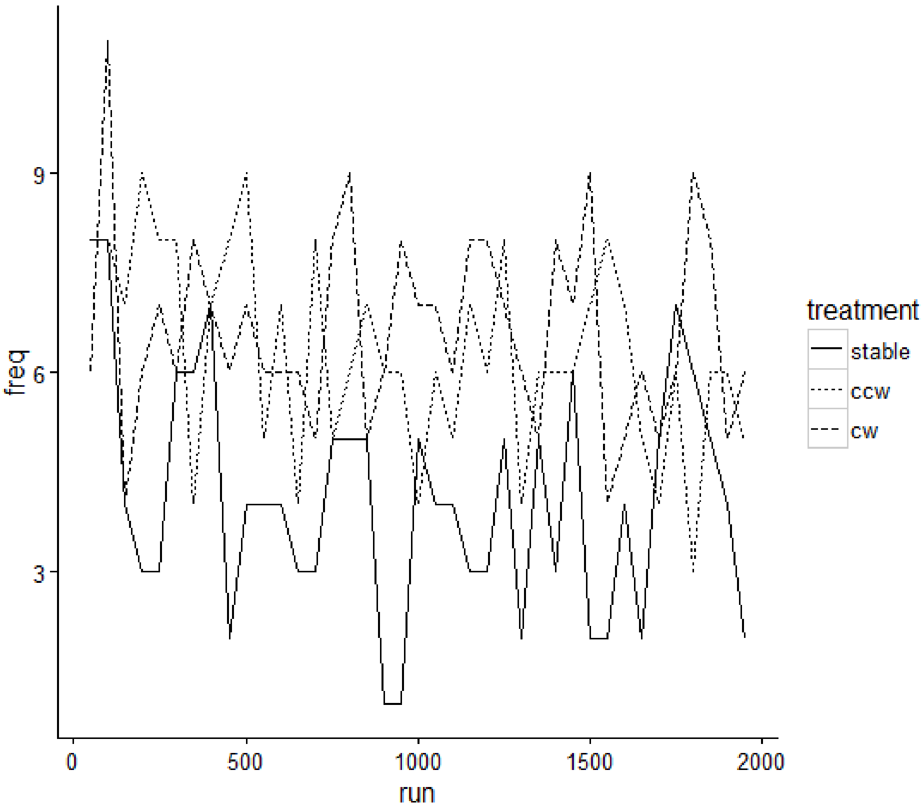

Figure 6.3 – *Frequency of changes between different behaviors (i.e. different values of ρ) as fifteen traders try to learn the best entropy-sensitivity. The curves are plotted at selected runs. The figure shows no signs of convergence.*

maximization) does not strongly dominated entropy-sensitive preferences and that heterogeneity is an ecologically rational result.

Table 6.7 – Learning entropy-sensitivity

treatment	-0.002	-0.001	0.000	0.001	0.002	length
stable		7	8			2,000
ccw			3	6	6	2,000
cw		1	9	2	3	2,000

Distributions of traders over different values of entropy-sensitivity. The initial distribution is the same for all three treatments; care has been taken to vary behavior across trader types. The stable treatment favors uncertainty aversion while the unstable treatments promote gambling.

Table 6.8 – Learning a rule for selecting a best opportunity

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

Final distributions of traders over rules for selecting a best alternative from a set of perceived opportunities and the length of the simulations. The maximum length is 2,000 runs. The initial distribution is the same for all treatments; care has been taken to vary behavior across trader types. The rules of thumb clearly emerge as ecologically rational.

6.2.7 Selecting the best opportunity

The calibration of choice from a set of perceived alternative actions favored rules of thumb, c.f. chapter 5. As it turns out, these rules of thumb are also ecologically rational, c.f. table 6.8. It is interesting that expected utility maximization survives in the unstable treatments while cumulative prospect theory does not. The latter better predicts human actions in the counter clockwise economy, c.f. table 5.2.

6.2.8 ESP versus CPT

The evidence with respect to whether ESP or CPT predicts human actions best is mixed. In the stable treatment, ESP does slightly better than CPT while the latter dominates ESP in the unstable ccw treatment. By letting them compete in an evolutionary competition we find that ESP is ecologically rational where CPT is not. The absence of neutral entropy-sensitivity tends to favor gambling in the stable and the cw unstable treatments. It also makes a comparison with table 6.7 difficult.[7]

Table 6.9 – ESP versus CPT

treatment	-0.002	-0.001	0.001	0.002	CPT	length
stable	6	3	6			2,000
ccw	3		8	3	1	2,000
cw		3	4	8		2,000

Distributions of traders over different values of entropy-sensitivity (columns "-0.002" to "0.002") and cumulative prospect theory (column "CPT"). The initial distribution is the same for all three treatments; care has been taken to vary behavior across trader types. Entropy-sensitive preferences clearly dominate cumulative prospect theory.

6.2.9 Arbitrage

Arbitrage is about recognizing opportunities for making a profit, but also about what happens after an opportunity has been seized. Appealing to the theory of

[7] The reason for excluding $\rho = 0$ is that we want to have a run-off between ESP and CPT, i.e. without EU.

Table 6.10 – Arbitrage

treatment	arbitrage					length
	0.1/0.95	0.2/0.95	0.3/0.90	0.9/0.95	no	
stable	6	4	4	1		2,000
ccw	2		4	4	5	2,000
cw		8		2	5	2,000

Distributions of traders over different values of risk attitudes. Values x/y in the column headers denote the arbitrage margin and threshold respectively. Column "no" refers to non-speculative behavior. The initial distribution is the same for all three treatments; care has been taken to vary behavior across trader types.

mental accounting we have assumed that traders want to make a profit on individual arbitrages. This induces endogenous constraints on subsequent trading, that depend on the risk attitude.

Although the prediction of arbitrage moves is below par, there exist calibrations of the risk attitude that render our model of arbitrage behavior a slightly better explanation of human trading behavior than non-speculative behavior. It is interesting to subject arbitrage behavior to evolutionary competition because the calibration based on prediction rates is not entirely satisfactory: if the margin is high then most arbitrages can be successfully completed, implying that the threshold for keeping accounts open is not a binding constraint. Furthermore, absent a solid understanding of how human traders perceive arbitrage opportunities, the induced constraints could be a real setback. Table 6.10, however, shows that arbitrage is an ecologically rational strategy.

This result presumably serves as a "worst case" in two respects. If one were to calibrate the granularity of mental accounts (and hence the degree of myopia) the rational of engaging in arbitrage will improve. Having more heterogeneity in price expectations most likely will also emphasize the relative attractiveness of arbitrage (here all traders have eGD-expectations).

Table 6.10 also illustrates that we should expect heterogeneity with respect to arbitrage behavior. In each treatment, after 2,000 runs, traders still switch between "strategies".

Learning in the stable treatment more or less agrees with the previous calibration based on prediction rates. That calibration suggested 0.2/0.95 for the stable treatment and 0.9/0.95 for the counter clockwise treatment. The unlikely high margin of 0.9 is not strongly dominated in the counter clockwise economy. Similar to other tests, preferred behaviors differ between the unstable treatments. It is not yet clear why this is the case.

6.3 Efficiency

Does learning increase efficiency? If traders can switch between strategies they can select the one which works best for them. If different traders choose different strategies utility levels could generally be higher than in homogenous configurations. To investigate this we look at efficiency; this measure compares the actual average utility to potential average utility, c.f. table 6.11.

Heterogeneity can lead to a positive symbiosis, as is the case for different attitudes towards target setting. However, it can also lead to a negative symbiosis, e.g. in the run-off between eGD-expectations and ZIP and while learning entropy sensitivity. The reasons for a negative symbiosis are not clear; especially not since ZIP strongly dominates eGD-expectations. One may also wonder about the high efficiency of GDW: why is GDW (i.e. price setting) not strongly dominant if its efficiency is so much higher than that of eGD (i.e. price taking)?[8] We leave these and other questions relating to heterogeneity for future research.[9]

6.4 Discussion

6.4.1 Capturing human trading behavior

Analyzing algorithms from the perspective of ecological rationality confirms some of our previous results, but it also leads to some surprises. Most notably, the out-performance of ZIP-expectations is unexpected. This may be due to the fact that eGD-expectations become inflexible over time, which poses restrictions on the set of feasible actions. ZIP-traders, on the other hand, adjust their expectations so as to remain competitive. The inflexibility of eGD-expectations can be reduced by calibrating the maximum number of observations that determine price expectations. What is also remarkable about ZIP is that it reduces efficiency in the run-off with eGD-trading.

Since we only have data of human trading in the stable and counter clockwise treatments, we cannot assess how well monopolistic competition predicts human moves in the clockwise economy. Here, apparently, monopolistic competition is an ecologically rational strategy, although it does not strongly dominate eGD-expectations. As expected, monopolistic competition is not successful in the stable and counter clockwise treatments. This difference between the clockwise and the counter clockwise treatments is also remarkable, because the unstable economies appear to be quite similar. In particular, in both treatments traders of type I and II operate in one market and agents of type III trade in two markets.

With respect to attitudes towards utility targets, the pre-calibration favored fixed feasibility. In this approach traders adjust their target so as to stabilize the probability of achieving it. Ambitious traders adopt targets that have a probability of less than 0.5 of being achieved, while more cautious agents choose targets that are easier to achieve. Such a strategy turns out to be not ecologically rational. Instead, optimized

[8]The answer could be due to specialization: traders of type III all prefer price-taking while the other traders opt for price-setting.

[9]For instance, why are the results of the unstable Scarf economies often very different? Or how does heterogeneity affect price formation and orbiting? In the case of speculative behavior, we have seen that the confidence intervals of the heterogenous population, consisting of seven eGD- and eight eEMA-traders, in the stable treatment are better than the corresponding homogenous confidence intervals. In the case of learning entropy-sensitivity, heterogeneity makes orbiting in the wrong direction in the ccw treatment less pronounced.

Table 6.11 – Impact on efficiency

experiment	case	pure	mix
human subjects	stb	0.92	
human subjects	ccw	0.92	
attitude toward target (stk / opt)	stb	0.24 / 0.25	0.30
attitude toward target (stk / opt)	ccw	0.35 / 0.38	0.50
monopolistic competition (eGD / GDW)	cw	0.39 / 0.81	0.75
expectation formation (eGD / ZIP)	ccw	0.49 / 0.68	0.07
expectation formation (eGD / ZIP)	cw	0.39 / 0.76	0.08
ESP ($\rho = 0$ / $\rho \neq 0$)	stb	0.30	0.31
ESP ($\rho = 0$ / $\rho \neq 0$)	ccw	0.30	0.05
ESP ($\rho = 0$ / $\rho \neq 0$)	cw	0.30	0.07
select best option (RoTh / EU)	ccw	0.49 / 0.07	0.49
select best option (RoTh / EU)	cw	0.39 / 0.07	0.39
arbitrage (no / yes)	stb	0.35	0.35
arbitrage (no / yes)	ccw	0.49	0.49
arbitrage (no / yes)	cw	0.39	0.38

(Average) efficiency levels per the end of a run. The scores of human subjects are per the end of the stable and ccw treatment. Homogenous averages are calculated over 1,000 runs and heterogenous averages over the last 100 runs. The names between brackets refer to strategies; column "pure" gives the corresponding averages of the homogenous populations. In case of "ESP" and "arbitrage", the pure configuration are $\rho = 0$ and "no" respectively. The mixed case of "monopolistic competition" assumes rules of thumb (c.f. table 6.4). The mixed cases of "expectation formation" and "select best option" come from tables 6.6 and 6.8. These tables report a competition between three strategies, of which the strongly dominated strategy is not mentioned in column "pure". Human traders do much better than robot traders, partly because they generate more transactions.

targets do well in the stable and clockwise examples, while sticky targets outperform in the counter clockwise treatment.

The rules of thumb not only predict human moves best, they are also ecologically rational in each of the three Scarf examples. Here, the only surprise is that the ranking of expected utility maximization and cumulative prospect theory is reversed. Since both are strongly dominated by the rules of thumb this is less relevant.

Orbiting is not the only phenomenon that can be used to discriminate between different behavioral hypotheses. Our simulations suggest that trading behavior is observationally distinct in the unstable economies. Once we understand how to explain this that knowledge perhaps may be used to improve our model of human trading behavior.

6.4.2 Disequilibrium theory

In this chapter, we have found that applying rules of thumb for prioritizing feasible actions is ecologically rational. We also found price setting in accordance with the theory of monopolistic competition to be strongly dominated by reservation prices that are anchored by price expectations.

As expected, competition between different strategies often results in heterogeneity. But persistent heterogeneity does not necessarily have a strong impact on price formation. Apparently, already much can be learned from simpler homogenous examples. Heterogeneity raises several different questions, that are best postponed until robot trading better captures human trading behavior. For instance, why would evolutionary learning often lead to different outcomes in the two unstable Scarf economies (they appear to be quite similar in relevant characteristics). Once we understand what causes these differences in robot behavior this knowledge perhaps may be applied to differentiate between rival behavioral hypotheses.

6.5 Conclusions

In this chapter we have tested the robustness of previous simulation results. We have subjected eGD-trading to an out-of-sample test and we have given robot traders the opportunity to switch between strategies, based on perceived relative merits.

Many of our earlier results are vindicated. For instance, that eGD-trading eventually becomes too insensitive to new information, but also that it is a rather strong algorithm. It dominates all other algorithms, except ZIP. Given the choice, robot traders unexpectedly strongly prefer the ZIP-algorithm. In line with previous results we also find that the rules of thumb for prioritizing feasible actions are ecologically rational, and that monopolistic competition is not. While previous evidence was mixed, entropy-sensitive preferences strongly dominate cumulative prospect theory.