



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.

Chapter 5

Fallible choice

How do people select a preferred action from a set of feasible alternatives? Do they disregard certain actions for "irrational" reasons? Both the selection rule and additional constraints can cause deviations from what economic textbooks consider to be rational behavior. Section 5.1 provides a motivation for treating choice as fallible. This chapter calibrates choice from the set of perceived alternatives and it models arbitrage as constraining the choice set.

For the calibration we use the data of Anderson et al.. We consider rules of thumb, expected utility maximization, prospect theory (c.f. section 5.2) and entropy-sensitive preferences (c.f. section 5.3) as alternative approaches to choosing between perceived opportunities (c.f. section 5.4). We appeal to the theory of mental accounting (c.f. section 5.5) for modeling arbitrage behavior. By assuming that traders seek profits on individual arbitrages, we can generate a maximal amount of endogenous constraints on feasible actions and induce myopic trading behavior. The calibration of arbitrage behavior can be found in section 5.5.2. Section 5.6 reflects on capturing human trading behavior and on disequilibrium theory while section 5.7 concludes.

5.1 Why choice is fallible

In the general equilibrium model, agents choose comprehensive consumption plans. By assumption, each agent has all the information that is required to make an optimal choice. In particular, everyone knows all of the available contingent commodities and their prices, which are treated as given. Choice behavior in the general equilibrium model consists of maximizing expected utility; it has been formalized by von Neumann and Morgenstern (1944) and by Savage (1954). With trading at all prices, on the other hand, choice is best understood as incremental and sequential rather than comprehensive. Instead of choosing integral consumption plans, agents repeatedly select actions from sets of perceived opportunities. Decisions like these also lend themselves to utility maximization, but does utility maximization characterize human decision-making?

One of the earliest criticisms of the idea of choice based on utility maximization, i.e. of (subjective) expected utility (SEU-) theory, was brought forward by Allais (1953). In a paradox that was named after him, people typically make choices that

appear to be inconsistent and contradicting SEU-theory.¹ The anomaly seems robust (c.f. Kahneman and Tversky (1979), MacCrimmon and Larsson (1979)). Nevertheless, many economists do not consider the Allais paradox to be a falsifier of SEU-theory. For instance, when Savage himself learned that he had chosen inconsistently, he simply replied that he had made a mistake. He *wanted* to choose consistently, and therefore he changed his choice. This seems to be the typical response, c.f. Binmore (2009, p.22-23). How people want to choose indeed is relevant to the extent that it reveals how behavior would change given the opportunity of learning.

Apart from the Allais' critique there are additional issues that make choice fallible rather than perfectly optimal. For instance, it is difficult to see why perceived choice sets would be complete. Alternative options have to be mentally constructed. According to Shah and Ludwig (2016), people often perceive only a few ways to navigate complex situations; their awareness of available options is limited by culture, previous experience, and also by unfounded implicit assumptions (e.g. with respect to feasibility). Especially when choice is sequential there are too many ambiguous contingencies. Furthermore, the perception of opportunities critically depends on expectations, that vary across individuals and over time.² Agents also may decide to constrain choice sets, e.g. by introducing budgets or other expenditure rules. Such artificial restrictions may serve to simplify choices or to manage self-control issues (c.f. section 5.5). Analyzing behavior on the assumption that decision-makers have actively explored and know what is possible in a particular situation seems simply heroic.

When the choice set is clear, alternatives may be difficult to compare. In 1935, Herbert Simon faced the problem of understanding the allocation of the city budget between maintenance by the parks department and programs run by the public schools in Milwaukee. He could not see how the marginal benefits of the two activities could be assessed and compared, much less equalized (recounted in Sunder (2003)). Furthermore, knowing an alternative in the choice set is not necessarily the same as understanding that option. People sometimes give inconsistent answers to the same question, depending on whether it is framed positively or negatively (e.g. Kahneman and Tversky (1979)).

Like choice sets, preferences also have to be mentally constructed. For instance by recollecting experienced utility. In case of doubt, a coin may be flipped to elicit one's preference.³ Marketing employs many techniques that influence the mental

¹The so-called Allais paradox has sparked research into different kinds of explanations. The apparent preference reversals have been attributed to indifference curves, which fan out (c.f. Machina (1982)); to cardinality specific utility, that depends on the number of outcomes (c.f. Neilson (1992); Humphrey (2001)); to overweighting small probabilities (e.g. Kahneman and Tversky (1979)) and to configural weight theories, which posit that people try to avoid the worst outcomes instead of obtaining the best outcome (c.f. Birnbaum (1999)). Weber (2007) suggests that the paradox may be due to a heuristic aimed at simplifying decision making. In section 5.3, we illustrate that the apparent preference reversals may also be due to differences in the complexity of the lotteries.

²In the general equilibrium model there are markets for all contingent commodities; this implies intersubjectivity with regard to states of Nature, which is absent if choice is sequential.

³Subjecting oneself to the result of flipping a coin may lead to the sensation of a loss and hence to a better understanding of one's true preferences: "I decided to flip a coin, like either heads or tails / Would let me know if I should go back to ship or back to jail. / So I hocked my sailor's suit an' I got a coin to flip. It came up tails, it rhymed with sails, so I made it back to the ship.", from:

construction of our preferences, like adding irrelevant alternatives or fixing a frame of reference that makes a particular alternative look good. While an entire industry thrives on it, these factors violate rational choice theory, or have no place in it.⁴

Anand (1993, p. 58-59) provides an interesting example of someone contemplating the question of whether people normally should have access to their own medical records. In a Socratic way, he advances different arguments illustrating that preferences can depend on the expected impact of counter-factual effects. Moreover, preferences may easily reverse as an understanding of the issue evolves, making them dependent on the depth of the analysis.⁵ Rabin's review of psychology for economists suggests that we tend to infer too much from too little: we overestimate the representativeness of small samples, see too many patterns and frequently ignore base-rates. Furthermore, we are biased toward confirmation of what we already believe and toward instantaneous gratification of our desires, c.f. Rabin (1998). Since recently, we also have ample opportunities to filter news coming our way.

(Good) information typically is costly to acquire; as a result it may be rational for economic agents to take decisions based on incomplete information instead of investing in complete information. For instance, Sims (2003), on rational inattention, considers individual agents learning from aggregate time series. Incomplete information raises issues similar to those associated with the Rational Expectations Hypothesis.⁶

If decisions are extremely context-sensitive then there is little that can be said about behavior. An external observer would have to model decisions as being random, if only for lack of relevant information. However, that would be overly relativistic. In the Scarf examples, there is more coherence in human trading behavior than is displayed by Zero-Intelligence traders. Coherence, whether it is based on rationality or on bias, is the subject of behavioral hypotheses. For the purpose of modeling disequilibrium behavior, we see no objection in assuming that agents behave consistently over time nor in agents trying to further their own interests.^{7,8, 9}

"Bob Dylan's 115th dream" by Bob Dylan.

⁴Thaler calls them SIFs: Supposedly Irrelevant Factors, c.f. Thaler (2015).

⁵Anand (1993, ch. 2) presents a careful review of some of the experimental evidence against rational preferences and expected utility maximization.

⁶Arrow (1990) argues that (i) most markets for contingent commodities do not exist, creating a gap that must be filled by conjectures; (ii) these predictions will have to depend on the forecasts of other agents; and (iii) they may affect behavior, so as to become self-denying or self-fulfilling; furthermore (iv) incomplete markets often imply a continuum of equilibria, which raises the question of how traders can coordinate on a single outcome among the infinite number of possibilities?

⁷There are infinitely many ways in which individual decision-makers can deviate from textbook rationality. Sims (1980) has called this "the wilderness of bounded rationality". In order to discipline research, it is sometimes proposed that economic agents should be modeled as having correct expectations about phenomena that they can observe, c.f. Sargent (1993). Obviously, this rules out that agents can have persistent biases, and therefore we reject specific prescriptions of how economic agents should or should not be modeled.

⁸However, bargaining may also involve other-regarding behavior, e.g. when people sacrifice opportunities in order to retaliate against behavior that they consider unfair.

⁹Boland (2014) argues that optimization precludes a proper disequilibrium theory; his argument, however, does not sufficiently distinguish between behavior and its results. If an agent, at some point in time, is in the interior of his consumption set, that does not mean that he hasn't tried to obtain an optimal choice. Indeed, with sequential choice all agents (including those who optimize)

Once we acknowledge that choice is fallible rather than perfectly optimal we face real challenges such as modeling expectation formation, the perception of opportunities and a choice between them.

5.2 Prospect theory

Expected utility theory explains choice based on its consequences. This, of course, is a simplification. The process of decision making can also have a profound effect on the outcome. For instance, according to Gneezy et al. (2003), people who are exposed to more information tend to take smaller risks. Prospect theory tries to capture the impact of the process on the outcomes of decision-making. A simple or reduced prospect (or lottery or gamble) $\mathcal{L} = (\mathbf{x}; \mathbf{p})$ is a list of prizes \mathbf{x} and a probability vector \mathbf{p} , with p_n the probability of winning prize x_n , $\forall n : x_n \in [L, H] \subset \mathbb{R}$ and $\mathbf{p} \in S^{N-1}$, with S^{N-1} the $N - 1$ -dimensional simplex. A prize can be different things; here, it will be either a money value or a change in utility (the context will make clear which interpretation applies). A prize $x \in \mathbb{R}$ can itself be interpreted as a degenerate gamble that yields x with probability 1, $\mathcal{L} = (x; 1)$. A compound lottery has one or more prizes that themselves are non-degenerate lotteries. If $0 < \sum_{j=1}^k \theta_j < 1$ then $(x_1, \dots, x_k; \theta_1, \dots, \theta_k)$ is short-hand for a lottery with a default prize of zero, $(x_1, \dots, x_k, 0; \theta_1, \dots, \theta_k, 1 - \sum_{j=1}^k \theta_j)$.

Kahneman and Tversky (1979) objects to expected utility theory, because it is not descriptive of human decision making. For instance, it fails to explain why people are sensitive to how decisions are framed. Furthermore, it cannot account for the Allais paradox nor for the strong experienced asymmetry between losses and gains. People try to avoid losses, but at the same time they can be risk seeking, e.g. they often prefer small probabilities of winning a large prize to the expected value. Moreover, their willingness to bet does not only depend on the degree of uncertainty, but also on its source (as illustrated by the so-called Ellsberg paradox¹⁰). The core propositions of prospect theory hold that people (i) perceive alternatives as consisting of gains and losses relative to a certain reference level; and (ii) that they interpret probabilities, overweighting small probabilities.

By coding options as gains and / or losses, people focus on changes rather than on absolute levels of wealth (as in expected utility theory). Furthermore, changes are often considered one at a time, implying that choice is incremental rather than comprehensive. Circumstances may influence the way how people perceive, or frame, choices.¹¹ According to prospect theory people value the possibility of owning \$1,000,000 differently, depending on whether they currently own, say, \$100,000 or \$1,900,000. Sometimes, a choice problem (implicitly) suggests a reference level. For instance, the effect of a drug can be described as saving m out of n lives, but also as failing to prevent death in $n - m$ out of n cases. The former, positive formula-

are in the interior of their consumption sets most of the time.

¹⁰See footnote footnote 14 on page 93.

¹¹We prefer to use framing in the more neutral sense of people determining for themselves what has to be decided, while Kahneman and Tversky (1979) use framing for the way tests are administered to subjects.

tion induces people to take zero as the reference level, leading to a perceived gain of m cures. The latter, negative formulation takes n lives as the reference point and emphasizes the loss, consisting of $n - m$ deaths.

A diminishing sensitivity to (monetary) gains and losses is expressed by an S-shaped value function that is convex in losses and concave in gains. Prospect theory stresses that losses loom larger than gains: it suggests that the above loss of \$900,000 roughly hurts twice as much as the gain of \$900,000 gives pleasure. Hence, the S-shaped value function is kinked and steeper in losses than it is in gains when it approaches zero.

Kahneman and Tversky (1979) argues that preferences are not linear in probabilities. In particular, the Allais paradox is taken to show that people tend to overweigh small probabilities. This is expressed by a transformation of probabilities into decision weights. In the original formulation, prospect theory mapped individual probabilities to decision weights. Tversky and Kahneman (1992) adopts the rank-dependent method of transforming probabilities, i.e. it maps cumulative probabilities to decision weights.¹² The idea here is to overweigh only unlikely outcomes. Given money holdings m , the value of a risky monetary prospect $\mathcal{L} = (x_1, \dots, x_n; p_1, \dots, p_n)$ can alternatively be expressed as

$$\begin{aligned} \sum_k u(m + x_k) p_k & \quad (\text{expected utility theory}) \\ \sum_k S(x_k) \pi(p_k) & \quad (\text{prospect theory}) \\ \sum_l S(x_l) \pi_l^- + \sum_g S(x_g) \pi_g^+ & \quad (\text{cumulative prospect theory}) \end{aligned}$$

with $S(\cdot)$ the value function and $\pi(\cdot)$ the transformation of probabilities; following Fennema and Wakker (1997), let $x_1 \leq \dots \leq x_r \leq 0 \leq x_{r+1} \leq \dots \leq x_n$ and let weighing functions w^- and w^+ for probabilities be associated with losses and gains respectively and

$$\begin{aligned} \pi_1^- &= w^-(p_1); & \pi_l^- &= w^-\left(\sum_{k=1}^l p_k\right) - w^-\left(\sum_{k=1}^{l-1} p_k\right) & 2 \leq l \leq r; \\ \pi_n^+ &= w^+(p_n); & \pi_g^+ &= w^+\left(\sum_{k=g}^n p_k\right) - w^+\left(\sum_{k=g+1}^n p_k\right) & r+1 \leq g \leq n-1. \end{aligned}$$

According to prospect theory, decision making consists of an editing and an evaluation phase. In the editing phase, people determine what has to be decided. This covers simplifications such as removing stochastically dominated alternatives from the considerations, or rounding of probabilities (e.g. people are expected to treat highly likely outcomes as certain) or collapsing prospects into equivalent, simpler versions. For instance, suppose that \mathcal{L} contains outcomes j and k with $x_j = x_k$ then people simplify \mathcal{L} by combining j and k , i.e. by assigning a probability $p_j + p_k$ to outcome x_j . Failure to remove dominated alternatives and / or to apply simplifications can result

¹²Tversky and Kahneman (1992) has also extended cumulative prospect theory to cover ambiguity, i.e. the case of unknown probabilities, see also Wakker (2010).

in inefficient and / or inconsistent choices. In cumulative prospect theory, neither is an issue because here cumulative probabilities are transformed.

We propose to test prospect theory based on actual trading behavior. This can be done by creating a context for those decisions, i.e. by simulating choice sets of perceived opportunities. In subsequent simulations (c.f. section 5.4), the specifications of the value function, $S(x)$, and of the decision weights, $\pi^-(p)$ and $\pi^+(p)$, are taken from Booiij and Van Praag (2009). These functionals of prospect theory all have to be estimated, which leads to a simultaneity problem. The typical econometric solution consists of using parametric forms for all three functions or of using a particular design that allows the value function and decision weights to be identified non-parametrically. Both approaches have drawbacks, that Booiij and Van Praag seek to avoid.¹³ They have estimated a complete prospect theory model, allowing for decision errors, and using a rich dataset that permits a non-parametric identification of the value function and the decision weights. Compared to other studies, Booiij and Van Praag (2009) finds its estimated value function to be closer to linear, and its decision weights to express more pessimism, and hence induce more risk aversion. Let $I(b)$ be the indicator-function: it is equal to 1 if b is true and equal to 0 otherwise. Following Booiij and Van Praag (2009), we use

$$\begin{aligned} S(x) &= x^{0.859} I(x \geq 0) - 1.576 (-x)^{0.826} I(x < 0); \\ \pi^+(p) &= e^{-0.772(-\log p)^{0.618}}; \\ \pi^-(p) &= e^{-1.022(-\log p)^{0.592}}. \end{aligned}$$

The value function takes monetary inputs; opportunities for action represented by prospects, as derived in section B.1.2, yield utility prizes. For being consistent, these utility prizes are converted to equivalent changes in budget before they are valued by $S(x)$.

5.3 Entropy-sensitive preferences

"... the sign of a crank or a half-baked speculator in the social sciences is his search for something in the social system that corresponds to the physicist's notion of entropy.", Samuelson (1972, p. 450).

5.3.1 Definition, justification and embedding

According to choice theory, if two prospects have the same expected value then they should be considered equivalent. But does expected value alone determine the ranking of opportunities? We propose that this is not the case; our current notion of rationality is too narrowly focused on the expected value of consequences. A richer view of consequences would also give us a broader perspective on attitudes towards

¹³In the parametric approach, misspecification of one function may lead to contamination of the parameters of the other functions. Allowing for unobserved heterogeneity also takes a parametric form, leading to a dependence on the choice of a stochastic error process. The non-parametric method estimates at the individual level, allowing for full heterogeneity. It requires the data to have a chained character, which may induce error propagation and an incentive compatibility problem.

risk. Consequences have other attributes as well, e.g. predictability. People appear to be sensitive to the trade-off between expected value and predictably.^{14,15}

This is captured by entropy-sensitive preferences defined as

$$\tilde{u}(\mathcal{L}) = \sum_{n=1}^N u(x_n)p_n + \rho H(\mathbf{p}) \quad (5.3.1)$$

with $H(\mathbf{p}) = -\sum_{n=1}^N p_n \log(p_n)$ the (Shannon) entropy of the discrete distribution \mathbf{p} .¹⁶ Entropy is a measure of the lack of predictability embodied in a probability distribution. If the distribution is continuous then the summations are replaced by integrals.

Observe that $\rho > 0$ implies a preference for gambling (since $H(\mathbf{p}) > 0$) and $\rho < 0$ means that risk aversion is also being driven by entropy. In both cases, agents consider the trade-off between expected consequences and the predictability of a lottery. The case of $\rho = 0$ corresponds with von Neumann-Morgenstern utility functions. If a trader has to assess the value of commodity bundles or of degenerate rather than non-degenerate lotteries, then the entropy term drops out of (5.3.1).

¹⁴ Many of the paradoxes of choice can be solved by introducing a trade-off between expected value and lack of predictability. One easily verifies that adding an entropy term to a utility function can solve the Allais paradox, the common ratio paradox and probabilistic insurance aversion. Ellsberg's paradox, which is usually taken as proof that people are averse to ambiguity, can also be explained in terms of predictability.

It is instructive to consider the one urn example: an urn contains 90 balls: 30 red, and 60 either black or yellow. A player must first choose one of two bets: A (Receive \$100 upon drawing a red ball) or B (Receive \$100 upon drawing a black ball). After having made this choice, the player must select another bet: either A' (Receive \$100 if you draw a red or a yellow ball) or B' (Receive \$100 if you draw a black or a yellow ball). Most people strictly prefer A to B and B' to A', which appears to be inconsistent. The strict preference is taken to signal a certain belief about the proportion of black balls. $A \succ B$ implies a belief that there are less black than red balls. On the other hand, $B' \succ A'$ signals a believe that there are strictly more black than red balls.

Instead of invoking an aversion to ambiguity, one can also explain the observed preferences based on how people can simplify the gambles: Let R , B and Y stand for the number of red, black and yellow balls respectively. The urn, which can be described by $[R, B, Y]$ can be in 61 different equi-probable states, $[R, B, Y] = [30, 0, 60], [30, 1, 59], \dots, [30, 60, 0]$. With each state, we have a particular version of the bets A and B, e.g. if $[R, B, Y] = [30, 22, 38]$, then bet A is equal to $\mathfrak{L} = ((\$100, \$0, \$0), (\frac{1}{3}, \frac{11}{45}, \frac{19}{45}))$ which can be simplified to $\mathfrak{L} = ((\$100, \$0), (\frac{1}{3}, \frac{2}{3}))$. For bet A, this is true whatever the state the urn is in. If each equi-probable state wins the same lottery, then we can simplify further and conclude that bet A is equivalent to $\mathfrak{L}_A = ((\$100, \$0), (\frac{1}{3}, \frac{2}{3}))$. Bet B also allows a simplification: $[R, B, Y] = [30, m, 60 - m]$ implies $\mathfrak{L}_m = ((\$100, \$0), (\frac{m}{90}, \frac{90-m}{90}))$, $m = 0, 1, \dots, 60$. But since these lotteries are all different, bet B is equivalent to the compound lottery $\mathfrak{L}_B = ((\mathfrak{L}_0, \mathfrak{L}_2, \dots, \mathfrak{L}_{60}), (\frac{1}{61}, \frac{1}{61}, \dots, \frac{1}{61}))$. This lottery is less attractive than \mathfrak{L}_A to entropy averse decision-makers. A similar reasoning gives $\mathcal{L}_{B'} \succ \mathcal{L}_{A'}$.

While \mathcal{L}_A and \mathcal{L}_B are equivalent in expectation, they differ in predictability. This is typical for the paradoxes of choice.

¹⁵In appendix B, we propose that it would be more rational for traders to prefer strategic offers over regular offers, c.f. section B.1.1. Interestingly, however, human traders prefer to submit regular instead of strategic offers. A trade-off between expected value and predictability potentially can explain this preference ordering of human traders, because strategic offers typically are less predictable.

¹⁶As early as 1976, J.R Meginniss added entropy to the utility function for capturing the utility of gambling. Being not satisfied with his ad hoc derivation of entropy-sensitive utility functions, Ng, Luce and Marley extended the approach of Meginniss, and gave an axiomatic derivation based on joint receipts, (c.f. Luce and Marley (2000); Luce et al. (2008); Ng et al. (2008)).

The application of $\tilde{u}(\cdot)$ to a lottery \mathcal{L} requires that the prospect is simplified, at least to the point where all prizes are different. To see this, note that a gamble like $\mathcal{L} = ((x, x), (\alpha, 1 - \alpha))$ is equivalent to getting x for sure; sensitivity to the entropy of $(\alpha, 1 - \alpha)$, however, would lead to $\tilde{u}(\mathcal{L}) \neq u(x)$. The axiom of independence avoids contradictions due to superficial differences between prospects by ruling out entropy terms.

Referring to a preference for gambling, Diecidue et al. argue that it "has become increasingly well understood that models of decision making, to be descriptively accurate, have to incorporate basic violations of rationality, based primarily on findings from the psychological literature.", Diecidue et al. (2004, p. 243). This, of course, underscores our argument that choice is fallible. Behavior, as displayed by real decision-makers, is different from what is implied by the axioms of choice theory. If we are to accept violations of rationality, however, then we (i) forsake an opportunity to scrutinize our current notion of rationality and (ii) we may be in for ad hoc theories. Instead we should contemplate whether our notion of rationality can be meaningfully generalized. We disagree with Le Menestrel (2001), which considers a utility of gambling to be a matter of process instead of consequence and therefore calls for excluding gambling from the consequentialist theory of expected utility. For demonstrating that this is misguided advice, it suffices to show that ESP fit into the framework of expected utility theory.

Below we show that the axiom of independence, after it has been slightly weakened, admits utility functions that include an entropy term. Expression 5.3.1 therefore is not just an arbitrary and ad hoc modification of preferences.

Preferences over lotteries are typically modeled as being transitive and complete, and among other things, as satisfying the so-called axiom of independence (or sure thing principle).

Definition 5.1 (Axiom of independence). $\forall \mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3 \in \mathfrak{L} : \exists \alpha \in (0, 1) : \mathcal{L}_1 \succeq \mathcal{L}_2 \iff \alpha \mathcal{L}_1 + (1 - \alpha) \mathcal{L}_3 \succeq \alpha \mathcal{L}_2 + (1 - \alpha) \mathcal{L}_3$.

In words, a preference between two lotteries should survive mixing with a third gamble, and the preference between lotteries should not depend on irrelevant alternatives (like \mathcal{L}_3).

The sure thing principle implies a linear dependence of preferences on probabilities, that leads to the well-known Expected Utility formulation, $\tilde{u}(\mathcal{L}) = \sum_{n=1}^N u(x_n) p_n$. This formula stipulates that agents rank lotteries according to their expected utility. The cancellation law, which is contained in the axiom of independence, $\mathcal{L}_1 \succeq \mathcal{L}_2 \iff \alpha \mathcal{L}_1 + (1 - \alpha) \mathcal{L}_3 \succeq \alpha \mathcal{L}_2 + (1 - \alpha) \mathcal{L}_3$, rules out sensitivity of preferences to the amount of uncertainty.

By requiring that lotteries be simplified to the point where all prizes are different, we can weaken the axiom of independence and thereby allow entropy to show up in utility functions. As a simpler alternative to the approach of Luce et al. we propose to restrict the application of the axiom of independence to lotteries \mathcal{L}_3 with prizes that do not already appear in either \mathcal{L}_1 or \mathcal{L}_2 .

Proposition 5.2. *Preferences $\tilde{u}(\mathcal{L}) = \sum_{n=1}^N u(x_n) p_n + \rho H(\mathbf{p})$ satisfy the axiom of independence, provided lotteries are mixed with another lottery that does not have any prizes in common with the original gambles.*

Proof. Let $\mathcal{L}_1 = (\mathbf{x}; \mathbf{p})$, $\mathcal{L}_2 = (\mathbf{y}; \mathbf{q})$ and $\mathcal{L}_3 = (\mathbf{z}; \mathbf{r})$ be lotteries with K, L and M prizes respectively, and such that \mathbf{z} does not have any element in common with either \mathbf{x} or \mathbf{y} . Lotteries \mathcal{L}_1 and \mathcal{L}_2 , on the other hand, may have one or more prizes in common. Furthermore, assume that preferences are entropy sensitive and that they can be described by $\tilde{u}(\mathcal{L}) = \sum_{n=1}^N u(x_n)p_n + \rho H(\mathbf{p})$, with $u(\cdot)$ a regular von Neumann-Morgenstern utility function. Consider

$$\begin{aligned} \alpha \mathcal{L}_1 + (1 - \alpha) \mathcal{L}_3 &\succeq \alpha \mathcal{L}_2 + (1 - \alpha) \mathcal{L}_3 \Leftrightarrow \\ (\mathbf{x}, \mathbf{z}; \alpha \mathbf{p}, (1 - \alpha) \mathbf{r}) &\succeq (\mathbf{y}, \mathbf{z}; \alpha \mathbf{q}, (1 - \alpha) \mathbf{r}) \Leftrightarrow \\ \alpha \sum_{i=1}^K u(x_i)p_i + \rho H((\alpha \mathbf{p}, (1 - \alpha) \mathbf{r})) &\geq \alpha \sum_{i=1}^L u(y_i)q_i + \rho H((\alpha \mathbf{q}, (1 - \alpha) \mathbf{r})). \end{aligned}$$

The first step compounds lotteries; after expanding $\tilde{u}((\mathbf{x}, \mathbf{z}; \alpha \mathbf{p}, (1 - \alpha) \mathbf{r}))$ and $\tilde{u}((\mathbf{y}, \mathbf{z}; \alpha \mathbf{q}, (1 - \alpha) \mathbf{r}))$ both sides contain a factor $(1 - \alpha) \sum_{i=1}^M u(z_i)r_i$, which has been canceled. Furthermore, we have

$$\begin{aligned} H((\alpha \mathbf{p}, (1 - \alpha) \mathbf{r})) &= \alpha \sum_{i=1}^K p_i \log \alpha p_i + (1 - \alpha) \sum_{i=1}^M r_i \log (1 - \alpha) r_i \\ &= \alpha \log \alpha + \alpha \sum_{i=1}^K p_i \log p_i + (1 - \alpha) \sum_{i=1}^M r_i \log (1 - \alpha) r_i. \end{aligned}$$

The first line applies the definition of the Shannon entropy; here we implicitly use the fact that \mathbf{x} and \mathbf{z} have no prizes in common. The factor $\alpha \log \alpha$ appears because $\sum_{i=1}^K p_i = 1$. For $H((\alpha \mathbf{q}, (1 - \alpha) \mathbf{r}))$ we have a similar expression. After substitution we have

$$\begin{aligned} \alpha \mathcal{L}_1 + (1 - \alpha) \mathcal{L}_3 &\succeq \alpha \mathcal{L}_2 + (1 - \alpha) \mathcal{L}_3 \Leftrightarrow \\ \sum_{i=1}^K u(x_i)p_i + \rho \sum_{i=1}^K p_i \log p_i &\geq \sum_{i=1}^L u(y_i)q_i + \rho \sum_{i=1}^L q_i \log q_i \Leftrightarrow \\ \mathcal{L}_1 &\succeq \mathcal{L}_2. \end{aligned}$$

The second line results after canceling common factors and dividing both sides by α . \square

5.3.2 Pre-calibration

The extent to which a trader is sensitive to the complexity of options to choose from is an individual characteristic that is a part of a trader's risk attitude. Because of heterogeneity and because it is of secondary importance compared to the expected value, entropy sensitivity is difficult to detect through the calibration of a random distribution of sensitivities (the method we have applied in previous calibrations). Therefore, we pre-calibrate entropy-sensitive preferences by optimizing the sensitivity per individual robot trader. We do this by comparing choices of each robot trader, conditional on entropy-sensitivity, with the alternatives selected by his human alter ego. Since the human subjects of the stable and counter-clockwise treatment of Anderson et al. are different, we consider both experiments.

Table 5.1 – Calibration of entropy sensitivity, ρ

trader \ treatment	stable	ccw
1	-0.002	0
2	0	0
3	-0.0005	0
4	0	0
5	0	0.0015
6	0	-0.002
7	0.003	0.0015
8	0	0
9	0	0.0025
10	0.0015	0
11	0.0005	0
12	0	0.003
13	0	0.002
14	0	-0.0005
15	0.002	-0.001

The values of ρ have been calibrated per individual trader by maximizing their individual prediction rates by means of a grid search over 1,000 runs. If $\rho = 0$, then trader i maximizes his expected utility. The values in the table are conservative in the sense that they are different from zero only if the best prediction rate of $\rho \neq 0$ is strictly better than the one of EU. Furthermore, if different values for ρ lead to the same prediction rate, then the one nearest to zero is reported. There is substantial heterogeneity and several traders appear to love gambling (since $\rho > 0$).

Our aim is not to calibrate the individual sensitivities to any precision, but rather to demonstrate that entropy-sensitive preferences apply to the context of trading. Therefore, we use a grid search for ρ ranging from -0.0045 to 0.003 with a step size of 0.0005. Table 5.1 presents conservative estimates of the individual sensitivities. The results are based on opportunities perceived as simple lotteries, because entropy-sensitive preferences require aggregation of identical prizes for correctly gauging entropy. Since entropy-sensitive preferences are a straightforward generalization of EU, any deviation from EU leads to a better explanation of human trading behavior (c.f. table 5.2).¹⁷

¹⁷This, of course, is to be expected; what matters is whether deviations are statistically significant. Instead of simulating confidence intervals for estimated sensitivities, we show that ESP are ecologically rational, c.f. chapter 6. That is, given the opportunity to switch between sensitivities, some robot traders prefer $\rho \neq 0$. Since entropy sensitivity is a deep parameter of the risk attitude that may not be subject to learning, this result should be understood as illustrating that heterogeneity with respect to entropy sensitivity can survive competition.

5.4 Calibration of choice between opportunities

The perception of opportunities critically depends on price expectations. In chapter 4, we found that "no arbitrage" prices, that are derived from Gjerstad-Dickhaut beliefs, capture human expectations best (although far from perfect). Here, we use these expectations to simulate choice sets of perceived opportunities. Furthermore, we use GD-beliefs to represent opportunities as lotteries, as explained in appendix B. There, we derive a valuation of opportunities that is based on an awareness that the economy is out of equilibrium and that as a result some planned transactions may not be completed.

The calibration consists of comparing different methods for selecting a best alternative from the same, given sets of opportunities.¹⁸ The data of Anderson et al. (2004) give us $5194 + 4605 = 9799$ decisions of human traders from session 414 (stable treatment) and 511 (counter-clockwise treatment) respectively. Here, we exclude the actions that can be considered as arbitrage, because our robot traders do not yet engage in arbitrage. This leaves us with $4835 + 4340 = 9175$ decisions. On average, $3801 + 3148 = 6949$ actions are considered to be a feasible choice by the eGD-algorithm (calculated over 1,000 runs each). The decisions leading to those actions form the basis of our calibration; the other actions are implemented in the simulation for maintaining the integrity of public information, but they are ignored in the calculation of recognition and prediction rates.

Taking the rules of thumb, as derived in appendix B, as a benchmark, we consider expected utility maximization, (cumulative) prospect theory and entropy-sensitive preferences. For a better understanding of the difference between expected utility maximization and prospect theory, we also add two intermediate selection rules: (i) the utility function replaced by the value function (keeping probabilities) and (ii) probabilities replaced by decision weights (keeping the utility function).

It turns out that the rules of thumb outperform both expected utility maximization and cumulative prospect theory, c.f. table 5.2. In over half of the cases, the subjects of Anderson et al. choose differently from expected utility maximization. Cumulative prospect theory also gives better predictions of human trading behavior than expected utility maximization. This can be attributed to mapping probabilities to decision weights. Evidence with respect to the relative merits of ESP and CPT is mixed.

Table 5.2 may appear to be inconsistent with table 4.6 on page 67, since both report on eGD-expectations in combination with rules of thumb for prioritizing feasible actions. For instance, table 4.6 reports an overall prediction rate of 68% in the stable Scarf economy while table 5.2 puts this percentage at 70%. The difference is due to

¹⁸One may perhaps expect that choice sets in two simulations with different selection rules will be synchronized, because the selection rules do not affect the choice set to which they are applied. That, however, is not necessarily the case. The set of perceived opportunities depends on price expectations and these are random. To synchronize price expectations across different simulations the random generator must be started with the same seed. This is necessary, but not sufficient. If a different selection rule leads to a different number of calls to the random generator, then we get random drift: at some point, the evaluation of the feasibility of the same human decision starts to differ because the state of the random generator is not the same across the simulations. To prevent the adverse impact from random drift, the code has been adjusted so that, for the purpose of the calibration of choice, all selection rules are applied to the same set of perceived opportunities *before* the next human move is taken into consideration.

table 4.6 taking all decisions into account and table 5.2 only the ones in which the human move is recognized as a feasible action.

Human trading achieved a high degree of efficiency in the stable Scarf economy because end-of-period allocations are concentrated around the midpoints of the sides of the triangles, i.e. around values that correspond to the Walrasian equilibrium, c.f. figures 3.6 and 3.7. Robot trading, however, is not equally efficient: the spread in end of period allocations is substantial and observations are relatively far removed from their Walrasian equilibrium values, c.f. figures 5.1 and 5.2.

The fact that traders of type *III* hold a lot of money at the end of a period shows that the speed of convergence in allocation of robot trading is lower than that of human trading. Given the same number of offers, human traders eliminate more Pareto improvements. This illustrates the importance of addressing the issue of quantity setting and modeling the submission of offers. If robot traders would have generated more transactions, then allocations most likely would have been more concentrated and closer to the midpoints of the sides of the triangles. This is because traders buy what they need and sell what they can spare; hence having more transactions generally means higher efficiency.

What is also remarkable is that end of period allocations of expected utility maximization resemble those of the rules of thumb in the stable treatment and those of cumulative prospect theory in the counter clockwise treatment. As expected, the spread of money exceeds the spread of commodities 2 and 3.

5.5 Mental accounting and arbitrage

5.5.1 Modeling of arbitrage behavior

Mental accounting is a description of the way individuals and households manage their finances, c.f. Thaler (1999). It covers the entire process of coding, categorizing and evaluating events for the purpose of making good financial decisions.¹⁹ Expenditures are grouped into budgets (e.g. housing, food, holidays); wealth is allocated into accounts (e.g. cash on hand, home equity, pension wealth, etc) and income is categorized (e.g. regular income or windfall gains). When money is tight, budget rules tend to be more explicit and budgets tend to be defined over shorter periods. Budgets may also serve the purpose of self-control, e.g. to curb an urge to buy on impulse.

The way accounts are organized affects behavior because they define how agents perceive decisions. A mental account is a decision context in which an agent tries to achieve a positive result. Suppose that an agent has acquired some stock and that he has opened a new mental account for this particular purchase. Suppose furthermore that the stock has depreciated. As long as the stock is not sold the mental account associated with the purchase is kept open, and the agent holds on to a "paper loss". Selling the stock means closing the account, i.e. adding up all gains and losses,

¹⁹Not every transaction is processed in the same way; some are simply treated as "normal expenditures". If the size of the transaction increases, or if a purchase or the situation around it is unusual, then it becomes more likely that the transaction will attract more attention, e.g. it may trigger an *ex post* analysis when the account is closed.

Table 5.2 – Correct prediction of human decisions (%)

Description	stable				counter-clockwise							
	RoTh	ESP	EU	val	wght	CPT	RoTh	ESP	EU	val	wght	CPT
propose bid 2	65	58	56	54	68	68	71	59	55	56	56	57
propose bid 3	64	48	48	48	63	63	49	40	41	41	54	54
propose ask 2	57	32	34	34	40	40	37	24	25	25	25	26
propose ask 3	68	35	32	32	46	46	69	41	40	40	53	53
accept bid 2	100	87	85	85	55	55	95	86	84	84	68	67
accept bid 3	94	85	70	70	42	42	97	98	98	98	98	98
accept ask 2	94	90	88	92	52	50	100	100	100	100	94	91
accept ask 3	93	82	76	77	27	24	80	87	86	87	73	72
cancel bid 2	100	5	5	5	31	31	100	62	62	62	79	79
cancel bid 3	100	50	50	50	50	50	100	40	40	40	60	60
cancel ask 2	100	0	0	0	67	67	100	53	60	60	60	60
cancel ask 3	100	94	94	94	98	98	100	63	63	63	88	88
total	70	51	49	48	50	49	66	50	49	49	56	56

Prediction rates, conditional on human actions recognized as a feasible action. Rules of thumb (column 'RoTh') predict human actions best. A value of 100 indicates perfect predictions. Entropy-sensitive preferences (column 'ESP') do better than normal preferences (column 'EU'). Cumulative prospect theory (column 'CPT') better predicts than expected utility maximization (column 'EU') in both the stable and the ccw treatment. This is mainly due to transforming probabilities into decision weights: replacing utility by the value function while retaining probabilities (column 'val') gives a score similar to EU; combining the utility function with decision weights (column 'wght'), however, causes the score to increase.




Figure 5.1 – Shares per trader type in the available commodities in the stable treatment: rules of thumb (row 1), expected utility maximization (row 2) and cumulative prospect theory (row 3). Each dot represents an allocation at the end of a period. The results of 10+1 periods from 1,000 runs have been plotted. End of period allocations are sensitive to choice; remarkably, the results of rules of thumb and expected utility maximization are similar. CPT results are more concentrated due to a consistent lack of trade in all of the runs.




Figure 5.2 – Shares per trader type in the available commodities in the counter clockwise treatment: rules of thumb (row 1), expected utility maximization (row 2) and cumulative prospect theory (row 3). Each dot represents an allocation at the end of a period. The results of 16+1 periods from 1,000 runs have been plotted. Equilibrium allocation selection is sensitive to choice; remarkably, the results of expected utility maximization and cumulative prospect theory are similar. CPT results are concentrated due to a consistent lack of trade in all of the runs.

and acknowledging the result; in this case, accepting the loss. The theory predicts that agents will try to avoid acknowledging losses and hence that the agent will sell winners and ride losses.²⁰ This idea can be applied to modeling arbitrage: if traders seek profits on individual arbitrages then this will generate endogenous constraints on the set of feasible opportunities and these will induce myopic trading behavior.

We only consider two operations on mental accounts: opening and closing. An arbitrage account will be opened if the trader commits to an arbitrage, i.e. (i) the margin between the floor and the reservation price is acceptable and (ii) his action results in a transaction. By assumption, the trader plans to reverse this transaction at a later stage. The account will be closed upon realizing a profit, or (if there currently is a loss) if the probability of realizing a profit in the future breaches a certain threshold.²¹ The theory of mental accounting covers many other aspects of individual financial decision making, c.f. Thaler (1999), some of which could be included in a model of arbitrage behavior.²² We will, however, abstain from complicating the model beyond opening and closing mental accounts for individual arbitrages.

5.5.2 Calibration of arbitrage behavior

In section 3.4.1, we found that the willingness to engage in arbitrage strongly differs between the practice period on the one hand and normal periods on the other. Most likely this difference is due to a change in risk attitude after trading starts "for real". Table 5.3 shows recognition and prediction rates of human actions in the practice periods, with and without arbitrage behavior. Although arbitrage behavior does help to explain arbitrage moves, the prediction rates of arbitrage moves are below par: 27% in the stable practice period and 14% in the ccw training period.²³ Arbitrage behavior also adversely affects the prediction of non-arbitrage moves, but very slightly.

Arbitrage is about observing and seizing additional opportunities, but also about managing them. While a trader anticipates a profit, arbitrage can also constrain his

²⁰To a certain extent this can also be achieved by re-arranging mental accounts. The theory proposes so-called rules of hedonic framing for maintenance of mental accounts, c.f. Thaler (1999). These rules are derived from the value function of prospect theory. For simplicity's sake, we will not let traders re-arrange their mental accounts.

²¹This is similar to the Value-at-Risk constraint in behavioral portfolio theory, c.f. Das et al. (2010).

²²Gamblers, who win money in a casino early on, often put that money in a different pocket (mental account) from their own money. They are willing to take more risks with the money they have gained from the casino ("house money") than with their own money. Prior losses do not stimulate risk seeking, unless there is a possibility to break even. One could treat the proceeds of arbitrage as "house money".

In principle, traders cannot short: for arbitrage they have to buy a commodity before they can sell it at a better price. However, they can short against their mental account of (actual or planned) consumption. Consider a trader of type *III* having four units of good 2 and ten units of good 3. This trader has two excess units of commodity 3. However, suppose that he has the possibility to accept a bid for one unit of good 2 at a favorable price. Seizing this opportunity would lower his actual utility, which can only be restored if the trader is able to buy back the unit at a later time. Subjectively, such trades can be just as risky as real shorts.

²³Note that the prediction rate is also low in the practice periods because traders are experimenting. There are examples of traders willing to sell at half the expected price or to buy at double this price, even though these actions would lower their level of utility.

Table 5.3 – Recognition and prediction in practice periods

action	with arbitrage				without arbitrage			
	stable		ccw		stable		ccw	
	rec	prd	rec	prd	rec	prd	rec	prd
arbitrage	61	27	35	14	40	0	30	0
non-arbitrage	73	45	74	43	72	49	75	47

Average percentages of human actions that are recognized as a feasible option and that are correctly predicted, as simulated over 1,000 runs. The prediction rates are conditional on the human action being recognized as a feasible option. Row "arbitrage" refers to human actions identified as arbitrage; "non-arbitrage" refers to all other actions. The results "with arbitrage" refer to robot behavior that admits arbitrage based on mental accounting (margin = 0.05; threshold = 0.7). Results "without arbitrage" serve as a benchmark. In the practice periods the prediction rates of both arbitrage and non-arbitrage moves are below corresponding values in the regular periods, c.f. footnote 24 below.

Table 5.4 – Calibration of arbitrage behavior in regular periods

thld \ mrgn	stable			ccw		
	0.3	0.2	0.1	0.9	0.7	0.5
0.99	68.08	68.10	67.91			
0.97	68.08	68.10	67.91			
0.95	68.08	<u>68.10</u>	67.91	<u>62.73</u>	62.14	60.69
0.90	68.07	68.08	67.89	62.73	62.14	60.69
0.85	68.05	68.07	67.88	62.73	62.14	60.69
0.75				62.73	62.14	60.69
0.50				62.73	62.14	60.69

The table shows prediction rates (conditional on recognition) by threshold and margin, on a coarse grid, in the regular periods for both the stable and the counter clockwise treatment. The table demonstrates that the model detects a difference in risk attitude between the stable and ccw treatments.

trading behavior. In our model, arbitrage behavior is captured by two parameters: (i) the minimum margin that a trader requires to engage in arbitrage and (ii) the threshold probability for keeping open a mental account that is suffering a loss. The former directly affects the perception of arbitrage opportunities; the latter indirectly constrains the set of actions to choose from. The lower the threshold probability, the longer a trader is willing to ride losses (in the hope of recouping them later). While open, the mental account can restrict non-arbitrage opportunities. For instance, suppose a trader has bought a unit of commodity 2 at a price of 35; if that commodity currently trades for 30 then he cannot sell that unit, even though he has no use for it and would be better off with the money (assuming that the likelihood of making a profit in the future is still deemed sufficiently high). Lower thresholds, therefore, could adversely affect the prediction rates.

Table 5.5 – Impact of arbitrage on price formation

description	avg distance (money)			confidence intervals (%)					
	stable	ccw	total	stable		ccw		cw	
				good 2	good 3	good 2	good 3	good 2	good 3
humans				74	78	17	17	17	26
arbitrage	5.89	6.57	6.30	75-90	79-95	27-46	23-65	28-85	27-53
no arbitrage	5.87	6.40	6.19	66-90	80-95	25-42	26-69	26-85	26-52

Statistics are averages, calculated over 1,000 runs in a population of seven eGD- and eight eEMA-traders. Arbitrage increases the concentration of the price of commodity 2 in the stable economy. In other cases, the impact is small. Interestingly, the confidence intervals of the heterogenous population without arbitrage in the stable treatment are better than the corresponding confidence intervals of both the homogenous eGD- and eEMA-population, c.f. table 4.8.

If traders apply a high threshold and / or if they require a high margin then effectively there is no arbitrage. In that case, behavior reduces to eGD-trading. Table 5.4 shows that there exists calibrations that make the overall prediction rate with arbitrage slightly better than the values of eGD-trading (i.e. of non-speculative behavior).²⁴ The overall prediction rate, however, is not very sensitive to changes in the margin and / or the threshold. The arbitrage margin of 0.9 in the ccw economy seems implausibly high. If the margin is that high, then rare cases of arbitrages are likely to be successfully completed and hence a low threshold does not effectively constrain trading behavior.

Although arbitrage behavior is not calibrated to great precision, we apply the calibrated model to gauge the impact of arbitrage on price convergence and on end of period allocations. For this purpose we create a heterogenous population, consisting of seven eGD- and eight eEMA-traders.²⁵

We find that arbitrage slightly improves concentration (c.f. table 5.5) and orbiting (c.f. figure 5.3). On the other hand, it adds to the average distance between robot and human trading prices and it affects the spread of end of period allocations (c.f. figures 5.4 and 5.5).

Speculative behavior supposedly reduces the spread of prices and increases liquidity, making markets more efficient. From figure 5.4, however, it is clear that arbitrage does not have a discernable impact on the simulated end-of-period allocations of the stable treatment. Although it does affect the spread in the unstable treatments it has a limited impact on distributive efficiency.

²⁴Observe that table 5.4 refers to regular periods while table 5.3 refers to the practice periods. The prediction rates in table 5.4 are substantially higher than those in table 5.3. This is precisely due to the difference between practice and regular periods. Consider the prediction rate in the stable treatment. In a single batch of 1,000 runs ($\text{thld} = 0.97 \times \text{mrgn} = 0.3$) we have a an overall prediction rate of 43% in the practice periods and a rate of 68% in the regular periods following the practice period. Hence, the subjects of Anderson et al. did not yet have good price expectations and / or they were experimenting in the practice period (also with non-arbitrage moves).

²⁵As expected, arbitrage does not have a material impact on price formation in populations with homogenous eGD-expectations.

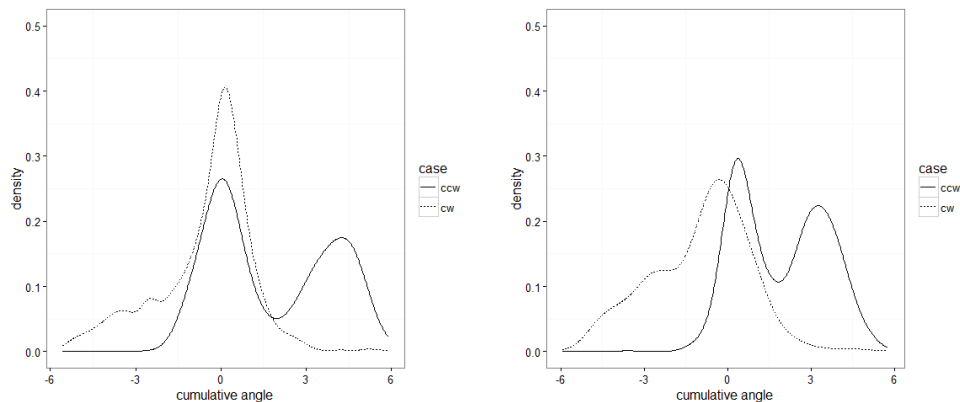


Figure 5.3 – Simulated probability densities of cumulative angles in the unstable Scarf economies, as generated by arbitrage (left) and non-arbitrage behavior in a heterogenous population. Heterogeneity makes the simulated densities less spiky compared to figure 4.8. Arbitrage moves the simulated densities in the right direction (i.e. ccw to the left and cw to the right), but not by enough.

5.6 Discussion

5.6.1 Capturing human trading behavior

Our calibration of choice between feasible actions leads to the surprising result that expected utility maximization does not do very well in explaining human trading behavior. Can this be due to eGD-expectations being inadequate?

In chapter 4 we found that, although the eGD-algorithm is best in capturing human expectations, its expected prices eventually become too inelastic. Any shortcomings, however, do not seriously affect the calibration of choice between feasible actions. This is due to conditioning on human moves. We restrict ourselves to decisions in which the actual human move is recognized as a feasible option. Having better expectations would lead to a different, presumably larger set of decisions with which to calibrate the choice of a best alternative. However, the current set already consists of 6949 test cases.

The eGD-expectations also play a role in the representation of opportunities as lotteries. Perhaps they indirectly contributed to the poor performance of expected utility maximization? Although possible, this does not seem very likely. In fully scripted runs, eGD-expectations are fairly close to the Walrasian equilibrium prices.

If it's not the eGD-expectations, can the implementation of expected utility maximization be at fault? With maximization of expected utility, proposing a new offer is always preferred to canceling a pending offer with an unfavorable price. This creates an incentive compatibility problem (c.f. footnote 15 on page 150). Hence, one can rightly argue that rational expected utility maximizers would use longer time-frames, that at least cover the period in which they foresee having incentives to postpone canceling their pending offer. But then uncertainty escalates beyond the point where it can be practically managed. Mitigating this self-control problem by using simple rules of thumb instead of a more complicated analysis seems much more plausible




Figure 5.4 – Shares per trader type in the available commodities due to arbitrage behavior (row 1 - stable and 3 - ccw) and non-speculative behavior (rows 2 - stable and 4 - ccw). Each dot represents an allocation at the end of a period. The results of 10+1 (stable) and 16+1 (ccw) periods from 1,000 runs have been plotted. Arbitrage appears to affect the spread, especially in the unstable treatment.

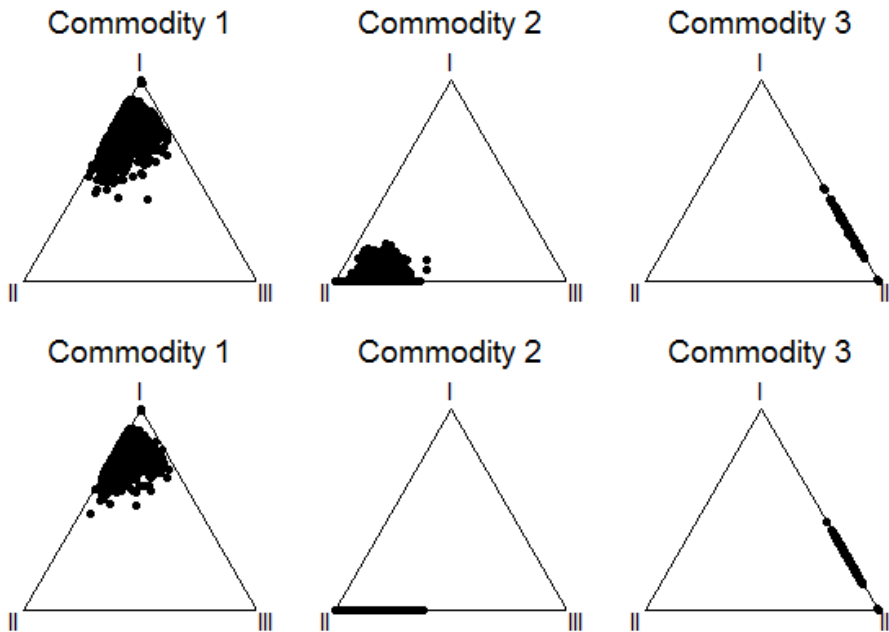


Figure 5.5 – Shares per trader type in the available commodities in the *cw* treatment due to arbitrage behavior (row 1) and non-speculative behavior (row 2). Each dot represents an allocation at the end of a period. The results of 10+1 (*cw*) periods from 1,000 runs have been plotted. Arbitrage contributes to the spread.

and effective. Other researchers have also found that simple rules of thumb, or combinations thereof, succeed in explaining behavior (e.g. Anufriev and Hommes (2012), Tuinstra (1999)).

In an environment that changes while one is meditating a course of action, there may not be enough time to analyze opportunities to the fullest extent. And even if there were time, an adequate analysis may prove to be too difficult. If snap assessments of opportunities and probabilities have to replace more careful estimates, then one has to wonder how people learn to improve themselves. Often there is insufficient feedback, because in each subsequent iteration only one of many possibilities is revealed. Learning rules of thumb by trial and error is far easier. Perhaps this is the preferred form in which people learn.²⁶ Tuinstra (1999) draws a similar conclusion.

Prospect theory gives better predictions of human moves than does expected utility maximization, especially in the unstable treatment. Our results show that this is due to replacing probabilities by decision weights. As a straightforward generalization of expected utility theory, entropy-sensitive preferences (ESP) constitute a simpler alternative to prospect theory. Instead of mapping estimated probabilities to decision weights, ESP leaves them as they are and focuses on the trade-off between expected value and complexity. It is interesting, therefore, to compare the relative

²⁶This is a major theme in the work of the psychologist Gerd Gigerenzer, c.f. Gigerenzer et al. (1999); Gigerenzer and Selten (2002).

merits of prospect theory and entropy-sensitive preferences in predicting human behavior. Here, the evidence is mixed. While ESP give better predictions in the stable economy, prospect theory clearly does better in the unstable treatment. If entropy-sensitive preferences play a role in learning which rules of thumb are best then that may explain why people prefer to propose regular offers over strategic offers.

Speculative behavior consists of recognizing and seizing seemingly profitable trades; but also of the maintenance of the arbitrage portfolio. Apparently there is more to recognizing arbitrage opportunities than observing an acceptable margin, because the recognition rate is below par. And while the prediction of arbitrage moves is also poor, our model of arbitrage behavior does give a slightly better explanation of human trading behavior than a model of non-speculative behavior. This means that having a maximal amount of restrictions on trading behavior is compatible with the data. Calibration of the granularity of mental accounts (and hence of myopia) should be expected to lead to better results.

5.6.2 Disequilibrium theory

The behavior of traders can be analyzed as perceiving feasible actions and, in case of multiple alternatives, selecting a preferred alternative. Although traders appear to further their own interest, we find that they do not apply expected utility maximization for selecting a best feasible action. Moreover, it seems that rules of thumb for prioritizing feasible actions may even lead to better results than expected utility maximization. The reason is that it is difficult to frame choice between perceived opportunities as an expected utility maximization problem. The hard part is the consistent valuation of different opportunities in an environment that is essentially uncertain, and also choosing a time horizon for decisions that removes any potential incentive compatibility problems.

Although speculative behavior can contribute to volatility, it is often considered to be an equilibrating force. It is thought to decrease the spread of transaction prices, and to increase liquidity which leads to higher efficiency. In the experiments of Anderson et al. speculative behavior plays a major role in the training periods only. We have considered arbitrage behavior, in which traders seek to make a profit on each individual arbitrage. To the extent that it plays a role after the practice periods, we find that arbitrage slightly improves convergence and also the direction of orbiting.

Non-speculative behavior seems to be a reliable driver of convergence in allocation. If traders buy what they need and sell what they can spare they are actively eliminating available Pareto-improvements and increasing efficiency.

5.7 Conclusions

In this chapter we have explained why we consider human choice to be fallible. We have calibrated the choice from sets of perceived opportunities. For each human move, that is recognized by robot traders as a feasible action, we often have a simulated context of alternative actions that are also feasible. This allows us to investigate how people choose. We compare rules of thumb, expected utility maximization, cumulative prospect theory and entropy-sensitive preferences. We find that rules of

thumb, rather than expected utility maximization, predict human actions best. It is unlikely that this finding is due to eGD-expectations being inadequate. The rules of thumb for prioritizing feasible actions and the representation of opportunities as lotteries are detailed in appendix B.

Entropy-sensitive utility functions generalize von Neuman-Morgenstern utility functions. The latter rank alternatives based on expected value. Consequences of feasible actions, however, have other attributes as well and we argue that people not only look at expected value, but that they are also sensitive to the lack of predictability of outcomes. This can explain different paradoxes of choice. We show how entropy-sensitive preference can be embedded in the consequentialist framework of choice theory, and we demonstrate that there exists a calibration that explains human choices better than does expected utility maximization.

Another way to induce fallibility in choices is to restrict the set of perceived opportunities. If traders engage in arbitrage and if they seek profits in individual arbitrages then this can induce myopic behavior. We appeal to the theory of mental accounting for modeling arbitrage behavior along this line. For assessing the impact of arbitrage we use a heterogenous population that consists of seven eGD- and eight eEMA-traders. Here we find that arbitrage slightly improves the concentration of trading prices and also the direction of orbiting (but not by enough). Arbitrage contributes to a greater spread of the end-of-period allocations.

Interestingly, heterogeneity of price expectations (without arbitrage behavior) also improves the concentration of trading prices. In the next chapter, robot traders are given the choice to switch between algorithms based on their relative merits. This introduces more heterogeneity and it gives us the opportunity assess the robustness of previous calibration results.