



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 1

Introduction

In a Walrasian equilibrium the plans of economic agents are both individually optimal at the given prices and mutually consistent, i.e. in all markets demand equals supply. In order to achieve such a state, trading has to be postponed until the equilibrium prices have been determined, because trading at non-equilibrium prices would distort the equilibrium. If someone pays too much for a commodity, or sells too cheap, then he loses wealth to another trader and that will typically shift the equilibrium.

The so-called tâtonnement price adjustment process postpones trading. Its price adjustments are straightforward: the price of a commodity is raised if demand exceeds supply; if there is excess supply then the price of the commodity is lowered. In the middle of the twentieth century, economists had high hopes that the tâtonnement process would always converge to the Walrasian equilibrium (provided it exists). If true, that would have provided a highly stylized but essential explanation of how the Walrasian equilibrium can be achieved.

Herbert Scarf, however, has demonstrated that tâtonnement may fail to converge to the Walrasian equilibrium. Scarf (1960) provides three examples consisting of a small exchange economy with three traders and three commodities. In the stable version, tâtonnement does converge to the Walrasian equilibrium. The other two examples are unstable; here, tâtonnement leads to prices orbiting around the Walrasian equilibrium values, in a clockwise or counter-clockwise direction. The three examples differ in the initial allocation only.

Algorithms that always converge to a Walrasian equilibrium (in any economy, from every starting point) are not realistic because they have to postpone trading. Processes with trading at non-equilibrium (or "all" or "false") prices, on the other hand, do not necessarily converge to the Walrasian equilibrium. This leaves us in a situation in which general equilibrium theory can state conditions for the existence, uniqueness and optimality of the Walrasian equilibrium, but it cannot satisfactorily explain how this equilibrium can be achieved.

The mainstream response to this lacuna in economic theory has been to downplay the importance of how an economy evolves when it is out of equilibrium. Samuelson considered this dynamics to be of secondary importance (c.f. Samuelson (1947)) while Friedman even declared the subject to be "a waste of time".¹ Others, like Hahn and

¹This view was offered in a private communication to Fisher, c.f. Fisher (2011).

Fisher, have argued the relevance of studying economies that are out of equilibrium and they have made significant contributions, c.f. Hahn and Negishi (1962), Fisher (1983). Looking back at the development of stability theory, Fisher concludes that "the search for stability (of the Walrasian equilibrium) at great levels of generality is probably a hopeless one", c.f. Fisher (2011, p. 43).

Experimental economists, on the other hand, claim that it takes only a few, uninformed, traders and a Continuous Double Auction (CDA) to obtain results that are approximately equal to the Walrasian equilibrium (e.g. Friedman and Rust (1993), Smith (2008), Anderson et al. (2004)).² This is remarkable because in order to prove the existence of a Walrasian equilibrium and its stability one has to make strong assumptions that go well beyond "a few, uninformed, traders". Furthermore, decision-making in the standard general equilibrium model and in a CDA are fundamentally different.³

The early experiments with CDAs have been conducted in an overly simple setting: in a single financial market, in which traders are told to execute specific limit orders (c.f. Smith (1962); Rust et al. (1993)). They can make a profit by selling above or buying below the limit price. This limit price is a reservation price, that anchors and greatly simplifies expectation formation. Furthermore, with only one market, the equilibrium is artificially prevented from shifting. Anderson et al. (2004) has taken this line of research to a higher level by letting participants trade in the famous examples of Scarf (1960). Their results are impressive: trading by human subjects closely approximates the Walrasian equilibrium in the stable example; in the unstable examples there is orbiting of prices in the direction that is predicted by tâtonnement theory. The latter appears to be robust since it has been replicated by Hirota et al. (2005) and Goeree and Lindsay (2016). This suggests that we may learn something profound about equilibrium discovery by studying experimental price formation.

This thesis aims to contribute to our understanding of equilibrium discovery by trying to replicate experimental trading with the help of algorithms. We want to know how human traders behave and how this affects equilibrium discovery. How do they propose prices? Which opportunities do they perceive? If traders recognize alternative actions, how do they select a preferred option? Are their strategies ecologically rational, i.e. do they survive competition with alternative strategies?

An important goal of this thesis is to find behavioral explanations that acknowledge disequilibrium and also the fallibility of human choice. In stability theory, traders typically behave as if they can trade quantities as they desire, but this presupposes that the economy already is in equilibrium. Furthermore, stability theory features many price adjustment processes that appeal to aggregate excess demand, i.e. to the difference between aggregate demand and supply. Such processes do not offer

²In a CDA both buyers and sellers can propose offers to exchange a certain quantity of a commodity at a specified price.

³In the terminology of Savage (1954), the general equilibrium model is a "small world", because agents are assumed to know all states of Nature and able to attach probabilities to them. The smallest example of an economy with trading at all prices, on the other hand, is a "large world", because here traders face essential uncertainty. They cannot differentiate between states of Nature, let alone attach probabilities to them. And, they do not know the (equilibrium) prices; instead agents are bound to make "mistakes" by selling and buying at prices that with hindsight prove to be unfavorable.

a behavioral explanation, because individual economic agents cannot observe aggregate excess demand. Finally, general equilibrium theory also assumes that individual agents possess more information than is plausible, and that they can process information seamlessly. But people make mistakes and they sometimes suffer from persistent biases that are not easily eradicated by learning nor by competition.⁴ We believe that fallibility matters.

The experiments of Anderson et al. are of great interest for our purpose because (i) trading at all prices in a CDA is sufficiently realistic; (ii) since there are two markets, trading at all prices can shift the stable state away from the Walrasian equilibrium; (iii) convergence of human trading to the Walrasian equilibrium is contingent on the initial allocation; and (iv) orbits (if any) provide an additional way of discriminating between rival explanations. Prof. Anderson has kindly provided the data of two sessions (that apply the stable and the counter clockwise treatment).

As part of this thesis, we have developed a simulation platform, called FACTS (short for Fallible Agents' Commodity Trading System). FACTS can import experimental data, allowing robot traders to virtually participate in a laboratory experiment with human subjects. Each robot trader is matched to a particular human subject and each has to predict the moves of its human alter ego, based on the same preferences, endowments and on the same public information. The data of Anderson et al. provide us with 9799 individual decisions for replication. In addition, FACTS can let robot traders interact freely, subject to the market protocol. This allows us to study the extent to which the aggregate results of robot and human trading are similar. Here, in particular we focus on the predictions of tâtonnement theory, i.e. that prices will converge in the stable Scarf economy and will orbit in the unstable examples, clockwise or counter clockwise depending on the initial allocation.⁵

Chapter 2 reviews different theories of price formation. For this, we largely draw on stability theory because our research takes place in the context of a general equilibrium model. Stability theory typically has consumers and producers making comprehensive plans that cover all markets (i.e. all goods, locations, dates and contingencies). In order to keep decision-making manageable, it is assumed that decision-makers act as price-takers or that they have very good expectations. Furthermore, price formation is often assumed to be driven by aggregate excess demand. We consider comprehensive plans, price taking and a crucial role for aggregate excess demand to be unrealistic and unsuitable for our purpose. We are interested in the discovery of stable states through trading at all prices. Therefore we interpret choice as sequential rather than comprehensive. Absent an auctioneer, traders have to propose

⁴The conjecture of Milton Friedman, that the market will drive out irrational behavior, has been shown to be too naive, because it underestimates the complexities of heterogeneity, c.f. De Long et al. (1990, 1991). Friedman (1953) puts the more sophisticated argument of Alchian (1950) on its head. This contends that economists should care less about intentions and the rationality of agents, because it is the environment which decides the success or failure of certain solutions / types of behavior. Becker (1962) used the idea of constrained, but otherwise random behavior to model impulsive households and to show that the most important economic insights indeed can be derived without invoking rationality. Here lies the origin of the so-called Zero Intelligence methodology (c.f. chapter 4 and appendix B).

⁵Since emergent aggregate phenomena, such as convergence and orbiting, are very sensitive to the specification of behavioral hypotheses at the micro level this provides an additional way to discriminate between rival hypotheses.

prices instead of taking them as given. And crucially, while making sequential choices and proposing prices, traders have to manage with whatever information is publicly available. We will therefore approach the subject matter of stability theory from the perspective of experimental economics and agent-based modeling.

In chapter 3 we discuss the experiments of Anderson et al. in greater detail. Here we also describe the data we have received from Prof. Anderson and use them to derive some stylized facts that characterize human trading. Among other things, these suggest that traders use their current estimates of expected prices as a reservation price. This is reminiscent of price taking, albeit in a more active form, because traders face incomplete and false signals.

Chapter 4 introduces FACTS, our simulation platform. Our calibration methodology proposes multiple criteria for assessing algorithms and defines how to implement them. We calibrate FACTS by considering how well different robot traders predict the moves of human subjects and by comparing the aggregate features of robot and human price formation. The main topic of this chapter is to examine how human traders propose prices. Do they set prices by maximizing expected utility against beliefs that a proposed price will be accepted? Or do they use reservation prices based on expected prices? We find that the subjects of Anderson et al. did not behave like price setters. Instead, the calibration confirms the stylized fact that traders use reservation prices anchored by expected prices. With respect to expectation formation we find that different criteria select different algorithms as being the best. The algorithm that best predicts human actions, eBAS, derives price expectations from bid / ask spreads. Algorithms that estimate so-called "no arbitrage" prices generate robust convergence. The ZIP-algorithm of Cliff (1997) causes prices to orbit systematically in the unstable Scarf economies in the direction that is predicted by tâtonnement theory.⁶ A meditated choice leads us to prefer the so-called eGD-algorithm that derives "no arbitrage" prices from Gjerstad-Dickhaut beliefs, c.f. Gjerstad and Dickhaut (1998): (i) algorithms that perform well in one-step-ahead predictions but fail to achieve convergence in the stable Scarf economy ignore an essential part of human behavior; (ii) the eBAS-algorithm is overly sensitive to haggling;⁷ (iii) theoretically, the notion of "no arbitrage" prices provides the best basis for price expectations; (iv) the eGD-algorithm can be improved whereas the eBAS-algorithm cannot.

Chapter 5 argues that human choice must be understood as fallible: (i) reflections on choice; (ii) an analysis of how traders choose a best action from sets of perceived opportunities; and, (iii) a model of arbitrage suggest that choice is different from how it is presented in text books. For ranking alternative feasible actions we consider expected utility maximization, cumulative prospect theory (c.f. Tversky and Kahneman (1992)), entropy-sensitive preferences (ESP) and simple rules of thumb for prioritizing feasible actions.⁸ We find that the latter predict human trading behavior

⁶In tâtonnement theory orbiting is an expression of a negative feedback mechanism; our simulations, on the other hand, suggest that orbiting of experimental prices is mainly due to a lack of feedback as a result of which one price increases beyond bounds. In experimental trading with humans, there is a similar increase in one price, although here that increase eventually is reversed.

⁷The success of eBAS, in terms of one-step-ahead predictions, is due to its flexibility and not to it capturing some essential characteristic of expectation formation.

⁸ESP introduce a trade-off between expected value and uncertainty. This trade-off provides a simple explanation of paradoxes of choice. We show that ESP fits the framework of choice theory

best. We model arbitrage by applying the theory of mental accounting, c.f. Thaler (1999); Das et al. (2010). Trying to realize a profit with each separate speculative action induces myopia. The overall prediction of human behavior slightly improves as a result of admitting arbitrage behavior based on mental accounting.

Chapter 6 tests the robustness of the results of selected calibrations by giving traders the opportunity to learn which strategy works best for them. Our approach to learning is a mixture of replicator dynamics and reinforcement learning. Here, the main results are that (i) monopolistic competition is strongly dominated in the stable and counter clockwise treatments by reservation prices based on expected prices; (ii) rules of thumb for prioritizing feasible actions strongly dominate other methods of selecting an alternative from a set of perceived opportunities; and, unexpectedly, (iii) ZIP is ecologically rational for the formation of price expectations.

Appendix A addresses market failure. This occurs if available Pareto improvements cannot be implemented through trading. Market failure can be due to trader behavior or to institutional constraints (e.g. that commodities have to be exchanged for money). Our initial simulations quickly ran into market failure. We analyze the likelihood of observing market failure with robot trading and find it to be high. This seems due to the initialization of price expectations (randomly selected from the price simplex) and to quantity setting based on expected utility maximization.

In appendix B, we provide more details with respect to FACTS. We explain how robot traders perceive opportunities for action and how these can be represented as lotteries. Here we also derive the rules of thumb for prioritizing actions. Furthermore, this appendix describes the algorithms for learning prices. Drawing on the literature, we use Zero Intelligence (ZI-) traders, adaptive expectations (eEMA), ZI-plus (ZIP) and Adaptive-Aggressive (AA-) traders, and agents endowed with so-called Gjerstad-Dickhaut (GD-) beliefs. These algorithms were developed for the simpler context of a single financial market in which traders execute limit orders. Where necessary, this thesis applies some changes to the algorithms in order to make them applicable to the context of a general equilibrium model, or to improve them.⁹ We have also added some variations (e.g. eRnd, GDW) and new algorithms (e.g. eBAS, eGD, eME, TU, MEA and MEW).

Appendix C describes a price adjustment process that was part of the development of FACTS. This process is based on an approximation of the aggregate excess demand function. Here, the auctioneer *assumes* that each trader has preferences that can be described by a Cobb-Douglas utility function. A trader's response to previously quoted prices suffices to identify these hypothetical preferences. The unique equilibrium prices of the associated Cobb-Douglas economy feed into the next iteration. We prove global convergence for CES economies in which traders have utility functions ranging from Leontief to Cobb-Douglas utility functions.

Looking back, we like to offer some reflections on our main results. Clower (1955) and Arrow (1959) argue that monopolistic competition is the appropriate way for understanding disequilibrium behavior. More recently, Schinkel (2001) ascribed to

after the axiom of independence is slightly weakened.

⁹For instance, Gjerstad-Dickhaut beliefs should not be used unconditionally for optimizing expected utility, as proposed by the authors; and our implementation of the AA-algorithm is computationally less demanding than the original as proposed by Vytelingum (2006).

the same position. This, however, is not how the subjects of Anderson et al. behave. In the Scarf economies the expected utility level of, say, buyers can increase if they are prepared to pay a *higher* price, because their expected utility benefits more from an increased likelihood that the offer will be accepted than it loses from the higher price.

We find that our rules of thumb for prioritizing feasible actions explain human trading behavior better than others methods, and also that these rules are ecologically rational. We propose that this makes sense, because learning which rules are best offers better opportunities for improving oneself compared to learning how to better value opportunities and how to better assess probabilities and decision weights.¹⁰

What have we learned with respect to the convergence of experimental prices in the stable Scarf economy?

- To the extent that price expectations are responsible for the convergence it is not sufficient for an algorithm to give a good summary of previously observed prices. Algorithms that derive prices from expected utility maximization and algorithms that manage reservation prices based on a utility target generate meaningful prices that reflect a trader's position. Unfortunately, the performance of these algorithms is not good enough.
- Traders learn from observing prices, but they can also learn a lot from scrutinizing the opportunities that are available to them: chapter 3 demonstrates that sophisticated traders in the unstable Scarf economies can deduce the Walrasian equilibrium prices without even having to trade. However, if traders use rules of thumb for prioritizing feasible actions then this explanation seems less likely.¹¹
- Preferences in the Scarf economies are highly symmetric. As a result, deviations from the equilibrium price in one market do not affect demand in the other market. A localized impact of "false" prices could also offer an explanation for the limited shift of the competitive equilibrium in the stable Scarf economy. However, our simulations suggest otherwise.
- Plott et al. (2013) proposes that prices in the stable Scarf economy converge because trading occurs along a Marshallian path, i.e. as if buyers and sellers have been ordered according to their reservation prices. It is unclear how a Marshallian path can be achieved: traders cannot coordinate on private reservation prices, and this may even not be incentive compatible. By varying the priority of acceptances of pending offers, we do find some support for the idea that trading between those who stand to gain the most improves convergence, but only slightly.

¹⁰The power of rules of thumb is a recurring theme in the work of the psychologist Gerd Gigerenzer. He argues that rules of thumb are not necessarily inferior to optimization, c.f. Gigerenzer et al. (1999). Our result supports this claim.

¹¹Interestingly, it may be possible to test whether human traders look inward by repeating the stable treatment with, say, Cobb-Douglas utility functions. If traders in that case have more difficulty in achieving convergence to the Walrasian equilibrium, then the subjects of Anderson et al. apparently have learned something from the relations between fixed complements.

- Marshall (1961) suggests another mechanism that seems more promising. According to Marshall, traders first determine how much they want to spend on each commodity, given (expected) prices. Then they determine proposed prices that take their previous transactions into account. If someone previously has paid too much on average, then he now wants to pay less than the expected price in order to compensate. This kind of behavior feeds corrections back into the market where the "mistake" was made (which is distinct from utility targets and monopolistic competition).

Perhaps the most important result of the calibration of expectation formation is an understanding of how the performance of the algorithms can be improved. Currently, robot traders do not generate enough transactions, and beliefs underlying the "no arbitrage" price expectations over time become too insensitive to new information. These two issues are related and they can be remedied, c.f. section 4.5. Algorithms that manage reservation prices based on a utility target are expected to benefit more than others from robot trading generating more transactions. This is due to the fact that human trading behavior is largely non-speculative; traders buy what they need and sell what they can spare. Having more transactions implies that traders learn that they can achieve higher levels of utility.