# Understanding Financial Crises:
# The Contribution of Experimental Economics

24.7.2012, forthcoming in the *Annals of Economics and Statistics*

Frank Heinemann[1]

Abstract

Several phases of financial crises contain strategic elements than can be directly tested by laboratory experiments. This paper summarizes what we can learn from experiments on the formation of bubbles, on herding behavior, bank runs, and on the effects of providing public information in environments with strategic complementarities. I will put a special emphasis on coordination games that are widely used to model bank runs, government debt crises, and currency crises.

Plusieurs phases de crises financières comportent des éléments stratégiques qui peuvent être directement testés par des expériences de laboratoire. Cet article résume ce que nous pouvons apprendre de ces expériences sur la formation des bulles, les comportements grégaires, les paniques bancaires et les effets de la diffusion d'information publique dans des environnements avec complémentarités stratégiques. L'article se concentre tout particulièrement sur les jeux de coordination qui sont largement utilisés pour modéliser les paniques bancaires, les crises de dette souveraine et les crises de change.

## 1. Introduction

In this paper, I summarize some of the experimental evidence that contributes to our understanding of financial crises. The patterns of financial crises are remarkably predictable. Minsky (1972) has described these patterns by phases that have later been refined by Kindleberger and Aliber (2005) and Aschinger (2001): I will start by summarizing these

---

[1] This paper is based on a talk given at the 2010 ADRES conference in Lyon. Author's address: TU Berlin, Chair of Macroeconomics (H 52), Strasse des 17. Juni, 135, 10623 Berlin, Germany, f.heinemann@ww.tu-berlin.de

phases and emphasizing, which aspects of financial crises have actually been analyzed by laboratory experiments.

1. *Trigger*: The chain of events leading to an eventual crisis is started by an exogenous event like a new technology, financial market innovation, or change in regulation. For example, financial market liberalization of the 1980s caused overinvestment in several emerging markets and eventually resulted in banking and currency crises in those countries. In the 19th century, overinvestment in the new railway industry caused several banking crises in particular in the United States. The dot.com bubble of the late 1990s is still in most people's memory. The subprime crisis was triggered by a combination of governmental home-owner programs, the upcoming instrument of collateralized debt obligations, and deregulation of banks in the U.S.

2. *Boom*: The innovations described above open new opportunities for profitable investment. Profits in the respective industries rise and capital is redirected from other industries towards investment in the new technology. Financing these investments requires credit.

3. *Credit expansion*: Banks are transforming short-term deposits into long-term credits. Credit expansion is attained by various instruments. Reducing downpayment or equity ratio in financing real estate or other investment projects allows more voluminous debt contracts. Means that are often observed are adjustments in credit lines granted by banks to portfolio holders. Usually, one can borrow up to a certain fraction of the portfolio's market value. In boom phases, banks raise this ratio. Mortgage loans in the U.S. sometimes rose to above 100% of a house's market value. Banks justified this by the expected increase in housing prices. In the late 1990s similar developments applied to stock markets, where banks fuelled the dot.com bubble by raising the credit lines of stock holders allowing them to purchase more shares on credit. Credit expansion can be measured by the relation of credit volume to GDP. Before crises, this ratio rises significantly (Kaminsky and Reinhart, 1999).

4. *Destabilizing speculation*: On average, speculation stabilizes prices. But, credit expansion activates funds that ask for investment opportunities. Attracted by rising asset prices, investors buy similar assets on credit to speculate on further price increases. Uninformed players invest in the same industries and follow those whom they perceive to be informed. We call such behavior herding. It leads to asset price bubbles and overinvestment in the respective industries. In experiments, we can simulate price bubbles and analyze under which conditions they arise. We can also analyze herding and whether irrational behavior amplifies or mitigates herding effects.

5. *Crash*: Eventually, profits do not live up to previous expectations, and banks must write off parts of the outstanding debt. The billion-dollar question is, of course, when does a bubble collapse? I will show you one experiment on a model with predictable (and rational) bubbles and crashes.

6. *Reversal of capital flow*: Once investors realize that the party is over, they try to withdraw their funds and reverse the capital flow. The outflow of capital eventually triggers bankruptcies of individual firms or countries or a currency crisis, when the central bank is running out of reserves. Excess supply of the formerly hip assets causes their prices to fall. Eventually, it becomes common knowledge that the bubble is bursting and more investors try to withdraw their funds.

7. *Panic*: The fear of further losses triggers panic sales that cause an even more rapid decline in asset prices. The bubble bursts. Panic sales can be analyzed in laboratory experiments by varying external conditions.

8. *Liquidity squeeze*: Since all market participants want to sell assets, liquidity is scarce. Banks compete for liquidity which drives up money demand and interest rates. Banks in distress may be run by depositors. Lack of trust in a bank's liquidity may trigger a bank run that brings forth the bank's illiquidity.

9. *Liquidity spirals*: To gather the necessary liquidity, banks sell long-term assets. Since many banks are on the same side of the market, asset prices fall and may even fall below their fundamental value that is given by the expected present value of future dividends. The subprime crisis led to such an undervaluation in spring 2009. The decline of asset prices reduces the value of banks. Liquidity spirals lead to contagion vie market prices from illiquid banks forced to sell assets to formerly liquid banks that subsequently face losses on their asset portfolios. The higher the losses, the more depositors withdraw from these banks as well, so that they may also become illiquid in the end. Eventually, the whole banking sector may be threatened with default.

Bank runs, refinancing games, speculative currency attacks, and liquidity spirals have in common that they are coordination games with strategic complementarities and multiple equilibria. Diamond and Dybvig (1983) formalize this aspect in a bank-run model. In this paper, I will provide experimental evidence on such coordination games and explore whether behavior can be predicted, how actual behavior can be described, and how behavior depends on payoff structure, information, and on uninformative extrinsic events.

## 2. Bubbles and Crashes: Rational Behavior?

One of the questions investigated by experimental economics is under which conditions bubbles can occur. The theoretical literature distinguishes rational and irrational bubbles.[2] In the built up of financial crises both types of bubbles may play a role. The first experiment on bubbles has been designed by Smith, Suchaneck, and Williams (1988): they put subjects in an environment, in which they can repeatedly trade an asset with an exogenously given fundamental value that is decreasing from period to period. The model has a unique equilibrium in which the asset price equals the fundamental and agents should not trade. In the experiment, however, subjects trade at prices rising quickly above the fundamental value in early periods and below the fundamental in later periods. The experiment has been replicated several times with varying conditions leading to similar result. This may be taken as evidence that irrational bubbles can occur.

Note that a subject can make an arbitrage profit by selling her or his own units of the asset when the price is above the fundamental and buying when the price is below. Thus, Dufwenberg, Lindqvist, and Moore (2005) invited subjects who had previously participated in the same experiment. If at least 1/3 of subjects are experienced in the sense that they participated in the experiment before, bubbles do not arise. Experienced traders seem to recognize opportunities of arbitrage and drive the price towards the fundamental. We may conclude that irrational bubbles cannot occur repeatedly on the same market and under the same circumstances. This fits very well to the observation that the trigger of a crisis is different each time a bubble builds up – and each time those who speculate on the bubble say "this time is different."[3] Hence, we should not expect a dot.com bubble to reappear, neither a bubble on tulips or railway companies, unless they come up with a very new business model that confuses market participants to such a degree that they do not recognize the similarity with previous bubbles. Whenever sufficiently many traders recognize a situation as similar to a previous bubble, they will speculate against it and thereby prevent it from becoming large enough to cause a crisis.

While this gives us some insight, under which circumstances irrational bubbles can be ruled out, the positive questions, what causes bubbles to arise and how large can they get, are more difficult to answer. The size of bubbles seems to be affected by uncertainty about fundamentals and by symmetry of information about fundamentals. Sutter, Huber, and

[2] Rational bubbles are equilibria in which assets are priced above their fundamental value, while irrational bubbles are non-equilibrium phenomena. Tirole (1985) analyses conditions for existence of rational bubbles.

[3] The recent book by Reinhart and Rogoff (2009) emphasises the different justifications of historical bubbles.

Kirchler (2012) build up on Smith et al. (1988) and compare treatments with symmetrically informed subjects to treatments in which some subjects receive more precise information than others. In treatments with asymmetric information, bubbles are significantly smaller than in treatments with symmetric information. The authors argue that this effect may be due to uninformed traders acting more cautiously to avoid exploitation by better informed traders, who reveal their information over time by affecting the price through their orders. However, it is not clear, why less cautious trading behavior should lead to a systematic *positive* bias between asset price and fundamental value.

Abreu and Brunnermeier (2002) provide a model of a rational bubble. Each subject possesses an asset and may sell this asset (to the experimenter) any time until the game is ending. If a player holds the asset until the end of the game, it yields a commonly known payoff (fundamental value). Before that, subjects can sell the asset to the experimenter at a price that is rising over time and eventually rises above the fundamental value, resembling a bubble. However, as soon as a certain fraction of subjects has sold the asset, the bubble bursts and remaining asset holders receive the fundamental value only. The payoff structure is illustrated in Figure 1. With perfect information, the unique equilibrium prescribes that all subjects should sell their assets at the first instance, when the price rises above the fundamental. However, the authors also introduce uncertainty about the fundamental value and about the time when the game ends. Here, the unique equilibrium prescribes holding the asset for some time after receiving a signal that indicates that the price is already above the fundamental. Reason is the lack of common knowledge. Traders compare the prospective gains from selling (which are small in the beginning) with the probability that the bubble continues to rise (which is large in the beginning). Only some time after the price is above fundamentals, the probability that others keep the asset and prevent the bubble from bursting is sufficiently small to be outweighed by the immediate gains from selling the asset.

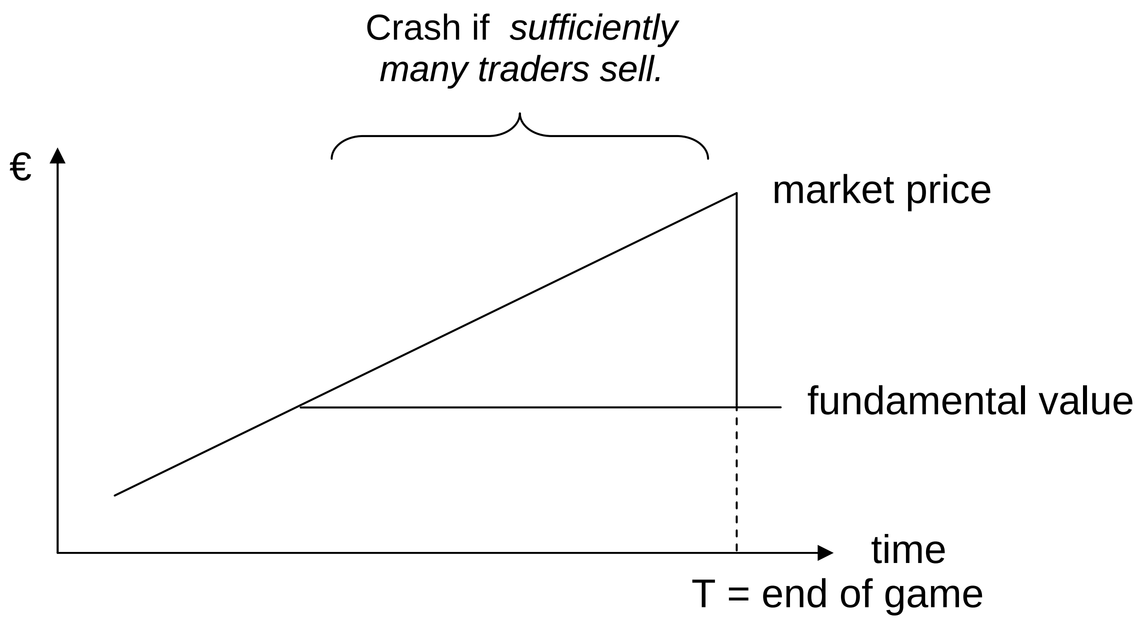


Figure 1. Payoff Structure in Abreu and Brunnermeier (2002)

Brunnermeier and Morgan (2010) report an experiment on this model. In sessions with perfect information, bubbles crash soon after the market price exceeds the fundamental value as predicted by theory. With rising uncertainty and asymmetric information about the fundamental value and the closing date, bubbles tend to persist longer and subjects hold on to the asset often beyond the point in time, where they can be sure that the price is above the fundamental. If subjects can observe the moves of other subjects, the first sale triggers a panic sale by the others. All of this is in line with theoretical predictions. We can conclude that uncertainty about fundamentals may lead to the occurrence of rational bubbles and crashes. A lack of common knowledge may be sufficient to sustain a bubble, even when most or even all traders are aware that there is a bubble. The data basis of this experiment is small and the question what may cause bubbles and crashes deserves a lot more investigation, both from theory and experimental economics.

## 3. Herding: limited levels of reasoning?

Herding is a frequent phenomenon on financial markets. Neoclassical theory postulates efficient markets, in which case a single small investor does not need much information to decide on his portfolio. If markets are strong-form efficient, all private information should be reflected in the price, and the best strategy is to hold a market portfolio. Hence, if others buy shares of internet companies (and drive up their prices), I should conclude that there is information in the market pointing to higher expected returns of these companies. Hence, I should follow the crowd and keep these stocks in my portfolio. In an experiment by Anderson and Holt (1997), subjects had to bet on either of two events that were (a priori) equally likely. Before deciding, subjects received private signals indicating the true event with probability 2/3. The first subject naturally bets on the event indicated by her signal. The second subject is informed about the decision of his predecessor and can conclude that she had a signal coinciding with her choice. If his own signal differs from the observed decision of subject 1, the posterior probability for both events is ½. One may assume that subject 2 prefers following his own signal, as he cannot be entirely sure that subject 1 made the rational choice. If both, subjects 1 and 2 make the same choice, then subject 3 can deduce that both received the same signal. Now, even if subject 3 gets the opposite signal, she should rather follow her predecessors, because their two signals outweigh her own signal. The same logic applies to all other subjects: if the first two choices coincide, all subsequent players should follow the same decision irrespective of their own private signals. More generally, in this game it is rational to follow the majority of predecessors and ignore one's own signal, whenever the difference in choices by predecessors is at least 2. Anderson and Holt (1997) confirm the occurrence of rational herding in the experiment: in about 71% of all cases where the difference in actions was at least 2, subsequent decisions followed the majority irrespective of their own signals.

In Kübler and Weizsäcker (2004), subjects can buy signals instead of getting them for free. If rationality prescribes herding, subjects should not buy any signals, as those should not affect their behavior anyway. In the experiment, however, subjects bought signals too often, and there were significant deviations from rational herding that can be explained by limited levels of reasoning: observations are consistent with assuming that subjects process their information in a rational way, but do not fully account for others processing their information rationally and neglect that others also account (to some extend) for the rationality of others. The consequences are that too much money is wasted in acquiring redundant information and the occurrence of herds is less likely than theoretically predicted. In experiments, it is a robust finding that subjects have a lower tendency to follow predecessors than prescribed by theories based on full rationality. The more effort agents put on getting signals about fundamentals and the less they follow their predecessors, the higher is the resulting correlation between fundamentals and assets prices. Thus, bounded-rational behavior may actually stabilize markets compared to fully rational behavior and may even yield a higher average payoff.

In events of bubbles and crises, herding behavior is usually blamed for destabilizing markets and reducing welfare. It is yet an open question under which conditions herding behavior destabilizes a dynamic economy and how the observed deviations from rationality can be exploited to construct mechanisms for providing information that raise the ex-ante expected welfare.

## 4. Bank Runs: Measures of Prevention

The classical theory paper on bank runs is Diamond and Dybvig (1983). Banks transform short-run deposits into credit for long-run investment projects. Long-run projects earn a higher dividend allowing the bank to promise higher returns to depositors than without maturity transformation. Depositors choose whether to withdraw deposits early (t=1) or late (t=2). A late withdrawal yields a higher payoff (interest rate), provided that the bank survives. Survival, however, depends on how many depositors withdraw early. The game has two equilibria. In one of them, all depositors withdraw early: they run the bank. This causes bankruptcy of the bank, who serves depositors in the order of appearance. Depositors are paid their promised return in t=1, provided that they are in front of the queue. Those who are further behind leave empty-handed. Before queuing up, depositors do not know their position in the queue so that they have a positive probability of being served: if a depositor waits until t=2, he gets zero for sure. Thus, the best response to expecting that others are running the bank is to run the bank yourself. In the other equilibrium, depositors withdraw late and earn a higher return than resulting from an early withdrawal. Here, the bank survives, so that no depositor has an incentive to run the bank early. Diamond and Dybvig (1983) show that another bankruptcy rule (suspension of convertibility) or introducing deposit insurance may alter incentives in such a way that the game has a unique equilibrium, in which all depositors leave their money in the bank. The nice feature of these rules is that they are costless in equilibrium. If the bank is not run, bankruptcy will not occur and deposit insurance will not be used. This feature rests on assuming that the bank has a pure liquidity problem, but is fundamentally solvent. In real crises, however, there is a positive probability that long-term projects yield lower returns than expected and the bank may become insolvent even if all depositors leave their funds in the bank.

Schotter and Yorulmazer (2009) test bank runs in a laboratory experiment, in which subjects play depositors of a bank and can decide in which of four periods they want to withdraw their deposits. The bank promises a fixed interest rate for each period. However, the bank's earnings are stochastic and it may be unable to pay all claims. In this case, depositors are paid according to the first-come-first-serve rule, with remaining funds being split among agents who simultaneously withdraw at the time of bankruptcy. Subjects have an incentive to leave their money in a liquid bank, because it earns interest, and to withdraw ahead of others, if they fear that the bank might become illiquid. The probability distribution of bank earnings and potential payoffs are common information. Schotter and Yorulmazer (2009) distinguish treatments with different earning distributions and with different information of depositors. Their main results are:

1. If some depositors have insider information about the bank‘s return, bank runs are less likely.
2. A higher mean of the bank’s earnings reduces the likelihood of bank runs especially if predecessors are observed.

These results indicate that banking crises are not purely self-fulfilling, but rather related to fundamentals. In addition, results indicate that more information reduces the probability of bank runs. This is in line with other experimental evidence (see below) and makes a case for timely disclosure of the risks in a bank’s balance sheet to depositors. While such information inevitably leads to a weak bank’s failure, it helps strong banks to avoid unnecessary liquidity crises. Since weak banks are bound to fail anyway, expected welfare is higher with disclosure. Thus, increasing transparency strengthens the correlation between fundamentals and depositors’ behavior, it reduces the likelihood of inefficient liquidation and thereby raises ex-ante expected welfare.

## 5. Refinancing Debt: a coordination game of strategic complements

The model of Diamond and Dybvig (1983) has shown that banking crises may result from coordination failure among depositors. Similar coordination games can be used to describe refinancing of firms, a collapse of the inter-bank market, currency crises, or a public debt crisis. A simplified bank’s balance sheet contains long-term credits on the active side and short-term liabilities and equity on the passive side as displayed in Figure 2.

| assets | liabilities and equity |
|---|---|
| long-term credits | short-term deposits |
| cash reserves | equity |

Figure 2. Simplified balance sheet of a bank

If many depositors withdraw at the same time (bank run) cash reserves are insufficient and the bank has to sell some of its long-term assets. Such “fire sales” are usually associated with prices below the present value of expected credit returns. This depletes equity and may lead to insolvency. On the other hand, if sufficiently many depositors roll over their claims, the bank can survive, pay off all deposits, and depositors are better off than by running the bank. The game has two equilibria in pure strategies: all withdraw and all roll over. If the bank is fundamentally sound, only the no-run equilibrium is efficient.

An even more dramatic logic applies to the inter-bank market: Banks decide whether or not to lend each other liquidity. If sufficiently many banks lend to each other via the inter-bank market, the banking system is stable and all fundamentally sound banks survive. If, however, banks withdraw from the inter-bank market and do not provide deposits, because they fear that other banks may collapse, banks in need of liquidity must sell some of their long-term assets in exchange for liquidity. If several large banks are selling long-term assets, the market prices of these assets decrease, which reduces the value of the remaining assets. Since balance sheets must evaluate assets with market prices, the decline in market prices depletes equity in all banks, even those that originally had excess liquidity.

Thus, if banks fail to coordinate on lending to each other, market effects lead to contagion that may (in the extreme case) wipe out the whole banking sector. To prevent this, the central bank may step in and act as a clearinghouse replacing the inter-bank market. This was actually done by the western central banks during the recent crisis, in particular in 2008 and 2009. Systemic crises can also be triggered by an unexpected decline in asset prices: suppose the market value of mortgage loans decreases. This reduces banks’ equity and may cause some banks to become insolvent. Others may just want to downsize their balance sheet and restore a proper relation between the value of outstanding credits and the amount of equity. Thus, they need to sell some of their assets, leading to a further decline in asset prices. Cifuentes, Ferruci and Shin (2005) and Brunnermeier and Pedersen (2009) provide more detailed explanations of these contagious liquidity spirals.

Currency crises are events in which foreign exchange markets show an excess supply of a currency that has a fixed exchange rate. In order to keep the exchange rate, the central bank must use reserves denominated in foreign currency to buy up the excess supply of its own currency. If excess supply is sufficiently large, central bank reserves are insufficient for sustaining the exchange rate. Thus, a coordinated attack by large traders may cause a devaluation that would not have occurred without the extra pressure on the foreign exchange market. Again, the underlying game has two equilibria in pure strategies: if all traders sell domestic currency, the currency is devalued and a single trader who holds on to it suffers a loss. However, if all traders keep holding domestic currency, there is no devaluation, and a single trader selling this currency suffers a loss from transaction costs (mainly: lower interest rate for foreign currency).

The final example is a public debt crisis. Borrowers on financial markets and rating agencies decide upon the soundness of a public debtor. If borrowers attribute a high probability to a sovereign default, they will ask for a high interest rate to cover the default risk. If the country is on the edge of being insolvent, higher interest rates will bring it over the edge and trigger the default. If borrowers believe, instead, that the country can service its debt, the interest rate remains low and the country may fulfill these expectations.

All of these crisis situations have been described as coordination games in which players have two possible actions: one is a safe alternative and the other one is risky in the sense that a higher payoff can be achieved, if sufficiently many other players take the same action, but a lower payoff results from too many others taking the safe action.

Consider a game with N players. All players simultaneously decide between A and B. The payoff for A is X. The payoff for B is $Y > X$, if at least a fraction k of the other players choose B. Otherwise, the payoff for choosing B is $Z < X$.

In refinancing games (bank run), action A is to withdraw early and cash in a safe payment X. Action B is to roll over debt with the prospect of higher returns Y, provided that other creditors roll over as well. Otherwise, the return is the final liquidation value Z.

This game has two equilibria in pure strategies: all players choose A and all players choose B. A player's optimal decision depends on expectations about decisions of others. Assuming rationality is not sufficient to determine a unique outcome. For getting any kind of predictions, we need a theory that is capable of dealing with strategic uncertainty. In experiments, we can test such theories and answer empirical questions about the likely behavior of players in a coordination game as the one described above.

For regulation and policy design, we would like to know, how players respond to changing incentives. Multiplicity of equilibria prevents comparative static exercises, while experiments can give us an answer. Transparency is a widely discussed topic, but theoretical results on the effects of transparency do not provide a clear guide to policy recommendations. Field empirics are even less helpful, because there are few reliable data on information and beliefs.

In experiments, we can control the information of subjects and, thus, answer how transparency will affect the outcome.

When there are multiple equilibria, then agents may coordinate on playing A or B, depending on an event that is unrelated to the game. In these so called sunspot equilibria the selected equilibrium depends on an extrinsic event. Imagine that traders attack a currency conditional on a sport result. If the condition is common knowledge, the best response is to condition one's own action on the same event. A recent paper by Fehr, Heinemann, and Llorente-Saguer (2011) shows how sunspot equilibria can be generated in the laboratory.

## 6. Predicting Behavior and Comparative Statics in Coordination Games

Heinemann, Nagel, and Ockenfels (2009) conduct an experiment in which they systematically explore how behavior depends on the relation between safe and risky payoffs and on the minimal fraction of agents who must follow the risky strategy in order to realize the higher payoff.

In this experiment, subjects are assigned to groups of 4, 7 or 10 subjects. Each subject plays 30 coordination games and 10 lotteries in which she decides between A and B. In coordination games, the payoff for choosing A is $X$ Euro, the payoff for B is 15 Euro, if *at least a fraction* $k$ of the other group members decide for B, and 0 Euro otherwise. $X$ varies from 1.50 to 15 Euro (in steps of 1.50) and $k$ = 1/3, 2/3 or 1. Each subject is in one group and plays 30 combinations of $X$ and $k$. With three different group sizes, the experiment yields data for 90 different coordination games. Figure 3 shows a sample screen with 10 binary coordination games. Lottery decisions extract subjects' certainty equivalents for a lottery paying 15 Euro with probability 2/3. At the end of a session, one of the 40 games was randomly selected to pay subjects. The experiment has been conducted at 4 different locations allowing out-of-sample tests.

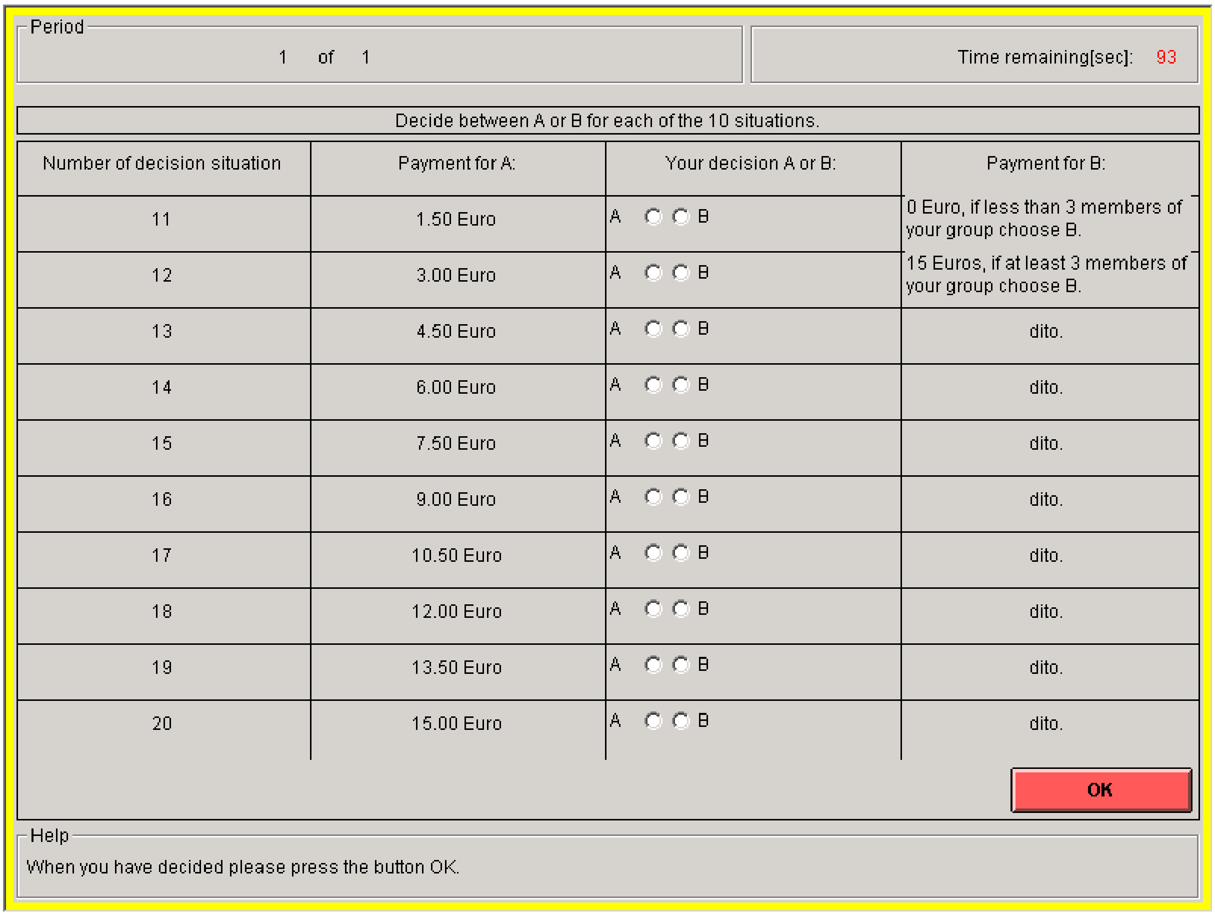


Figure 3: Sample screen.

The results show that the proportion of subjects choosing B decreases in the safe payoff $X$ and in the fraction $k$ of others required for success with B, but does not depend on group size $N$. Thus, comparative statics follows the intuition that risky actions are chosen less likely if their opportunity costs or the coordination requirement for success are rising. Figure 4 displays the proportion of B-choices conditional on $X$ and $k$ for one data set.

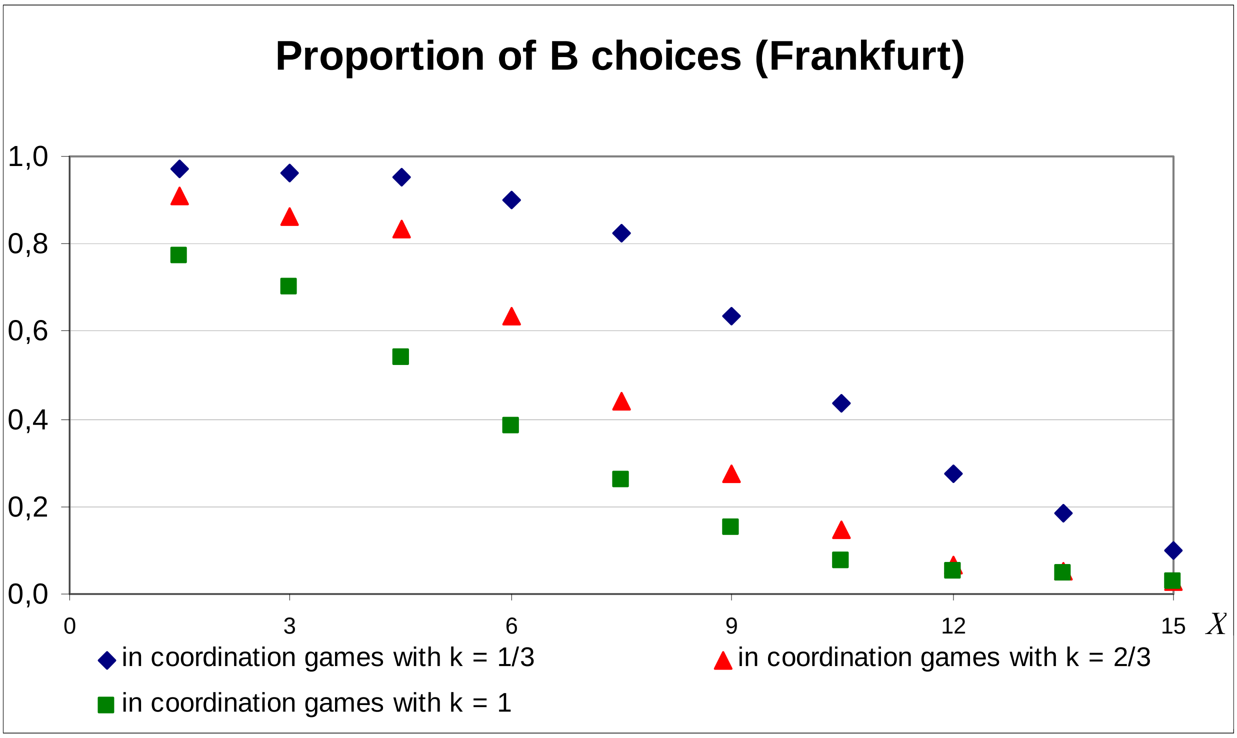


Figure 4: Proportion of B-choices conditional on $X$ and $k$ for data set Frankfurt.

How likely are "crises"? If each player chooses B with probability $p$, the probability that at least $K$ out of $N$ subjects choose B is given by 1–Bin($K$–1, $N$, $p$), where Bin is the cumulative binomial distribution. Table 1 displays the probability that sufficiently many players choose B conditional on $X$, $K$, and $N$ for the data set in Frankfurt[4]. In 58 (64%) of all games, the outcome can be predicted with an error rate of less than 5%. In 44 coordination games (49%), behavior can be predicted across subject pools with an error rate below 5%.

| $N$ | $K$ | $k$ | $X$=1.50 | 3.00 | 4.50 | 6.00 | 7.50 | 9.00 | 10.50 | 12.00 | 13.50 | 15.00 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Frankfurt | | | | | | | | | | | | |
| 4 | 2 | 1/3 | 1.00 | 1.00 | 1.00 | 0.99 | 0.98 | 0.86 | 0.59 | 0.29 | 0.16 | 0.06 |
| 7 | 3 | 1/3 | 1.00 | 1.00 | 1.00 | 1.00 | 0.98 | 0.92 | 0.64 | 0.27 | 0.10 | 0.10 |
| 10 | 4 | 1/3 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 0.94 | 0.73 | 0.49 | 0.22 | 0.12 |
| 4 | 3 | 2/3 | 0.92 | 0.92 | 0.88 | 0.66 | 0.27 | 0.06 | 0.01 | 0.01 | 0.00 | 0.00 |
| 7 | 5 | 2/3 | 0.90 | 0.68 | 0.73 | 0.36 | 0.13 | 0.04 | 0.00 | 0.00 | 0.00 | 0.00 |
| 10 | 7 | 2/3 | 0.95 | 0.95 | 0.88 | 0.22 | 0.04 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| 4 | 4 | 1 | 0.37 | 0.27 | 0.14 | 0.04 | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| 7 | 7 | 1 | 0.09 | 0.04 | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| 10 | 10 | 1 | 0.02 | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |

[4] Note that $k = (K-1) / (N-1)$.

Table 1: Probabilities for success of B. The underlined numbers indicate situations in which success or failure of coordination can be predicted with an error rate of less than 5% in *all* four locations.

To predict behavior in more complicated and applied models, we need a theory of strategic uncertainty. The theory of global games, introduced by Carlsson and Van Damme (1993) and extended by Morris and Shin (2003) can be used for this purpose. The theory of global games embeds the original coordination game in a stochastic environment, in which a random state of the world determines payoffs. Agents get noisy private signals about the state of the world and, hence, beliefs are private information. Using this as a theory of strategic uncertainty postulates that players behave as if payoffs are uncertain and as if all players have private information about payoffs. In equilibrium, beliefs about payoffs and about the likely behavior of other players are stochastic and their distribution is determined by the distribution of private signals. Given some technical requirements, the theory predicts a unique equilibrium with a common threshold, such that players choose B, if their private signals are on one side of the threshold, and A otherwise.

Heinemann et al. (2009) estimate a global game in which the payoffs are assumed to be $15+y$ if at least $K$ players decide for B and $y$ otherwise, where $y$ is a random variable. The true realization is of course $y$ = 0. Subjects receive private signals $y_i$, so that their posterior beliefs about $y$ are normal with mean $y_i$ and variance $\sigma^2$. The estimation procedure reveals estimated values for $\sigma$ around 4 to 8 in the different subject pools. The estimated global game delivers a good fit of the data (see Figure 5), but predicts too strong responses to the fraction $k$ of players needed for success with B.

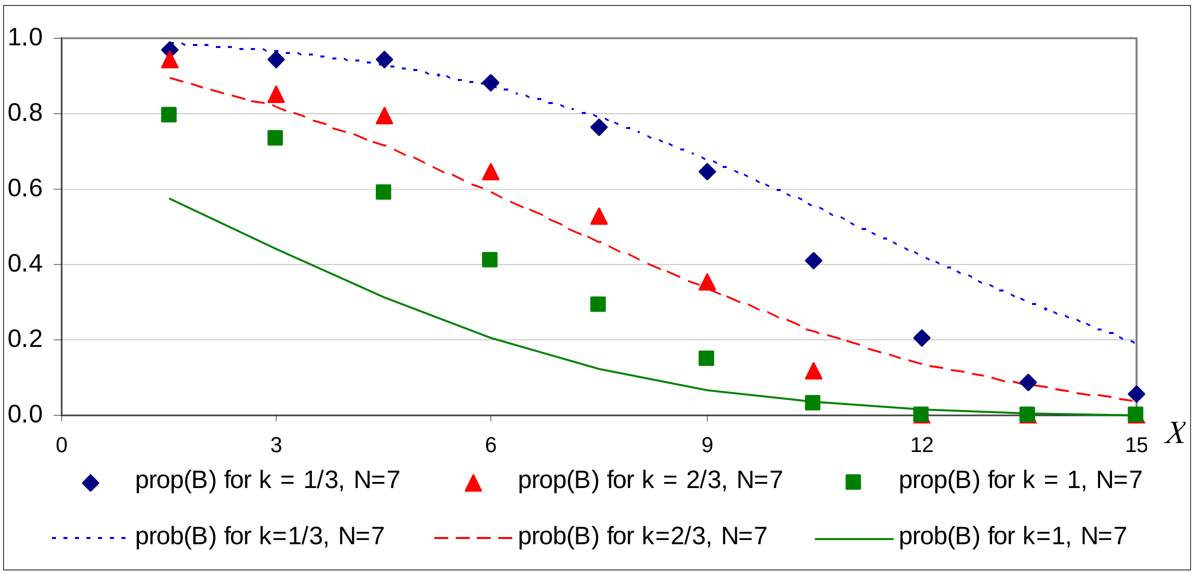

Figure 5: Dots represent the relative frequency of B-choices given $X$ and curves indicate the probability of B-choices in the estimated global-game equilibrium for $N = 7$ (Frankfurt data).

Heinemann et al. (2009) also elicit beliefs by directly asking subjects about their subjective probability of another randomly selected subject choosing B. Stated beliefs about others' behavior are on average correct. Figure 6 shows how close average stated beliefs and average actions are. However, there are vast differences between stated beliefs of different subjects. In particular in situations, in which we would not be able to give a reliable prediction about success or failure of action B, i.e. situations that have success probabilities between .05 and .95 in Table 1, the variance of stated beliefs is large, while stated beliefs vary less in situations, where we can predict the outcome. Figure 7 displays the variance of stated beliefs for all games.

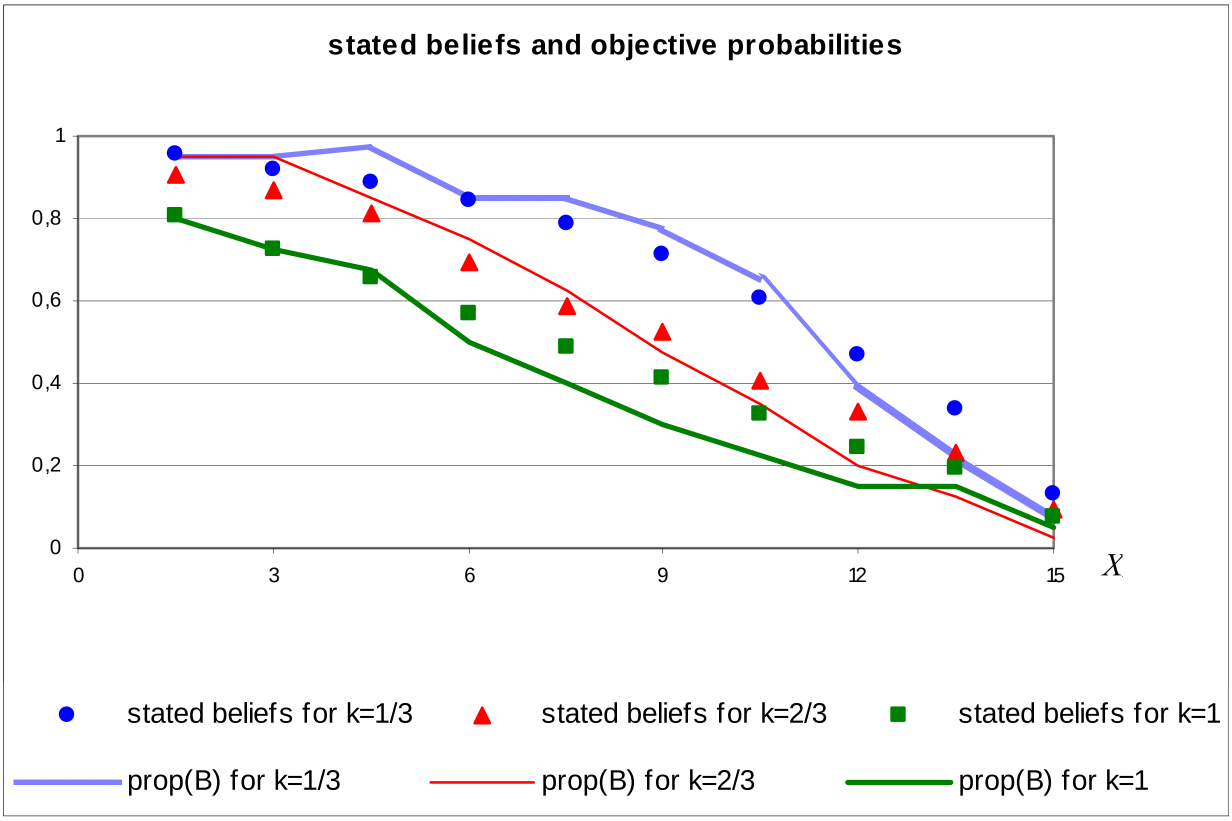


Figure 6: Stated beliefs and relative frequency of B-choices.

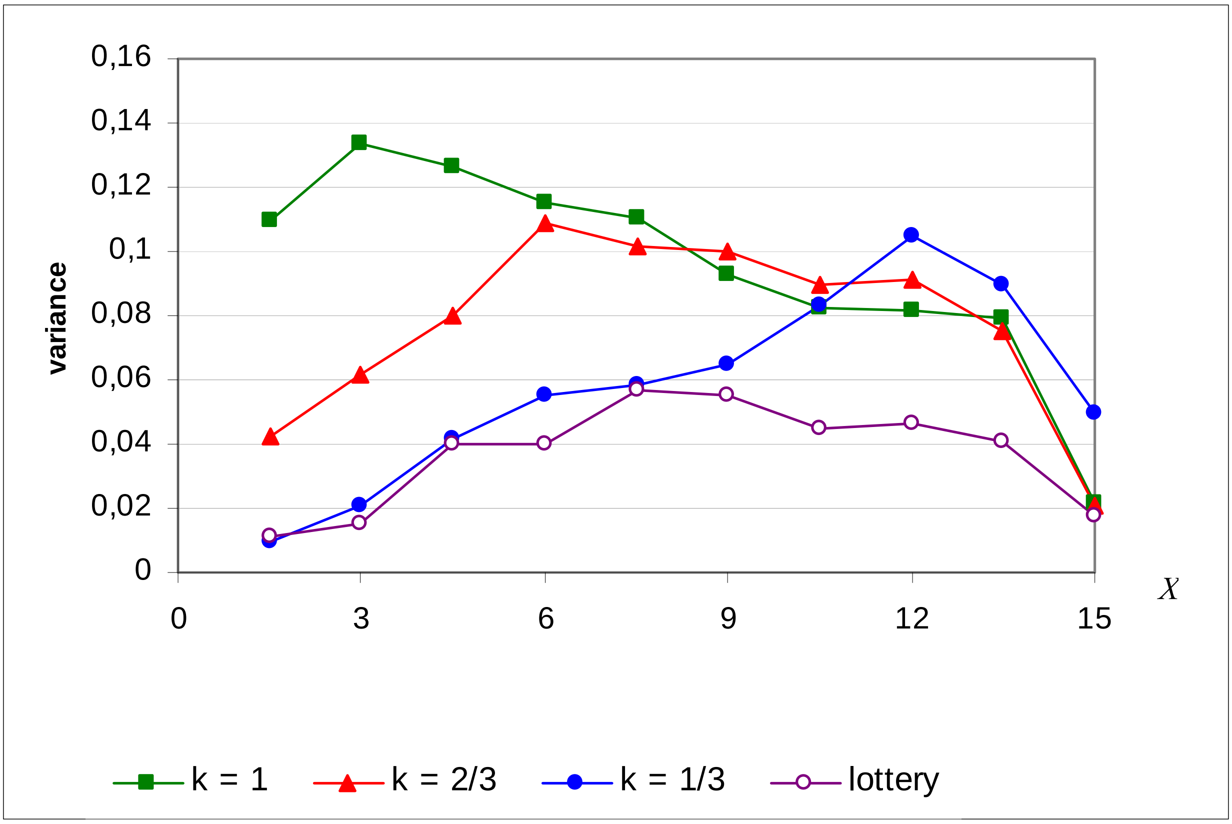


Figure 7: Variance of stated beliefs for each game.

For a player who happens to play coordination games, it is good to have a simple strategy to follow. There are a couple of normative refinement theories for games with multiple equilibria. A refinement selects one out of many equilibria according to general criteria. The payoff dominant equilibrium (PDE) selects B, whenever $X < 15$, the maximin criterion always selects A. Both strategies yield low payoffs when pursued against the distribution of actual choices in the experiment. Hence, they are bad recommendations for an individual player.[5]

Heinemann et al. (2009) compare the expected payoffs arising from four refinements: the global-game selection (GGS), the risk-dominant equilibrium (RDE), the best response to another subject choosing B with probability $p = 2/3$ (P2/3), and the limiting-logit equilibrium (LLE). Payoffs are compared with those arising from a best response to actual behavior and with average payoffs that were achieved by subjects' actual choices. Figure 8 displays the expected payoffs for a subject in Frankfurt following either of these strategies. All four refinements yield higher expected payoffs than subjects actually achieved on average. GGS

[5] The payoff-dominant equilibrium is by definition the best recommendation if it is followed by all players, but it may be a bad recommendation for an individual in a given environment.

and RDE provide the best recommendations also in response to data from the same experiment carried out at other locations.

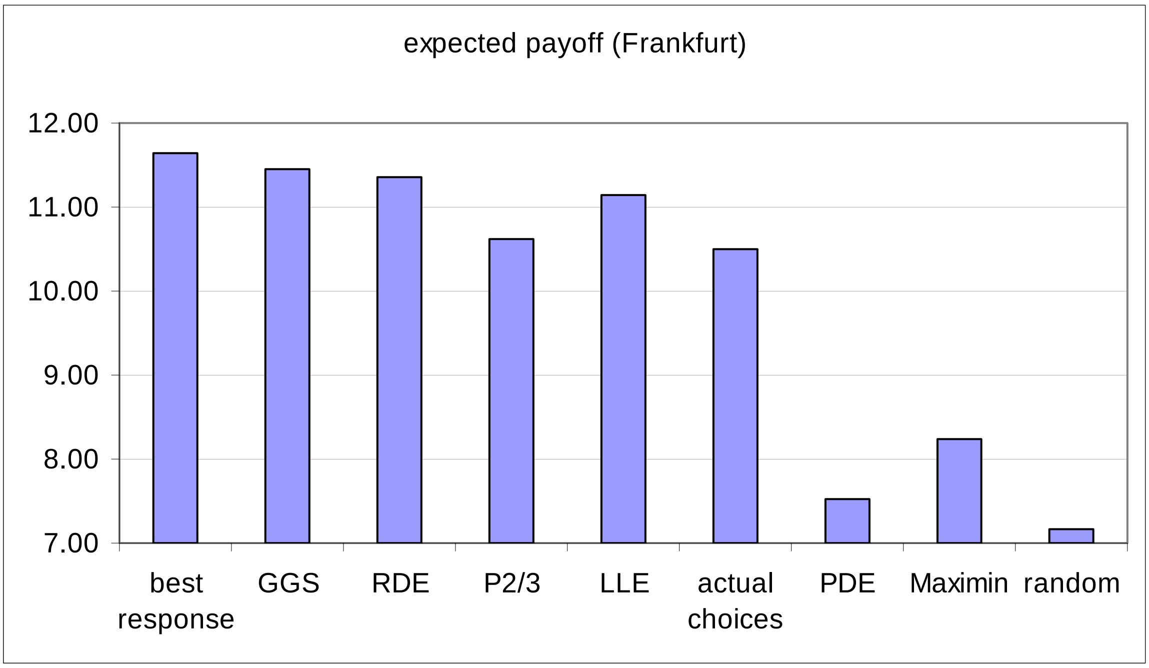


Figure 8: Expected payoffs from different strategies against actual distribution of choices.

A nice feature about the global-game selection is its simplicity. The global-game selection is the limiting equilibrium of a global game with diminishing variance of private signals. In the limit, the marginal player who is indifferent believes that the proportion of other players choosing either action has a uniform distribution. This defines a simple equilibrium condition. If at least $K$ out of $N$ players must chooses B to get $Y$ Euro (0 otherwise) and the safe payoff is $X$, the global game selection implies that a subject should choose B if and only if

$$X < Y\left(1 - \frac{K-1}{N}\right).$$

So, for example, if the group size is $N$=7 and $K$=5 players are needed to get $Y$=15 Euro for choosing B, a player should choose B if and only if the safe payoff for A is lower than $X^*$ = 6.43.

## 7. Managing Information Flow

Transparency is a widely discussed topic, and theoretical results on the effects of transparency in coordination games do not provide a clear guide to policy recommendations. Public information may cause multiple equilibria, almost perfect private information guarantees uniqueness. The comparative statics properties with respect to the precision of signals are mixed and minor in size. Due to the possibility of controlling information, laboratory experiments are the best environment for empirical studies on the effects of transparency.

Heinemann, Nagel, and Ockenfels (2004) present an experiment on a speculative-attack game by Morris and Shin (1998). Subjects are assigned to groups of 15 and simultaneously decide between two options A and B. The payoff for A is a constant $T$, the payoff for B is a random number $Y$ if sufficiently many subjects $a(Y)$ decide for B in the same situation (with $a' < 0$), and zero otherwise. In Morris and Shin (1998), B is a speculative attack, $Y$ is the difference between pegged and shadow exchange rate, and $T$ are transaction costs. The function $a(Y)$ is the amount of pressure on the foreign exchange market at which the central bank gives up defending the exchange rate.

In the experiment, $Y$ is a random number with uniform distribution in [10, 90]. In sessions with common information, subjects know $Y$ when deciding. In sessions with private information, subjects get private signals $X_i \in [Y-10,\ Y+10]$. These signals are independently drawn with a uniform distribution. In each round, subjects see 10 games on one screen and decide for each of them. When all are done, they learn the outcome for each of the games. In the next round, they get 10 new games to play. Subjects coordinate their strategies over time and should settle on an equilibrium. After 8 rounds, the safe payoff $T$ is changed from 20 to 50 or vice versa for another 8 rounds.

Subjects converge to using threshold strategies in which they choose B (attack), if $Y$ is above the threshold, and choose A otherwise. In each group, the individual thresholds converge over time and Heinemann et al. (2004) estimate the group specific thresholds from data of the last 4 rounds.

The game with private information is a global game and has a unique equilibrium. The game with common information has multiple equilibria. In threshold strategies, any threshold from $T$ to the value at which $a(Y) = 1$ is an equilibrium. Heinemann et al. (2004) show that group specific thresholds can be widely explained by the parameters of the payoff function ($Y$, $T$, and $a(Y)$), the order in which $T$ is changed, and the information condition (common versus private). Responses to parameters of the payoff function are as predicted by the theory of global games or the risk-dominant equilibrium.[6] With common information, subjects

[6] These comparative statics properties have also been found by Schmidt et al. (2003) and many other recent experiments on coordination games.

coordinate on lower and more efficient thresholds than with private information. The average values of residuals are about the same for common and private information and there is no indication that public information destabilizes the economy by causing multiple equilibria.

In private information treatments, group specific thresholds are around or below the unique equilibrium. In common information treatments, nearly all thresholds are below the global-game selection and often very close to the payoff-dominant equilibrium that prescribes a threshold of *T*. This result supports the impression from other experiments (e.g. by Van Huyck et al. 1991) that players of a coordination game with strategic complements deviate from the risk-dominant or global-game selection towards more efficient strategies.

Qu (2011) builds up on Heinemann et al. (2004) and compares the private-information treatment as a baseline with two treatments, in which private signals are aggregated before subjects decide between the two actions. In a market treatment, subjects first trade a contingent claim paying conditional on whether action B will be successful or not. The market price of this asset is publicly observable and aggregates private signals. In a second stage of the treatment, subjects can decide between A and B conditional on the asset's market price. This is contrasted with a cheap-talk treatment, in which subjects first announce, whether they intend to choose A or B. The total number of subjects announcing B is publicly revealed before subjects actually decide between A and B. Since announcements are not binding, it is a weakly dominating strategy to announce B independent of one's signal. Hence, announcements should be uninformative. Qu (2011) shows, however, that announcements are correlated with behavior. Both mechanisms, market and cheap talk, improve coordination, but only cheap talk raises efficiency compared with the baseline treatment, because the market mechanism often leads to coordination on the inefficient equilibrium. It is surprising that cheap talk leads to a more efficient outcome than the market. There is no explanation yet, but the result may be robust: for 2-player coordination games, Cooper et al. (1992) have already shown that non-binding pre-play communication can lead to a dramatic improvement in efficiency.

The social value of information has been explored by Cornand and Heinemann (2010), who present an experiment on a game by Morris and Shin (2002) with strategic complements and a unique equilibrium. Subjects are matched in pairs and have a motive to align their actions with some random fundamental, but also have an incentive to coordinate their actions. In particular, subjects face a payoff function $U_i = const. - r\,(a_i - a_j)^2 - (1-r)\,(a_i - Z)^2$, where $a_i$ is the action of the subject, $a_j$ is the action of her partner, and $Z$ is the fundamental state that is randomly drawn with a uniform distribution on [50, 450]. Thus, deviations of players' actions lead to a quadratic loss, as do deviations from the fundamental. The relative importance of coordination is governed by parameter $r$ that was varied in the experiment. Each subject receives two signals about the unknown state, one is a common signal $Y$ that is shown to both players, while the other signal $X_i$ is private. Private signals are drawn independently from each other and from the common signal. All signals have a uniform distribution in $[Z - 10, Z + 10]$. The game has a unique equilibrium, in which players choose $a_i = \frac{1-r}{2-r} X_i + \frac{1}{2-r} Y$. Note that, although both signals are equally precise, the equilibrium weight on the common signal exceeds ½, if $r > 0$, i.e. if coordination matters.

In the experiment, subjects put a larger weight on the common than on the private signal, but the difference is smaller than prescribed by equilibrium. Cornand and Heinemann (2010) argue that behavior can be better explained by limited levels of reasoning: subjects put a larger weight on the common signal, because they realize that it is more informative about the likely action of the other player, but they do not account for other subjects also putting a higher weight on the common signal. By eliciting higher-order beliefs directly, Cornand and Heinemann (2010) find evidence that subjects also underestimate the importance of common information for predicting the likely beliefs of other players. For explaining data, limited levels of reasoning and non-Bayesian updating of higher-order beliefs need to be combined.

These results have interesting implications for the welfare effects of public announcements. Morris and Shin (2002) show that public announcements may reduce expected welfare, because the large weight that players put on common signals may divert their choices from the expected fundamental. In financial markets, we often have situations in which predicting the likely behavior of others raises one's payoff by lowering the payoff of those who do a worse job in predicting others' beliefs. In such environments, coordination affects single players' payoffs, but not welfare; the coordination part is a zero-sum game. Here, public information reduces expected welfare if the precision of public signals is sufficiently low. Cornand and Heinemann (2010) show that for the low weights that actual players attribute to common signals in the experiment, the welfare effects of public announcements are always positive. Thus, bounded rationality does not only reduce the probability of inefficient herding (see above) but also prevents negative welfare effects of transparency.

## 8. The Power of Sunspots

In games with multiple equilibria, agents may condition their actions on events that are unrelated to the game. If all traders believe that a bank will collapse if a sport event brings a certain result, the best response is to run the bank when the result actually occurs. Cass and Shell (1983) show: whenever there are multiple equilibria, there are also sunspot equilibria, in which agents condition their actions on publicly observable but intrinsically uninformative signals. Even though these signals are uninformative they may serve as focal points for agents' beliefs. A common belief that a selected equilibrium depends on some unrelated but commonly observable event (sunspot) is self-fulfilling. The experiments described in the previous section did not have unrelated events in the controlled conditions. To test, whether "sunspots" may affect behavior, we need a design in which an unrelated event is semantically sufficiently strong to raise common beliefs among players that they should condition their actions on the event.

Fehr, Heinemann, and Llorente-Saguer (2011) present an experiment that reliably generates sunspot equilibria. They analyze how the likelihood that unrelated signals affect behavior depends on the structure and distribution of signals. The basic game is a coordination game repeatedly played by groups of 6 subjects.

In each round, each subject is randomly matched with another subject from the same group and chooses a number from 0 to 100. A subject's payoff is higher, the closer his number is to the number chosen by his partner. Let $a_i$ and $a_j$ be the chosen numbers of the two subjects. The payoff for each is given by

$$\pi_i(a_i, a_j) = 200 - \frac{1}{50} \cdot \left(a_i - a_j\right)^2 .$$

Thus, the payoff is at most 200. It is reduced by the quadratic difference between choices. In this game, any number in [0, 100] is an equilibrium and all equilibria are efficient. They differ in the degree of strategic risk. The least risky strategy is to choose 50. If you choose 50, the difference to your partner's choice can be at most 50 giving you a payoff of at least 150. This strategy is maximin, it is risk dominant, and it is the only symmetric equilibrium. In the experiment, subjects are re-matched within their group each period and repeat the game 80 times. In this baseline treatment, however, they quickly converge to all choosing 50.

The other treatments provide subjects with extrinsic information ("sunspots"). Subjects receive signals that are randomly drawn and may be 0 or 100. Since signals are in the same space as actions, they are salient in the sense that participants can easily relate them to the problem of choosing a number. Signals provide focal points for actions that may overcome the incentive to minimize strategic risk. In this game strategies can be ordered by risk dominance: a strategy is riskier if it is further away from 50. This order is used to measure the power of sunspots. If an extrinsic signal draws actions further away from 50 than another signal, we may say that the former is more powerful.

Signals are introduced as in the following example for Treatment CP. The computer randomly selects a number Z that may be 0 or 100 with probability ½. You do not know Z, but you receive two hints X and Y that are also either 0 or 100. Each of these hints is “correct“ (equal to Z) with 75% probability. With 25% probability it indicates the “wrong“ state. Hint Y is the same for you and your partner. Hint numbers X are drawn independently for you and your partner and may therefore differ from each other.

The payoff function is the same as before. Instructions and a quiz make sure that subjects understand that their payoffs only depend on their chosen numbers and not on any of the random numbers. Although the payoff function does not depend on signals, the set of equilibria gets larger. Any function mapping the common signal Y into the action space [0, 100] is an equilibrium. If all 6 subjects within a group coordinate on the same mapping, they are in equilibrium. Strategies that condition choices on private signals are no equilibrium and should be eliminated over time, because they do not survive iterative elimination of dominated strategies. Before analyzing behavior in Treatment CP, let us review the other treatments in this study.

In Treatment C, subjects receive just one common signal Y. This treatment is played by six groups. Five groups quickly converge to choosing the number indicated by Y. One group also coordinates on this equilibrium but needs more than 70 rounds to achieve this. In Treatment CC, subjects receive two common signals X and Y and subjects know that their respective partner gets the same two signals. This treatment is also played by six groups. Five groups quickly converge to choosing the average of the two signals (X+Y)/2. The remaining group comes close to this but is never perfectly coordinated. Results from Treatments C and CC indicate that pure public signals with a salient message are strong enough focal points to overcome the higher strategic risk associated with deviations from 50.

In three treatments subjects receive private signals only. Recall that conditioning actions on private signals is not an equilibrium. When the probability that one’s partner receives the same signal is as low as 62.5%, subjects quickly learn to ignore their signals and converge to playing 50. In two treatments the probability that the respective partner receives the same signal is 90%. Here, most groups converge to 50, but 4 out of 12 groups coordinate on sunspot-driven strategies, either following their signals or choosing 10 [90] conditional on the signal being 0 [100]. This show that salient messages may affect behavior in coordination games even if this is not an equilibrium. The power of extrinsic signals seems to be continuously increasing in the likelihood that the partner gets the same signal and reaches its maximum if signals are publicly provided.

While treatments with pure common signals reliably generate sunspot equilibria and imprecise private signals are neglected, the combination of common and imprecise private signals in Treatment CP (described above) leads to rather diverse patterns of behavior. Treatment CP is played by 12 groups. Figure 9 displays the average actions over 10 periods for each of these groups conditional on the 4 possible combinations of common (public) and private signal. Only one group (26) converges to the risk dominant equilibrium 50. Five groups converge to follow the public signal and choose Y, while two groups (24 and 29) coordinate on another sunspot equilibrium and choose 25 if the common signal is 0 and 75 if the common signal is 100. Three more groups (28, 32, and 34) are not fully coordinated, but close to this interior sunspot equilibrium. In one group (25), private signals affect behavior until the end. Subjects in this group do not manage to coordinate their choices.

These results indicate that extrinsic private signals have larger effects on behavior, if they are combined with public signals of similar salience. Comparing Treatments CP and C shows that the power of extrinsic public signals is reduced, if they are accompanied by private signals. Since players quickly converge to an equilibrium for pure public signals or pure private signals with low precision, these information settings have no effect on average payoffs. With highly correlated private signals, however, or when private and public signals are combined, subjects are less coordinated and achieve lower payoffs. Hence, sunspots may also have negative welfare effects.

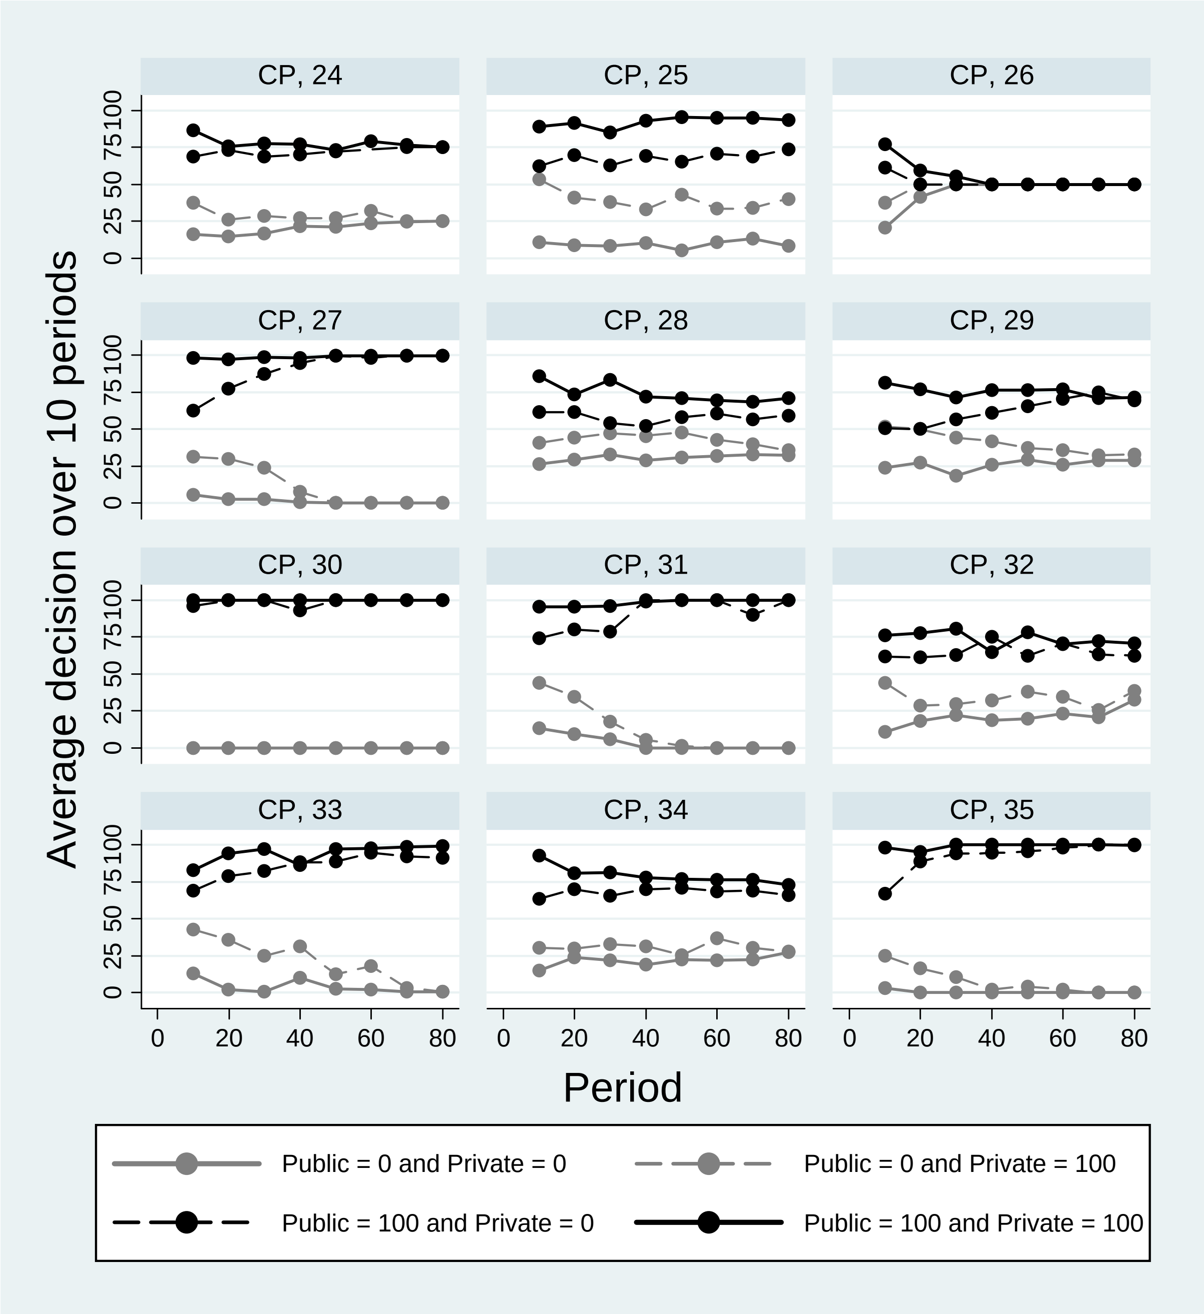


Figure 9: Average actions conditional on realized signals.

This experiment shows that sunspot equilibria are not just a theoretical curiosity. While global-game experiments, discussed in the previous section, give the impression that actions are fairly predictable by parameters of the payoff function even in games with multiple equilibria, the sunspot experiment recalls worries about destabilizing features of announcements. If salient messages may cause a swing of beliefs, they may counteract political measures that are taken to exploit comparative statics. If, for example, an increase in transaction costs is being communicated as a sign of weakness, it may actually trigger a speculative attack, although an increase in transaction costs should rather reduce the probability of an attack. Note, however, that in the sunspot experiment all equilibria yield the same payoffs, because in each equililbrium all players are perfectly coordinated. Equilibria differ by their levels of strategic risk, and sunspots are sufficiently strong to overcome risk dominance. In the refinancing games, discussed before, equilibria could be ordered by payoff dominance. It is yet an open question, whether sunspots can be sufficiently strong to coordinate players in a laboratory on a payoff-dominated equilibrium.

## 9. Conclusion

Laboratory experiments are a strong scientific tool to analyze basic properties of behavior and complement our theoretical knowledge. In situations of multiple equilibria, experiments can tell us which equilibrium is selected and how the selection can be influenced towards a socially desired outcome. Systematic deviations from rationality are frequently observed and the lab allows testing behavioral theories that provide potential explanations for these deviations. The interaction of theory and experiments advances both branches of research and, of course, there is a fruitful relation to field empirics as well. Here, I concentrated on those aspects of financial crises, where experiments provide us a better view and understanding of theoretical results. We may learn how the different theories should be used to appraise a given situation.

In particular, we may conclude:

1. Bubbles are unlikely to arise in a market in which traders experienced a bubble before.
2. Bubbles are more likely to arise if there is a lack of common knowledge about the time when a bubble is bound to burst.
3. People apply limited levels of reasoning. This reduces the likelihood of herding behavior and the relative weight attributed to public as compared to private signals.

Each financial crisis is related to a failure in refinancing debt. Refinancing games are coordination games, for which we can develop a good understanding if we combine theory and experiments.

4. Behavior in coordination games with multiple equilibria is fairly predictable.
5. Actual behavior deviates from the global-game selection in direction to more efficient strategies.
6. Comparative statics follows the global-game selection and is intuitive.
7. The global-game selection gives a good recommendation for individual behavior.

8. Better information leads to more efficient coordination.

Provided that different equilibria yield the same payoffs:

9. Extrinsic information may affect behavior and even lead to coordination on non-equilibrium strategies. Sunspots matter!
10. Extrinsic private information may reduce the ability to coordinate and lead to welfare losses.

There is a large gap to bridge from laboratory environments to financial markets and macroeconomic settings. Theories bridge this gap. With experiments we can test theoretical predictions in the lab and develop behavioral theories that may also apply to large markets.

## References


Abreu, Dilip, and Markus Brunnermeier (2002), Bubbles and Crashes, *Econometrica* 71, 173–204.

Anderson, Lisa R., and Charles A. Holt (1997), Information Cascades in the Laboratory, *American Economic Review* 87, 847-862.

Aschinger, Gerhard (2001): *Währungs- und Finanzkrisen*, Vahlen, München.

Brunnermeier, Markus, and John Morgan (2010), Clock Games: Theory and Experiments, *Games and Economic Behavior*, *Games and Economic Behavior* 68, 532-550.

Brunnermeier, M., and L. Pedersen (2009), Market Liquidity and Funding Liquidity, *Review of Financial Studies* 22, 2201-2238.

Carlsson, Hans, and Erik van Damme (1993), Global Games and Equilibrium Selection, *Econometrica* 61, 989–1018.

Cass, D. and K. Shell (1983), Do sunspots matter? *Journal of Political Economy* 91, 193–227.

Cifuentes, R., G. Ferruci, and H.S. Shin (2005), Liquidity Risk and Contagion, *Journal of the European Economic Association* 3, 556-566.

Cooper, Russel, Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross (1992), Communication in Coordination Games, *Quarterly Journal of Economics* 107, 739-771.

Cornand, C., and F. Heinemann (2010), Measuring Agents' Reaction to Private and Public Information in Games with Strategic Complementarities, CESifo Working Paper 2947, http://anna.ww.tu-berlin.de/~makro/Heinemann/download/ch_3.pdf .

Diamond, D., and Ph. Dybvig (1983), Bank Runs, Liquidity and Deposit Insurance, *Journal of Political Economy* 91, 401-419.

Dufwenberg, Martin, Tobias Lindqvist, and Evan Moore (2005), Bubbles and Experience: An Experiment on Speculation, *American Economic Review* 95, 1731-37.

Fehr, Dietmar, Frank Heinemann, and Aniol Llorente-Saguer (2011), The Power of Sunspots, SFB Discussion Paper 2011-070, http://www.macroeconomics.tu-berlin.de/fileadmin/fg124/heinemann/downloads/sunspots-2011.pdf .

Heinemann, Frank, Rosemarie Nagel, and Peter Ockenfels (2004), The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information. *Econometrica* 72, 1583–99.

Heinemann, Frank, Rosemarie Nagel, and Peter Ockenfels (2009), Measuring Strategic Uncertainty in Coordination Games, *Review of Economic Studies* 76, 181-221.

Kaminsky, G. and C. Reinhart (1999), The Twin Crises: The Causes of Banking and Balance-of-Payments Crises, *American Economic Review* 89, 473-500.

Kindleberger, Charles P., and Robert Z. Aliber (2005), *Manias, Panics, and Crashes: A History of Financial Crises*, 5th edition, Palgrave Macmillan.

Kübler, Dorothea, and Georg von Weizsäcker (2004), Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory, *Review of Economic Studies* 71, 425-442.

Minsky, H.P. (1972), Financial Instability Revisited: the Economics of Disaster, Boards of Governors of the Federal Reserve System, http://fraser.stlouisfed.org/historicaldocs/dismech/download/59037/fininst_minsky.pdf .

Morris, Stephen, and Hyun Song Shin (1998), Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks, *American Economic Review* 88, 587-597.

Morris, S., and H.S. Shin (2002), Social Value of Public Information, *American Economic Review*, 92, 1522-1534.

Morris, Stephen, and Hyun Song Shin (2003), Global Games: Theory and Applications, in *Advances in Economics and Econometrics, the Eighth World Congress* (edited by M. Dewatripont, L. Hansen and S. Turnovsky), Cambridge University Press, 56-114.

Qu, Hong (2011), *How do market prices and cheap talk affect coordination*? Dissertation, Carnegie Mellon University.

Reinhart, Carmen, and Kenneth Rogoff (2009), *This Time is Different*, Princeton Univ. Press.

Schmidt, David, Robert Shupp, James M. Walker, and Elinor Ostrom (2003), Playing Safe in Coordination Games: the Roles of Risk Dominance, Payoff Dominance, and History of Play, *Games and Economic Behavior* 42, 281-299.

Schotter, Andrew, and Tanju Yorulmazer (2009), On the Severity of Bank Runs, *Journal of Financial Intermediation* 18, 217-241.

Smith, Vernon L., Gerry Suchanek, and Arlington W. Williams (1988), Bubbles, Crashes and Endogenous Expectations in Experimental Spot Asset Markets, *Econometrica* 56, 1119-51.

Sutter, Matthias, Jürgen Huber, and Michael Kirchler (2012), Bubbles and Information: An Experiment, *Management Science* 58, 384-393.

Tirole, Jean (1985), Asset Bubbles and Overlapping Generations, *Econometrica* 53, 1499-1528.

Van Huyck, John B., Raymond C. Battalio and Richard O. Beil (1991), Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games, *Quarterly Journal of Economics* 106, 885-910.