MANIPULATION AND (MIS)TRUST IN PREDICTION MARKETS

Lawrence Choo, Todd R. Kaplan, and Ro'i Zultan

Discussion Paper No. 20-12

July 2020

Monaster Center for
Economic Research
Ben-Gurion University of the Negev
P.O. Box 653
Beer Sheva, Israel

Fax: 972-8-6472941
Tel: 972-8-6472286

Manipulation and (mis)trust in prediction markets[*]

Lawrence Choo Todd R. Kaplan Ro'i Zultan[†]

April 7, 2020

Abstract

Markets are increasingly used as information aggregation mechanisms to predict future events. If policy makers use markets to guide policy and managerial decisions, interested parties may attempt to manipulate the market in order to influence decisions. We experimentally find that, despite successful manipulation of prices, policy makers could still benefit from following the information contained in the market prices. Nonetheless, manipulation is detrimental to the policy decisions in two ways. First, manipulators affect market prices, making them less informative. Second, when there are manipulators, policy makers often ignore—or even act against—the information revealed in market prices. Finally, mere suspicion of manipulation erodes trust in the market, leading to the implementation of suboptimal policies—even without *actual* manipulation.

Keywords: prediction markets, policy, experiment
JEL codes: C92, D53, D8, G14

[*]We thank seminar participants at Samford University, University of Exeter, University of Haifa, Technion, and conference participants at the Economic Science Association international meeting in Berlin, Experimental Finance meeting in Heidelberg, 9th Workshop on Theoretical and Experimental Economics in Munich, the Israel Behavioral Finance meeting, Society for the Advancement of Economic Theory (SAET) in Ischia and the Econometric Society European Meeting in Manchester for useful comments and discussion. Financial support from the International Foundation for Research in Experimental Economics is gratefully acknowledged.

[†]Choo: FAU-Erlangen-Nuremberg, Chair of Economic Theory, Lange Gasse 20, 90403, Nuremberg Germany. E-Mail: lawrence.cy.choo@fau.de. Tel: +49 911-5302-222. Kaplan: University of Exeter, Exeter, EX4 4PU, UK, and University of Haifa, Mount Carmel, Haifa 31905,Israel. E-mail: toddrkaplan@gmail.com, Tel: +44-4392-263237. Zultan: Ben-Gurion University of the Negev, P.O.B. 653, Beer-Sheva 84105, Israel. E-mail: zultan@bgu.ac.il, Tel:+972(0)86472306.

1. Introduction

Prediction markets, where traded assets yield payoffs based on the future realizations of uncertain events, are able to aggregate dispersed information.[1] Predictions based on asset prices in such markets overwhelmingly outperform conventional forecasting methods (e.g., Arrow et al., 2008; Palan, Huber, and Senninger, 2019; Wolfers and Zitzewitz, 2004). It is not surprising, then, that governments as well as private corporations are increasingly using prediction markets as basis for policy decisions (e.g, Chen and Plott, 2002; Cowgill and Zitzewitz, 2015; Gillen, Plott, and Shum, 2017). Moreover, trading prices in natural financial markets can be used to inform policy making if it is difficult, or even impossible to design a dedicated artificial prediction market.[2] For example, if there is no clear future resolution of uncertainty, or the variables of interest are unobservable and hard to measure, there is no straightforward way of fixing the redemption value of the traded assets.

Consider, for example, legislation geared towards different energy technologies. Which one of the traditional or many alternative energy technologies is most efficient—and should therefore be supported by appropriate legislation—depends on a myriad of unknown variables. Increasing stock prices of sustainable energy technology firms may lead legislators to believe that the state of the world is favorable to such technologies, and vote accordingly.

If policy makers "listen" to the market, parties with vested interest in the policy decision may have an incentive to manipulate the market prices (Hanson, 2004).[3] In the example above, if energy companies expect stock prices to influence future legislation, they might artificially inflate their stock prices, incurring short-term market losses, in order to influence the policy making in their favor.

Such situations also arise naturally when private firms use prediction markets. For example, when a firm runs a prediction market to forecast future sales of a new product, competitors as well as parties within the firms whose interests do not perfectly align with the firm's may try to manipulate the market prices in

[1]The information aggregation properties of markets were first formally noted by Hayek (1945) and central to the *efficient market hypothesis* (e.g, Fama, 1970). See for example, Radner (1979), Muth (1961) and Ostrovsky (2012) for the theoretical properties of market aggregation. More recently, economist have argued that well-designed markets can be utilized as tools to gather information (e.g., Arrow et al., 2008; Plott, 2000).

[2]Future markets can sometimes be interpreted as natural markets for information.

[3]Listening to the market can sometimes be inefficient if the market also listens to policy makers. For example, central bankers may use bond markets prices to guide their policy decisions. However, bond market participants are also reacting to the central bankers' decisions. See *The Economist* article "Can central bankers talk too much?" (24th October 2019).

order to influence the firm's strategy.

We study the effects of such manipulation attempts in an experimental asset market, where the value of the traded assets is contingent on an underlying state of the world. In each market period, each trader receives a private signal that, in itself, is not sufficient to deduce the state. Nonetheless, the combination of all signals fully reveals the true state. Policy makers observe all transaction prices and vote on a policy, the outcome of which depends on the true state. In some periods, a *minority* of traders stand to gain from a policy that is, on the whole, harmful to the other traders and to the policy makers.

Our experiment findings can be summarized as follows. If it is common knowledge that there are no *manipulators* (i.e., traders with incentives to deceive the policy makers) in the market, the market prices are able to aggregate the information dispersed in the market to almost perfectly match the true values. When the non-existence of manipulators is not common knowledge, however, prices are more volatile, but are still sufficiently discriminating to (almost) fully reveal the true state. Nonetheless, the observing policy makers are reluctant to trust the market, leading to substantial loss of efficiency. When manipulators participate in the market, they have a substantial effect on market prices. While policy makers are still better off always implementing the policy suggested by market prices, often choose to ignore this information, or even support an alternative policy. These results highlight the important role of *trust* in the market. That is, manipulation hinders the ability of markets to inform optimal policy making via two distinct channels. One, by influencing market prices; two, by eroding trust in the market.

The success of prediction markets in forecasting various outcomes is generally taken to indicate that manipulation attempts are unsuccessful (Wolfers and Zitzewitz, 2004). Identifying actual manipulation attempts in the field is, however, a challenging endeavor. Predictions of political elections outcomes, for example, provide a situation where parties or candidates have a natural incentive to manipulate the market prices in order to create a 'bandwagon effect' (Rhode and Strumpf, 2004). Yet, opposing parties are generally equally likely to engage in manipulation attempts, which may cancel each other out. Instances where manipulation can be identified are rare.[4] Hansen, Schmidt, and Strobel (2004) found one such case in political markets, where a party sent an email to its supporters asking them explicitly to affect the market in order to gain politi-

[4]A notable exception is Camerer (1998), who actively placed temporary bets on horse races, to find that these had a negligible effect on closing betting odds.

cal traction, manipulation indeed influenced the market prices. Notwithstanding, if the researcher is aware of manipulation attempts, so are other traders in the market. When all the traders are fully aware of attempt to manipulate the marking, the attempt will likely prove futile since the traders can profit by gathering information of their own and being a counter-party to the manipulator (Hanson and Oprea, 2009). Even if there is a distortion, decision makers will take it into account when making decisions.

One can also study manipulation in financial markets. For instance, there has been analysis of price manipulation in Bitcoin exchanges (Gandal et al., 2018). The fraudulent transactions run by two bots caused volumes to spike and price to increase from $150 to $1000 in two months. The objective of at least one of the bots was to hide a loss of stolen bitcoins from the exchange by acquiring bitcoins to replace them and thus converting a bitcoin deficit into a less fragile fiat currency deficit. Since these transactions were apparently by the exchange itself, the ability to analyze the manipulation was thanks to a breach in the private data of the exchange and the carelessness in hiding the manipulation (under the presumption that transactions would remain secret). It is rare to find such clear examples due to the incentives to hide manipulation, the difficulty in knowing the true valuations, and not being able to know the objectives and information of the traders.

In contrast to field studies, laboratory experiments provide a controlled environment where the ability of markets to aggregate disperse information can be studied (Deck and Porter, 2013; Plott and Sunder, 1982, 1988). In an experiment, some traders can be endowed with incentives to manipulate the market. Thus, manipulation attempts can be directly measured, and the market outcomes are fully observable. Several studies looked at manipulation in single asset markets, where the value of the asset depends on an unknown state of the world. In Hanson, Oprea, and Porter (2006), some traders received additional payment to their market earnings based on the median transaction price, incentivizing them to push prices up. Manipulation attempts did not have a significant effect on prices. Apparently, traders in the market—who know that some traders are incentivized to inflate prices—were able to counter the manipulation attempts. Similarly, in Veiga and Vorsatz (2010), a robot trader that created artificial demand and supply was not able to manipulate prices.[5] However, in both those studies—even without manipulators—average prices did not

[5]Such manipulation did lead to higher prices if the true value was low in a different setting, where some traders have perfect information regarding the true value (Veiga and Vorsatz, 2009, 2010).

move substantially with the true value of the asset, indicating that the market did not fully aggregate the information. The question of whether manipulators are able to impede efficient information aggregation therefore remains open. Furthermore, these studies did not address the issue of whether policy makers trust markets that are susceptible to manipulation, which is at the heart of the current study.

A recent study by Maciejovsky and Budescu (in press) highlighted the importance of trust. Their experiments compared the ability of group communication and markets to aggregate information (see also Maciejovsky and Budescu, 2013). Under manipulation incentives, markets outperformed the groups. Nonetheless, participants reported more trust in the groups, as did third-party observers.

Deck, Lin, and Porter (2013) introduced forecasters, who observe the market activity and make costly investments. Their experiment involved two possible states of the world, with each trader receiving an independent stochastic signal about the true state. Manipulators knew the true state with certainty, but did not receive their market earnings and were only paid based on forecast errors. Without manipulators, prices did not converge to the benchmark levels, but were informative enough to improve forecasts. With manipulators, prices were completely non-informative, and forecasts made by inexperienced forecasters were even negatively correlated with the true state.

The current research goes beyond the existing literature in addressing several open issues. First, most existing studies used a single-asset setup, in which markets fail to aggregate information efficiently even without manipulation (cf. Corgnet, DeSantis, and Porter, 2015; Plott and Sunder, 1988). We study manipulation in markets with state-contingent assets, which are able to aggregate information efficiently, and test whether manipulation is able to undermine this ability. Second, previous studies focused on market behavior. Our main focus is on the policy makers' response to market prices in view of potential manipulation. Accordingly, we go beyond the existing results in differentiating between two possible outcomes of manipulation: obscuring the true state, and promoting a false state. We do so by including a third, neutral, state, which is neither the true state nor the one favorable to the manipulators. Furthermore, we allow the policy makers to vote for a status-quo policy, which is not state specific. This allows us to estimate trust in the market, and draw a distinction between ignoring the market and voting against the market. Finally, we vary whether the existence of manipulators is commonly known and test the effects of this variable on information aggregation and policy decisions.

We study a market with Arrow-Debreu securities, each corresponding to one possible state of the world.[6] These types of assets, dubbed by Wolfers and Zitzewitz (2004) "Winner-take-all contracts", are able to efficiently aggregate disperse information even in complex situations (Choo, Kaplan, and Zultan, 2019). After the end of the trading period, policy makers—who observe all transactions—vote on multiple policies, each optimal in a different state of the world. Voting for a "safe" status quo option is also allowed, which is implemented if none of the policies receive a majority of votes. The introduction of a status quo option allows us to estimate the trust that policy makers place in the market prices, and to study how this trust varies according to market activity and the policy makers' awareness of manipulation attempts. We introduce manipulators by varying the incentives of two of the traders across market periods. With probability 0.5, these traders substantially gain from the implementation of a policy that they know is *not* socially optimal. Importantly, while in previous laboratory studies the existence of manipulators in the market was common knowledge, we compare situations with and without common knowledge. This comparison serves two purposes. First, it affects the ability of other traders actively can counteract manipulation effects. Second, it allows us to estimate the effect of knowledge of manipulation on the policy makers.

2. Experimental design and procedure

Each session included twelve participants, who participated in fourteen experimental market periods. The participants were randomly allocated to roles of eight traders and four policy makers. Two of the traders—the potential manipulators—are designated as *red* ($\mathscr{R}$) traders, and the other six traders as *blue* ($\mathscr{B}$) traders. To facilitate comprehension, all roles ($\mathscr{R}$ traders, $\mathscr{B}$ traders and policy makers) were fixed across all periods.

Each period consisted of a trading stage and a voting stage, with different subsets of participants (traders or policy makers) active in each stage. We manipulated two independent variables in a 2×2 mixed between-within design. First, the existence of manipulators varied within subjects across the market periods, as a random uniform draw determined independently for each period whether the market included manipulators (*Man*) or not (*NoMan*). Second, the

[6] To continue our example above, one may think of these securities as stocks of firms specializing in different energy technologies.

Table 1: Summary of the experimental design.

Treatment	Manipulator traders?	Existence of manipulators announced?	Number of sessions
CK-NoMan	No	Yes	7
CK-Man	Yes	Yes	
NCK-NoMan	No	No	7
NCK-Man	Yes	No	

results of the random draw were announced in the Common Knowledge (*CK*) sessions, whereas in the No Common Knowledge (*NCK*) sessions, the other traders and policy makers (i.e., anyone who is not a manipulator) only knew that there is a 50-50 chance that there are manipulators in the market. Table 1 summarizes the four resulting treatments. In the following, we describe in detail the market procedure, followed by the detailed design.

2.1. Market procedure

Each market involves eight traders (active in the trading stage) and four policy makers (active in the voting stage). At the beginning of the market stage, the eight traders are randomly allocated into two information groups of four traders each, with the $\mathscr{R}$ traders always placed in the same information group. The $\mathscr{B}$ traders do not know if they are grouped with the $\mathscr{R}$ traders or not.

2.1.1. Trading stage

Before the trading stage commences, nature selects one of three possible states of the world, X, Y, and Z, with equal probabilities. Each information group of four traders is then informed that one of the other two states is not the true state of the world. For example, if nature selects state Y, one information group is informed that state X is not true and the other that state Z is not true. The policy makers do not receive any private information.[7]

Traders trade three Arrow-Debreu securities x, y, and z (corresponding to the three possible states, X, Y, and Z) in three concurrent markets. Trade takes place using the continuous double auction mechanism as follows. At the beginning of trade, each trader is endowed with 200 ECU (experimental currency

[7]That is, the policy makers only know that the true state is X, Y or Z with equal probabilities.

units) and 5 units of each security type. During the trading duration of 120 seconds, traders can place bids and asks (in the range of 0–20 ECU)—and accept open bids and asks—for each of the three securities. Short-sales are prohibited. When the markets close, each security pays a dividend of 10 ECU if it corresponds to the true state of the world and 0 ECU otherwise.[8] The paid dividends are added to the traders' capital balances to determine their trading stage earnings.

2.1.2. Voting stage

The policy makers observe all of the transaction prices (but not open bids and asks) in the trading stage and proceed to the voting stage. Each policy maker casts a vote for one of three policies $\mathscr{X}$, $\mathscr{Y}$, and $\mathscr{Z}$ (corresponding to the three possible states, X, Y, and Z), or for the status quo $\mathscr{Q}$. The policy ($\mathscr{X}$, $\mathscr{Y}$, $\mathscr{Z}$ or $\mathscr{Q}$) that receives the most votes is implemented. In case of a tie, the status quo $\mathscr{Q}$ is implemented by default.[9] This tie-breaking rule is simpler and arguably more realistic than the alternative solving the tie by randomizing over the tied policies. In principle, this means that a voter may vote strategically for a policy she *does not* prefer in order to induce a tie to prevent the implementation of another policy. This situation, however, only arises under implausible beliefs and is negligible.[10]

2.1.3. Payoffs from the implemented policy

Independent of the true state, implementing the status quo $\mathscr{Q}$ yields a payoff of 100 ECU for each trader and policy maker. Participants' payoffs from implementing any of the three policies $\mathscr{X}$, $\mathscr{Y}$, or $\mathscr{Z}$ depend on the state of the world, their role and the market type.

The $\mathscr{B}$ traders and policy makers gain 400 ECU from the implementation of the policy that corresponds to the true state, and lose 400 ECU if any of the poli-

[8]For example, if the true state is Y then security y pays a dividend of 10 ECU and the other securities 0 ECU.

[9]For example, if policies $\mathscr{X}$, $\mathscr{Y}$, $\mathscr{Z}$ and $\mathscr{Q}$ receive 2, 1, 1 and 0 votes, respectively, then policy $\mathscr{X}$ is implemented. Alternatively, if two votes go to policies $\mathscr{X}$ and $\mathscr{Y}$ each, the status quo $\mathscr{Q}$ is implemented.

[10]The only situation where a policy maker may want to vote strategically is if she believes that the other votes are split 2:1 among two policies that she *does not* want implemented. Even if a policy maker hold the unlikely beliefs that the other voters may support two different policies that she herself does not support, there is no reason to expect that the split will be exactly 2:1. Empirically, such ties occurred in less than 4% of all markets in the experiment.

Table 2: Implemented policy payoffs.

	NoMan markets			Man markets		
Implemented poicy	Policy maker	$\mathscr{B}$ trader	$\mathscr{R}$ Trader	Policy maker	$\mathscr{B}$ trader	$\mathscr{R}$ Trader
True Policy	400	400	400	400	400	−400
Fake Policy	−400	−400	−400	−400	−400	1,000
Neutral Policy	−400	−400	−400	−400	−400	−400
Status quo	100	100	100	100	100	100

cies that correspond to the other two states is implemented.[11] Note that in the absence of the trading stage, voting for the status quo maximizes the expected payoff for a policy maker unless she assigns a probability of at least 0.625 to one of the three states.[12]

The payoff for the $\mathscr{R}$ traders depends on the market type. In the *NoMan* markets, the $\mathscr{R}$ traders receive the same payoff as the other participants in the market. In the *Man* markets, they receive a high payoff of 1,000 ECU if the implemented policy is the one that corresponds to the state they know not to be true (i.e., a policy that harms the other participants) and lose 400 ECU from the implementation of either of the other two policies. This payoff structure incentivizes the $\mathscr{R}$ traders to manipulate prices in the *Man* markets, in order to influence the policy makers' beliefs and consequently the implemented policy.

Henceforth, we refer to true state of the world as the *True* state, the state that the $\mathscr{R}$ traders know not to be true as the *Fake* state and to the remaining state as the *Neutral* state. For convenience, we maintain this terminology for the corresponding policies and securities.[13] Table 2 summarizes the payoffs from the implemented outcome by market type and role.

2.1.4. Total payoffs

Writing π_i for the payoff to individual i from the implemented policy, the payoff of each policy maker is $650 + \pi_i$. The corresponding payoff for trader i is given

[11] Suppose that the true state is X, the $\mathscr{B}$ traders and policy makers receive 400 ECU if policy $\mathscr{X}$ is implemented, −400 ECU if policies $\mathscr{Y}$ or $\mathscr{Z}$ is implemented and 100 ECU if the status quo $\mathscr{Q}$ is implemented.

[12] Recall that policy makers do not receive any private information about the true state. This implies that in the absence of the trading stage, the policy maker should assign equal posteriors to each possible state of the world.

[13] For example, if the true state is Y and the $\mathscr{R}$ traders are informed that X is not true, then the True state, True security, and True policy are Y, y, and $\mathscr{Y}$, respectively; The Fake state, security, and policy are X, x, and $\mathscr{X}$, respectively; and the Neutral state, security, and policy are Z, z, and $\mathscr{Z}$, respectively.

by

$$400 + [\underbrace{L_i + d(x)e_i^x + d(y)e_i^y + d(z)e_i^z}_{\text{Trading stage earnings}}] + \pi_i,$$

where $L_i \geq 0$ is the trader's cash balance at the end of trading stage, e_i^j is her inventory of security $j = x, y, z$ at the end of trading stage and $d(j)$ is the dividend of security j. The difference in base payment between traders and policy makers makes up for the value of the trader's endowment and therefore, their average trading stage earnings.

2.2. Treatment design and experimental procedure

The first part of the experiment was a training phase consisting of one practice and five experimental periods, in which participants could learn the trading mechanism and information structure.[14] Each period followed the design and procedure of the *NoMan* market described above, with the exception that there was no voting stage. Instead, there were no policy makers, and all twelve participants participated in the role of traders, divided into two information groups of six traders each.

The main part of the experiment consisted of one practice period and fourteen experimental periods. Each period included either a *NoMan* or a *Man* market design with equal probabilities. For efficient between-treatment comparisons, we pre-generated a sequence of states and market types, which we implemented in all sessions. We ran seven sessions for each of the *Common Knowledge* (CK) and *No Common Knowledge* (NCK) treatments.

***CK*:** At the beginning of each period, all participants were informed as to whether they are participating in a *NoMan* or a *Man* market.

***NCK*:** At the beginning of each period, only the $\mathscr{R}$ traders were informed as to whether they are participating in a *NoMan* or a *Man* market.

The experiment was conducted at the University of Exeter FEELE laboratory in 2018 and 2019. The student subjects were recruited through ORSEE (Greiner, 2015). The experiment was programmed with z-Tree (Fischbacher, 2007). At the end of each session, one period (out of five) from the training phase and two periods (out of fourteen) from the experimental phase were randomly chosen

[14]The instructions are detailed in Appendix Appendix B and the experiment data is available upon request.

for payment. Payoffs were converted to cash at the rate of 100 ECU equals 1 GBP, and added to a show up payment of 5 GBP.

3. Theoretical analysis

We maintain the terminology introduced above to denote the true state (and corresponding security and policy) as *True*; the state that the $\mathscr{R}$ traders know not to be true (and corresponding security and policy) as *Fake*; and the remaining state, security, and policy as *Neutral*. We denote the $\mathscr{B}$ traders who are in the same and different information groups as the $\mathscr{R}$ traders as $\mathscr{B}^1$ and $\mathscr{B}^2$ traders, respectively.

We evaluate the market's success at aggregating information by comparing market prices against two benchmark models: the *rational expectations equilibrium* (Radner, 1979, henceforth REE) and the *prior information equilibrium* (Choo, Kaplan, and Zultan, 2019; Plott and Sunder, 1982, 1988, henceforth PIE). The REE and PIE are both static models, and differ with respect to whether beliefs are exogenous or endogenous to the market activity. Given these beliefs, both models assume that the standard principles of supply and demand determine the market prices.

In the REE, beliefs are Bayesian-rational given the market prices, and the prices are consistent with utility maximization given the traders' beliefs. Without manipulators (as in our *NoMan* markets), REE prices reflect the aggregate information held by all traders (Radner, 1979). In our setting, the true state of the world is fully revealed, and the REE prices match the true values of the securities.[15] Let us call the REE in which all traders are informed about the true state of the world and the securities are traded at their true values as the *fully revealing equilibrium* (henceforth FRE).

The PIE model, in contrast, assumes that traders form beliefs based on the exogenously provided information, and do not condition expectations on observed prices (*Unsophisticated equilibrium* in the language of Radner, 1979). In addition to providing a starting point for the belief updating dynamics, the PIE model provides a better fit than the REE model in single-asset markets (Corgnet, DeSantis, and Porter, 2015; Corgnet et al., 2018; Plott and Sunder, 1988) and with inexperienced traders (Choo, Kaplan, and Zultan, 2019).

Applying the REE model to markets with manipulators (as in our *Man* mar-

[15] There is an extensive literature studying how REE prices can result from dynamic behavior of traders (e.g., Dubey, Geanakoplos, and Shubik, 1987; Hellwig, 1982; Ostrovsky, 2012).

kets) is not straightforward, as the general characterizations of the REE assume that traders trade according to their true valuations, which we expect **not** to be the case with manipulators.[16] Consequently, new equilibria exist, including, as we show below, an equilibrium in which the true state of the world is not revealed in prices.

We therefore complement the static equilibrium analysis with a dynamic myopic reasoning model (henceforth, MRM) to generate predictions for which of the equilibrium prices will converge to. This model considers a simplified discrete-time trading process, wherein supply and demand correspond to traders' beliefs, and beliefs are updated in each period based on the market clearing prices.[17]

In the following subsections, we first analyze the PIE and REE in our setting, followed by the MRM analysis. The analysis assumes that policy makers are risk neutral, and use Bayes' rule whenever updating their beliefs.

3.1. Equilibrium analysis

First, a note on how supply and demand determine the market prices is in place. The standard assumption in analyzing experimental markets is that prices converge to the highest valuation in the market (e.g., Plott and Sunder, 1982, 1988). The rationale is that, while the number of securities in the market constrain supply when short-sales are prohibited, demand is (in principle) unlimited. In the continuous double auction, however, each trader is limited to buy or sell orders of one unit of security at any given point in time. We therefore differentiate between the short-run and the eventual supply and demand. In the short-run—which is relevant for learning—we define the supply and demand at a given price to be the number of traders willing to sell and buy a security, respectively, at that price. Eventually, all traders but those with the highest valuation will sell their complete inventory, at which point the prices will converge towards the highest valuations.[18]

[16]Strictly speaking, the payoffs from the market activity include the potential policy payoffs, which far outweigh the value of the securities.

[17]This model generates predictions not only for final prices, but also for the dynamic dissemination of information in the market. See Choo, Kaplan, and Zultan (2019) for evidence that traders that learn the true state first according to the MRM indeed bought more of the valuable asset and were instrumental in price convergence to equilibrium.

[18]To illustrate, consider the situation without manipulators. Each asset is valued at either zero or 5.00 (being worth 10.00 with probability 0.5). With unlimited demand, all prices go to 5.00, and there would be no learning. It is more plausible—and is indeed clear from the results—that prices during the adjustment period reveal information, in line with our short-run assumption.

3.1.1. Rational expectations equilibrium

With this note in mind, we now proceed to the equilibrium analysis. The analysis of the REE with manipulators should take into account the effect prices have on the policy makers, and incorporate the gains from the implemented policy into the traders' utility. In the following, we make the simplifying assumption that the policy makers vote for policy $\mathscr{X}$, $\mathscr{Y}$ or $\mathscr{Z}$ if and only if the price of the associated security is higher than those of the other two securities. Otherwise, they vote for the status quo.[19] Moreover, we assume that the effect on the implemented policy trumps any potential gains from the market.

Proposition 1. *With manipulators, the fully revealing equilibrium (FRE) is an REE.*

Proof. The FRE prices fully reveal the true state of the world. The informed $\mathscr{B}$ traders trade truthfully, as it both maximizes their market payoffs and informs the policy makers of the true state. Hence, they do not supply (demand) the securities below (above) their true values. The manipulators ($\mathscr{R}$ traders) are therefore unable to influence the prices, and are better off trading according to their (correct) valuations. □

Moreover, if prices discriminate between the True security and the other two securities, it follows that in equilibrium traders are informed of the true state of the world, leading to the FRE prices. Nonetheless, there also exists a non-revealing REE.

Proposition 2. *There exists a non-revealing rational expectations equilibrium (NRE) in which the True and Fake securities are traded at the same (positive) price, and the Neutral security is traded at a price of zero.*

Proof. The price of the Neutral security reveals that the Neutral state is not the true state of the world, but prices do not discriminate between the True and Fake securities. This fully reveals the true state of the world to the $\mathscr{R}$ and $\mathscr{B}^1$ traders, while the $\mathscr{B}^2$ traders maintain their prior beliefs (which assign equal

[19] Recall that the expected payoff from implementing policy $\mathscr{X}$, $\mathscr{Y}$ or $\mathscr{Z}$ is higher than 100 ECU, the payoff from implementing the status quo, if the probability of the associated state being the true state is at least 0.625. This is trivially true if price discrimination can only favor the true asset. In this case, voting for the policy is weakly dominant. Note that it is *not* otherwise dominant to vote for the status quo. If two other policy makers vote for one policy while the third votes for another policy, a policy maker who prefers the status quo should join the minority voter rather than vote truthfully. implementing the status quo through split votes, however, requires unlikely coordination between policy makers, and only happened in four of 140 markets.

probabilities to the True and Fake states being true). Furthermore, the policy makers hold similar beliefs, and vote for the status quo. The $\mathscr{B}^2$ traders, therefore, must pose the same supply and demand to both securities. The manipulators ($\mathscr{R}$ traders) are willing to forgo any market payoff in order to obscure the true state of the world, which is attainable if they mirror the supply and demand of the informed $\mathscr{B}^1$ traders. That is, they match their supply and demand for the Fake security to the supply and demand of the $\mathscr{B}^1$ traders for the True security, and vice versa. □

Note that in the NRE, the short-run price of the True and Fake securities is 5.00. In the long-run, it is possible that the informed $\mathscr{B}^1$ traders run the price of the True security up to 10.00, with the manipulators doing the same for the Fake security. Therefore, any price between 5.00 and 10.00 is consistent with the analysis, as long as it does not discriminate between the True and Fake securities. While these prices can leave arbitrage opportunities, this can be explained by the $\mathscr{R}$ and $\mathscr{B}^1$ traders willing to take losses in the market in order to the influence the voting stage. The $\mathscr{R}$ traders are effectively subsidising the market, from their policy-based profits.

3.1.2. Prior information equilibrium

The PIE describes the market clearing prices when traders update their beliefs about the true state given their own private information and condition their demands for securities upon such posteriors, but do not update their beliefs any further based on the observed prices. In the *NoMan* markets, all traders believe the True security to be true with probability 0.5, and therefore value it at 5.00 ECU. For each the other two securities, there is one group that values it at 5.00 ECU, whereas the other group values it at zero.[20] Thus, the market-clearing prices of the True, Fake and Neutral securities will be 5.00, 2.50, and 2.50 ECU, respectively.[21] Note that, although the true state of the world is only partially revealed in prices, there is sufficient information for the policy makers to infer the true state and vote for the True policy.

In the *Man* markets, we maintain the assumption that traders take into account the potential policy payoffs when trading. Specifically, manipulators may

[20]The $\mathscr{R}$ and $\mathscr{B}^1$ value the Neutral security at 5.00 ECU and the Fake security at 0.00 ECU, and vice versa for the $\mathscr{B}^2$ traders.

[21]Any price strictly between zero and 5.00 ECU will clear the markets for the Fake and Neutral securities. Taking the midpoint for simplicity, as we do here and in the following, does not affect the analysis.

Table 3: The REE and PIE equilibrium prices.

	Security prices			Implemented policy
	True	Fake	Neutral	
Prior information equilibrium				
NoMan	5	2.5	2.5	True policy
Man	5	5	0	Status quo
Rational expectations equilibrium				
NoMan	10	0	0	True policy
Man (FRE)	10	0	0	True policy
Man (NRE)	5^+	5^+	0	Status quo

Note. 5^+ Prices for the True and Fake securities may increase above 5 but by the same amount.

distort their supply and demand to obscure the true state of the world. Recall that the manipulators are only informed which is the Fake state, so they cannot discriminate between the True and the Neutral securities. Nonetheless, they can obscure the true state by trading as if they know the Fake state to be true. This creates excess demand (resp. supply) for the Fake security at any price above (resp. below) 5.00 ECU, the value assigned by the $\mathscr{B}^2$ traders. On the other hand, as only the two $\mathscr{B}^1$ traders are willing to buy the Neutral security at any positive price, excess supply will drive its prices down to zero. The supply and demand for the True security remain as in the *NoMan* markets. The market-clearing prices of the True, Fake and Neutral securities are, therefore, 5.00, 5.00, and zero ECU, respectively, as in the NRE. As prices do not distinguish between the True and Fake securities, risk-neutral policy-makers should vote for the status quo policy. The following proposition summarizes the PIE equilibrium analysis.

Proposition 3. *In the* NoMan *market, the PIE prices reveal the true state of the world. In the* Man *market, the PIE prices only eliminate the Neutral state as a possibility and are equivalent to the NRE prices.*

Table 3 summarizes the REE and PIE equilibrium prices and implemented policies.

3.2. Dynamic myopic reasoning model

The MRM assumes that trade takes place over $t \in \{1, 2, ...\}$ hypothetical periods. In each t, traders proceed according to the following four stages:

Stage 1 Traders observe period $t-1$ prices of all securities.

Stage 2 Traders update their beliefs about the true state.

Stage 3 Traders set their supply and demand for each security according to their updated beliefs.

Stage 4 The market clears at a price that equates supply and demand.

Proposition 4. *The dynamic myopic reasoning model predicts that prices converge to the FRE without manipulators, and to the NRE with manipulators.*

Proof. In the *NoMan* markets, period 1 beliefs are set by the prior information that the trader holds. That is, the starting point is the PIE. In the *Man* markets, we continue to assume that the $\mathscr{R}$ traders set their supply and demand of the securities *as if* they know the Fake state to be true (i.e., as if they value the Fake security at above 5.00 ECU and the other securities at zero ECU).

***No Manipulation (NoMan)* markets.** At period $t = 1$, the analysis is identical to that provided above for the PIE. Traders value the securities corresponding to the states that they know to be possible at 5.00 ECU. The resulting market clearing prices of the True, Fake and Neutral securities are 5.00, 2.50 and 2.50 ECU, respectively. This price profile uniquely identify the true state of the world. Therefore, at period $t = 2$, all traders value the True security at 10.00 ECU and the other securities at zero ECU. The resulting market clearing prices of the True, Fake and Neutral securities are hence 10.00, 0.00 and 0.00 ECU, respectively. Since traders are fully informed about the true state, there will be no further revisions to prices in period 3. That is, prices converge to the FRE.

***Manipulation (Man)* markets.** At period $t = 1$, the $\mathscr{B}^1$ and $\mathscr{B}^2$ traders behave as in the *NoMan* market. In contrast, the $\mathscr{R}$ traders will only demand the Fake security. The resulting market clearing prices of the True, Fake and Neutral securities are hence 5.00, 5.00 and 0.00 ECU, respectively, as in the PIE. These prices reveal that the Neutral state is not the true state of the world. Therefore, the $\mathscr{R}$ and $\mathscr{B}^1$ traders—who can also rule out the Fake state—have sufficient information to deduce the true state. The symmetry between the True and Fake security prices, however, imply that the $\mathscr{B}^2$ traders are still uninformed about the true state. This symmetry persists in the next period $t = 2$. The $\mathscr{B}^2$ traders, who value both the True and the Fake securities at 5.00 ECU, form a majority of the market. Hence, there is excess supply (resp. demand) above (resp. below) the price of 5.00 ECU for both securities. The resulting market-clearing prices of the True, Fake and Neutral securities remain at 5.00, 5.00 and 0.00 ECU, respectively, and an equilibrium is reached.[22] Thus, the MRM predicts that prices converge to the NRE prices in the manipulation markets. □

4. Results

4.1. Trading stage

We commence with the analysis of security prices in the trading stage. After establishing the effects of manipulators and common knowledge on market prices and dynamics, we proceed to look at voting behavior in the voting stage. Finally, we combine the market and voting data to reveal the effects of manipulators and common knowledge on the policy makers' strategies, and more specifically on the level of trust and mistrust in the market when casting a vote.[23]

We define the *market closing price* to be the average price over the last five transactions of a security.[24] Figure 1 presents violin plots of the market closing prices by security type and treatment.[25] The findings in the *CK-NoMan* treatment, where it is common knowledge that there are no manipulators, are striking. Prices converge almost perfectly to the FRE prices. Both the median and mode of the True security equal to the true value of 10.00. The modal price of the Fake and Neutral securities is zero, with the median prices not far above zero, at 0.46 and 1.20, respectively. Thus, we state our first result:

Result 1. *When it is common knowledge that there are no manipulators in the market, Arrow-Debreu markets are successful at aggregating information about the true state of the world into prices.*

The comparison with the *NCK-NoMan* treatment reveals the importance of common knowledge, which was maintained in all previous studies. We see that—although there are no manipulators in the market—the mere suspicion of manipulation is enough to impede price convergence, with many transaction prices

[22]As noted in the REE analysis, eventually the prices of the True and Fake securities will increase towards the long run equilibrium price of 10.00 ECU.

[23]See Appendix A.1 for detailed transaction prices and votes by session.

[24]Our interest in market closing prices is consistent with the theoretical analysis of information aggregation as a dynamic process. We take the last five transactions to "smooth out" the volatility in prices. All of the results are robust to using the last ten transactions or the transactions taking place in the last 60 seconds of trade.

[25]Our data comprised of 14 sessions with 14 periods in each session. This resulted in 196 markets. There were at least one transaction for each security type in 98.9% of all markets. The market closing prices involve around 48%, 36%, 44% and 33% of all transactions in the *CK-NoMan*, *NCK-NoMan*, *CK-Man* and *NCK-Man* treatments, respectively.

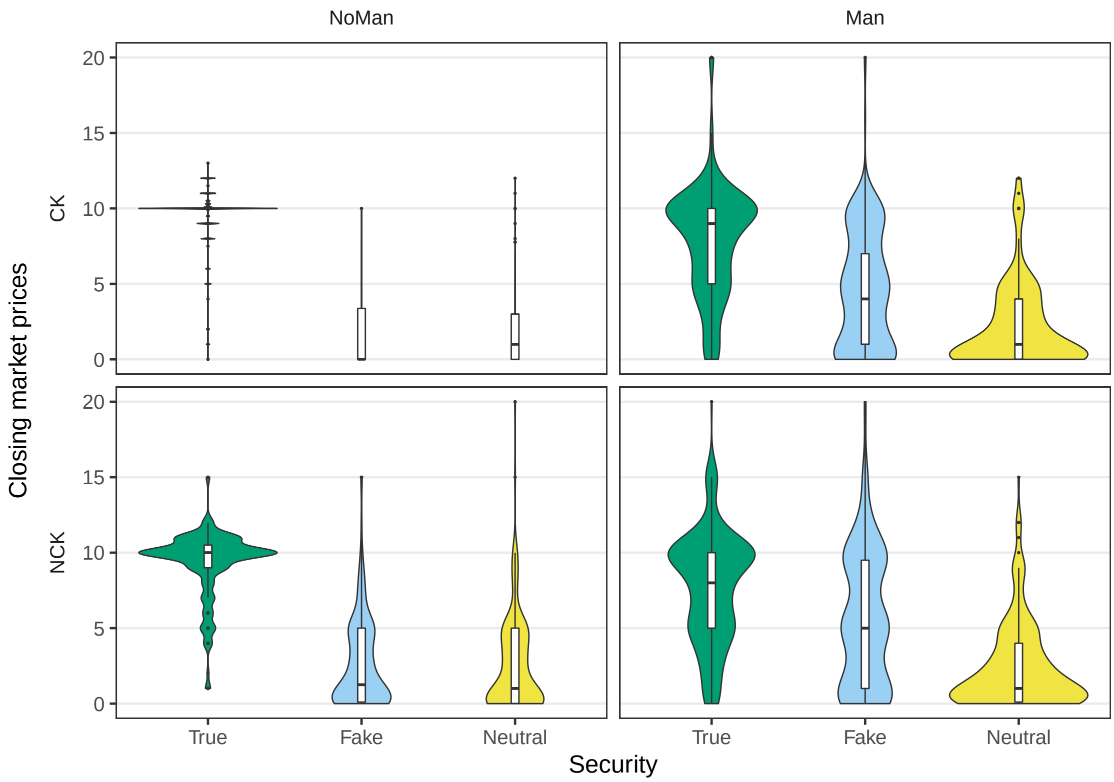

Notes. Market closing prices are defined as the average price in the last five transactions for that security. The violin plots (shaded areas) present kernel distributions of market closing prices. The boxplots (candlesticks) present the median, interquartile region, and outliers.

Figure 1: Market closing prices.

substantially above or below the true values of the securities. Although the median transaction price for the True security still reflects its true value, the prices of the two other securities are mostly distributed around the PIE price of 2.50 ECU. Our next result reflects the importance of common knowledge in the market:

Result 2. *The mere suspicion of manipulators—even when there are none in the market—impedes the information aggregation properties of Arrow-Debreu markets.*

Moving to the *Man* markets, we see that manipulators have a substantial influence on prices. he median price of the True security is now below the true value of $10,00$ ECU, while the prices of the Fake security vary around the PIE price of 5.00 ECU. This pattern is more pronounced in the *NCK-Man* treatment compared to the *CK-Man* treatment, to the extent that the price distributions of the True and Fake securities, taken over the manipulation-market periods, are not significantly different from each other (Kolmogorov-Smirnov $p = 0.541$ in *NCK-Man* compared to $p = 0.012$ in *CK-Man*). This result suggests that, when there is common knowledge that manipulators are active in the market, non-manipulator traders are able to counter the manipulation attempts, albeit only to a small extent.[26] Our third result summarizes the effect of manipulation in the market:

Result 3. *A minority of manipulators are able to substantially harm the information aggregation properties of Arrow-Debreu markets. Common knowledge of manipulation attempts somewhat mitigates the effect of manipulators on market prices.*

Results 1-3 are confirmed by testing price convergence to the theoretical equilibrium predictions. To do so, define for each market the variables M^F and M^P as the mean square deviations of the market closing prices of each security from the FRE and PIE prices, respectively. Note that separate comparisons to the NRE are redundant, as it is only relevant in the *Man* markets, where it is equivalent to the PIE. Figure 2 plots the means and 95% confidence intervals of M^F and M^P and reports nonparametric tests comparing the two measures across treatments.[27]

[26]The manipulators are able to gain from manipulating the markets, obtaining mean payoffs of 605 ECU and 667 ECU in the CK-Man and NCK-Man treatments, resepctively, compared to 250 ECU obtainable by doing nothing, assuming that the True policy is then implemented.

[27]The M^F and M^P values are computed for each of the 196 markets. This resulted in 42 observations each in *CK-NoMan* and *NCK-NoMan*, and 56 observation each in *CK-Man* and *NCK-Man*.

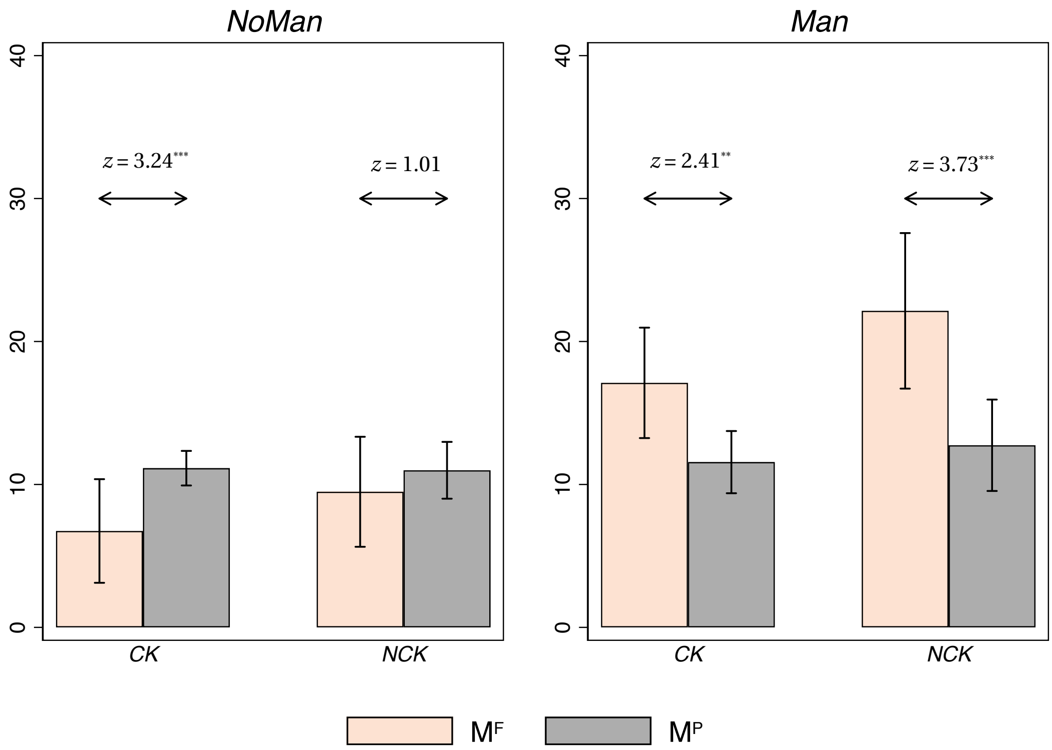

Notes. The M^F and M^P are the mean square deviations of closing market prices from the FRE and PIE prices, respectively. Note also that the NRE is equivalent to the PIE in the *Man* markets. the z-scores for comparisons between treatments are based on Mann-Whitney tests. The z-scores for comparisons between market types are based on Wilcoxon signed-ranks tests. ***, ** and * indicate significance at $p < 0.01$, $p < 0.05$ and $p < 0.10$, respectively.

Figure 2: Mean and 95% CI of M^F and M^P by treatment.

In the *CK-NoMan* treatment, closing market prices are significantly closer to the FRE than to the PIE predictions, indicating successful information aggregation. The picture changes in the *NCK-NoMan* treatment, where the traders are not informed that there are no active manipulators in the market. Here, the deviation of prices from the FRE prices is larger, and similar to the deviation from the PIE prices.[28] Manipulators in the *CK-Man* and *NCK-Man* treatments are successful at impeding information aggregation in prices, as the deviations from the PIE prices are now significantly smaller than the deviations from the FRE prices.

In Appendix A.2, we provide additional details on the trading behavior in the different treatments. In particular, we indeed observe that the manipulators ($\mathscr{R}$ traders in the *Man* markets) create artificial demand for the Fake security.

4.2. Voting stage

Figure 3 presents the distributions of votes (panel A) and implemented policies (panel B) by treatments. Statistical tests reported below are based on a multinomial logistic regression of the policy voted for based on treatment and standard errors clustered on groups. In the *CK-NoMan* treatment, policy makers exhibit high trust in the market, voting for on of the policies $\mathscr{X}$, $\mathscr{Y}$ or $\mathscr{Z}$ in over 97% of the time. As prices fully reveal the true state, the policy makers are able to learn from the market, with close to 90% of the votes cast for the True policy, which is consequently implemented in 93% of all markets.

Result 4. *When policy makers know that the market is free of manipulators, they trust the market, and are able to implement the True policy with high probability. Conversely, uncertainty regarding the existence of manipulators substantially impedes policy decisions—even when there are no manipulators in the market.*

Common knowledge of active manipulators has a substantial effect on policy making, with the True policy implemented in only around three quarters of all markets in the *NCK-NoMan* treatment, a decrease of 16.1 percentage points ompared to *CK-NoMan* ($p = 0.065$). The difference is mostly due to an increase of 10.7 percentage points in status quo votes ($p = 0.045$), but also a nonsignificant increase of 5.4 percentage points in votes cast for the Fake and Neutral

[28] The difference between M^F and M^P in *NCK-NoMan* remain nonsignificant in higher-powered regression analysis with markets as the unit of observation and fixed effects for sessions.

A. Votes.

B. Implemented policies.

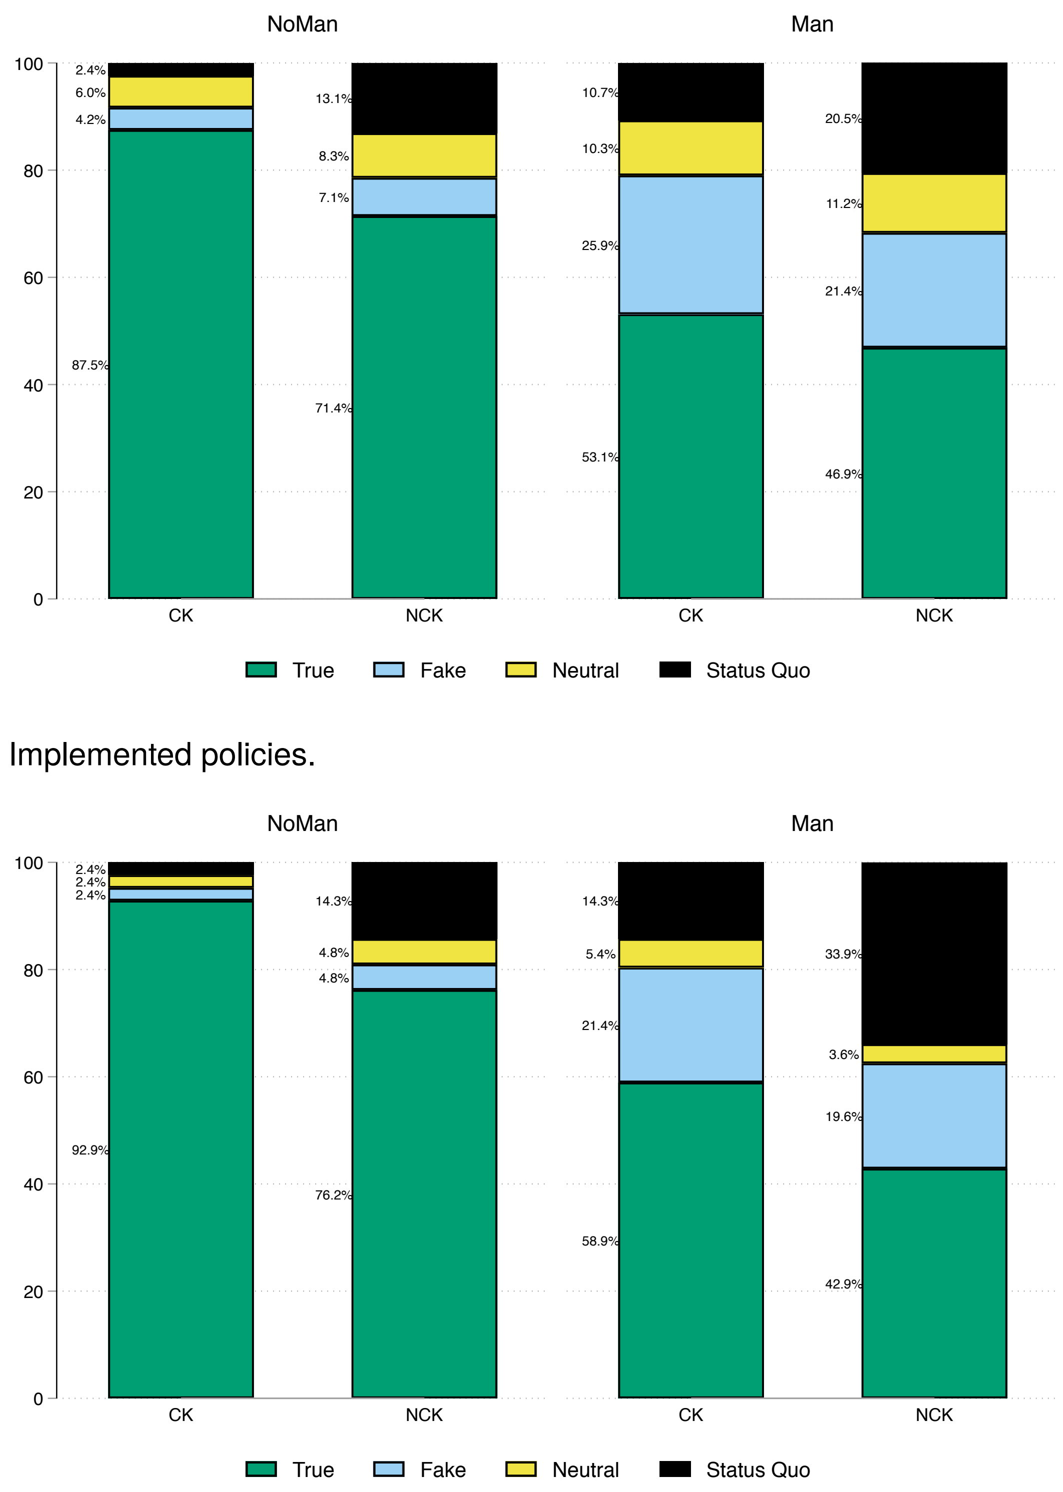

Figure 3: Distributions of votes and implemented policies.

policies. The effect of (lack of) common knowledge on policy making could be attributed to the higher variance of market prices evident in Figure 1, but may also arise from the erosion of trust in the market. We explore this issue in Section 4.3. The following result summarizes the results regarding policy making without manipulators:

As we saw, manipulators were successful in undermining the information aggregation properties of the market. This result carries over to the voting stage, with only around half of the votes cast for the True policy in the *Man* compared to the *NoMan* treatments ($p < 0.001$, separately by common knowledge or combined), and an increase of roughly 15–20 percentage points in votes cast for the Fake policy ($p = 0.002$ with common knowledge and $p = 0.027$ without). Lack of common knowledge appears to lower trust in the market, almost doubling the share of status quo votes in the *NCK-Man* treatment compared to the *CK-Man* treatment, albeit not significantly ($p = 0.171$). Accordingly, the status quo policy was implemented in one third of the markets in the *NCK-Man* treatment compared to one in seven markets in the *CK-Man* treatment ($p = 0.103$).

Once more, we may ask to what extent the effects of manipulators and common knowledge on voting behavior are mediated by the level of information aggregation in market prices, and to what extent these effects are due to engendered mistrust in the market. To address these questions, we now turn to an analysis of the voting strategies. First, we state the result concerning the effect of manipulators on policy:

Result 5. *Manipulators are successful in manipulating around 25% of the votes. Uncertainty about the existence of manipulators leads to less trust in the market, as reflected in more votes cast for the status quo policy in the no common knowledge of manipulators (NCK) markets.*

4.3. Voting strategies

In analyzing the voting strategies, we consider whether policy makers vote in line with the observed market prices. To do so, write S_1 and S_2 for the securities with the highest and second-highest market closing prices, respectively.[29] We categorize all votes into three categories accordingly:

Following the market. Voting for the policy corresponding to S_1 security.

[29]If the market closing prices of two securities are equal, we break the tie according to the average prices in the last ten (rather than five) transactions.

Table 4: Shares of policy makers who follow, oppose and ignore the market.

	CK-NoMan	NCK-NoMan	CK-Man	NCK-Man
Follow the market	89.3%	75.6%	58.9%	58.0%
	(4.4%)	(6.4%)	(6.8%)	(6.3&)
Oppose the market	8.3%	11.3%	30.4%	21.4%
	(3.6%)	(3.1%)	(5.0%)	(5.8%)
Ignore the market	2.4%	13.1%	10.7%	20.5%
	(1.2%)	(5.2%)	(3.1%)	(6.4%)
n	168	168	224	224

Note. Robust standard errors clustered on groups based on a multinomial logistic regression in parentheses.

Opposing the market. Voting for the policy that is not associated with the S_1 security.

Ignoring the market. Voting for the status quo.

Table 4 reports the proportion of policy makers in each treatment who follow, oppose and ignore the market. All statistical tests reported in this section are based on a multinomial logistic regression predicting the vote category based on the treatment with standard errors clustered on groups.

Policy makers must choose a voting strategy based on the observed market activity on the one hand, and their trust in the market on the other hand. Trust, in turn, is influenced by the observed market activity and by the policy makers' prior information regarding manipulation. Given the infinite trading profiles in continuous time, the full strategy space is non-tractable. To estimate the extent to which policy makers can extract information from the market, we therefore consider the expected payoffs obtained if all policy makers follow a *simple heuristic* based on the ability of closing market prices to differentiate between the securities.

Let P_1 and P_2 be the corresponding market closing prices of securities S_1 and S_2, respectively. The heuristic, henceforth α*-strategy*, takes one parameter, $0 \leq \alpha \leq 1$, which can be interpreted as the degree to which the voter trusts the market. An α-strategy dictates following the market if and only if $P_2/P_1 \leq \alpha$, and otherwise ignore the market and vote for the status quo.[30] Note that $\alpha = 0$ implies always ignoring the market, for the status quo payoff of 100 ECU. As α increases, the policy maker trusts the market more, and is willing to follow the

[30]For example, if the closing prices of securities x, y and z are 10, 4, and 2 ECU, respectively, then an α-strategy dictates voting for policy $\mathscr{X}$ for any $\alpha \leq 0.4$, and vote for the status quo $\mathcal{Q}$ otherwise.

market for lower price differentiations. At the upper end, we have full trust in the market at $\alpha = 1$, for which the policy maker always follows the market.[31]

Figure 4 plots the mean payoff (red circles) obtained by the policy makers in each treatment if all follow an α-strategy with the parameter α varying from zero to one in steps of 0.05. In the *NCK* treatments, it is not clear ex-ante whether the policy makers are able to distinguish whether there are manipulators in the market. The figure therefore includes separate panels by manipulation treatments as well as a panel showing the combined results. The solid vertical line marks the mean payoff obtained in a treatment by determining the payoff of each policy maker based solely on her own actual vote (i.e., the payoff a policy maker would receive if her vote were always pivotal). Finally, the dashed vertical lines mark the status quo payoff of 100 ECU.

In the *CK-NoMan* treatment, we see that even a little trust in the market can lead to substantial gains. The highly conservative strategy of following the market only if P_1 is at least 20 times larger than P_2 (i.e., $\alpha = 0.05$) leads to a mean payoff of 188 ECU, almost twice the status quo payoff. Moderate to high trust in the market ($\alpha \geq 0.5$) yields mean payoffs of more than 300 ECU, maximized at full trust in the market ($\alpha = 1$), with a mean payoff of 343 ECU.[32] Policy makers indeed trust the market, following the market in 89.3% of cases, ignoring the market in 2.4% of cases and opposing the market in only 7.7% of cases, for a mean actual-vote payoff of 312 ECU.

The comparison to the *NCK-NoMan* treatment is illuminating. We see that—despite the lessened information aggregation in prices—prices are highly informative, with α-strategy payoffs as high as obtained in the *CK-NoMan* treatment. Blindly following the market (i.e., $\alpha = 1$) yields a high payoff of 374 ECU, not much less than the 400 ECU policy makers could obtain if they knew the True state for certain! Nevertheless, the actual voting behavior reveals low trust in the market. As noted above, the share of status quo votes (ignoring the market) increases from 2.4% in the *CK-NoMan* treatment to 13.1% in the *CK-NoMan* treatment ($p = 0.045$). The share of policy makers who oppose the market also increases, from 8.3% in *CK-NoMan* to 11.3%in *NCK-NoMan*, although the difference is not statistically significant ($p = 0.532$). Consequently, the pay-

[31]To confirm that the family of α-strategies plausibly approximates actual voting strategies, we use a logistic regression to predict whether the policy makers vote for the policy corresponding to S_1 on P_2/P_1 and the treatment, with robust standard errors clustered on groups. The coefficients for P_2/P_1 are highly significant ($p < 0.001$) for all four treatments, and take values between -0.359 and -0.433.

[32]The payoff curve flattens above $\alpha = 0.5$ because the price ratio P_2/P_1 mostly falls below 0.5 in the *CK-NoMan* treatment markets.

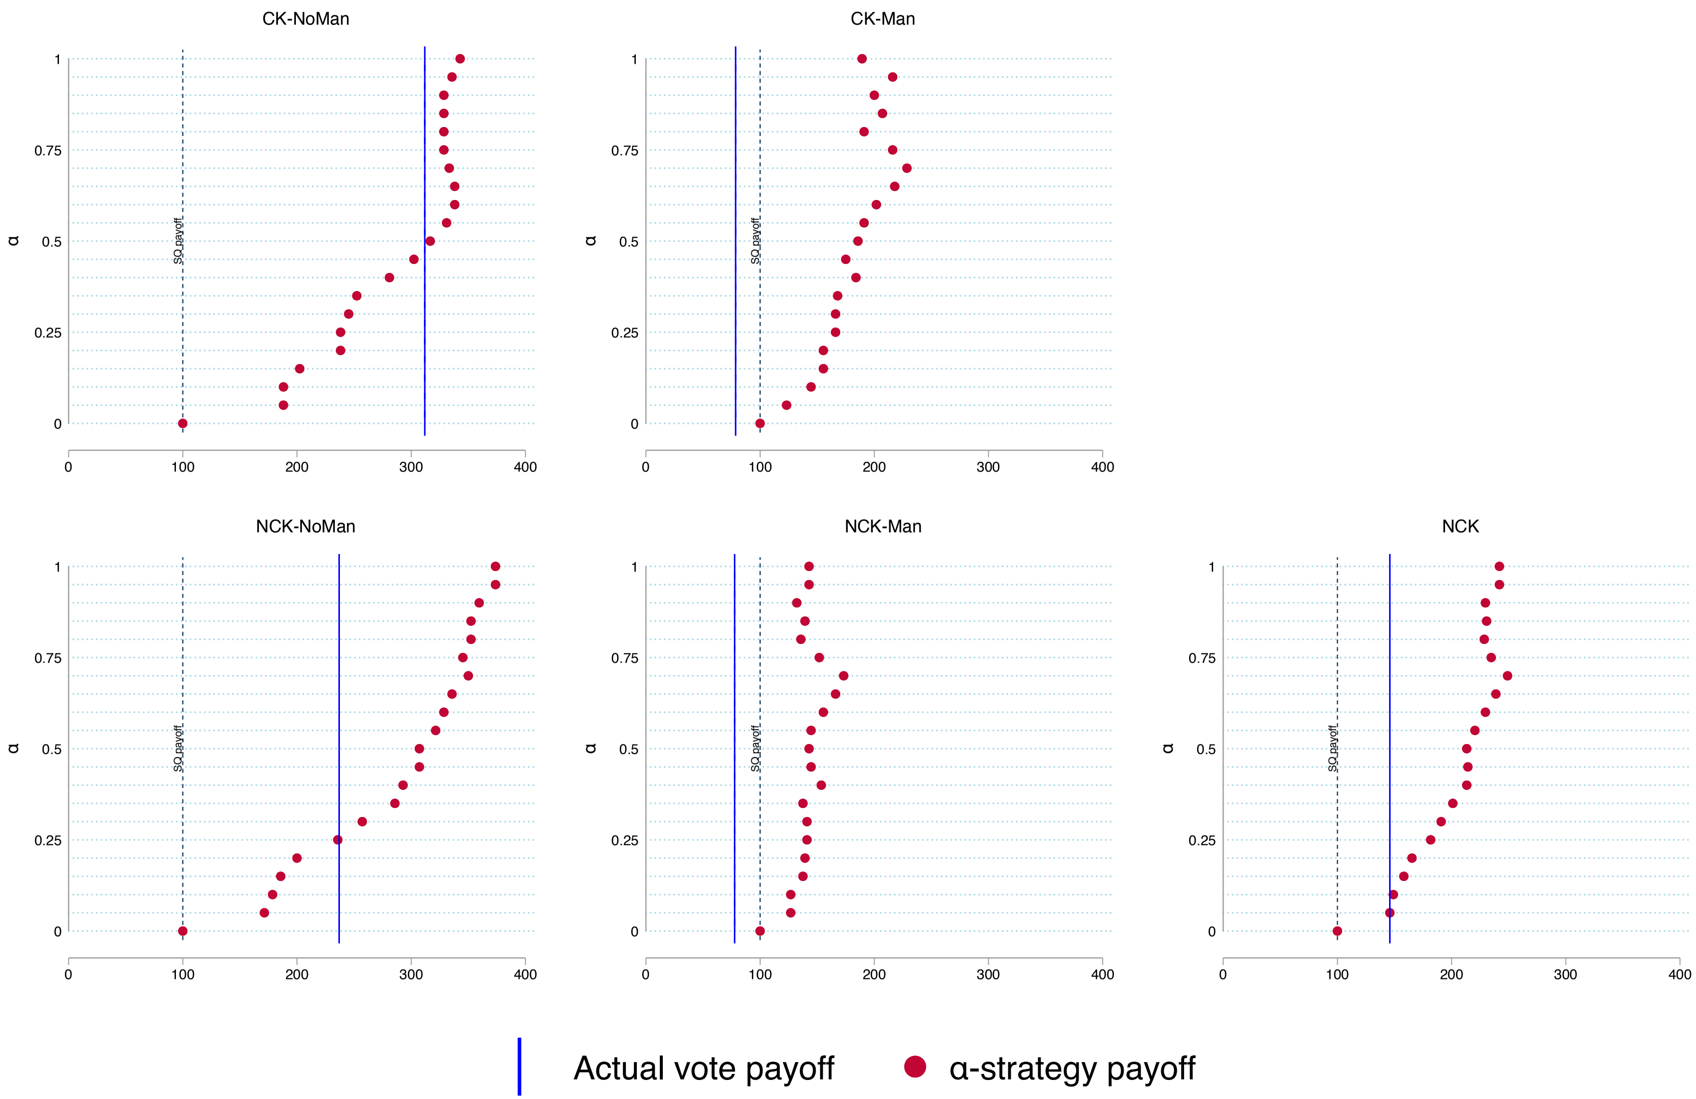

Figure 4: Expected payoff for decisive votes.

off based on actual votes is only 237 ECU, only 63% of the payoff obtainable by trusting the market fully, and a payoff comparable to the 235 ECU obtained from following the market only if P_1 security is at least *four* times larger than P_2 ($\alpha = 0.25$). We therefore conclude that the suboptimal voting observed in the previous section is not due to volatility in prices, but to mistrust in the market due to uncertainty regarding price manipulation. We first state the result regarding the ability to learn from markets without manipulators:

Result 6. *When there are no manipulators in the market, voting according to the the security with the highest closing price extracts around* $80 - -90\%$ *of the possible gains with respect to the status quo, regardless of whether the non-existence of manipulators is common knowledge.*

Next, we state the result regarding (mis)trust in the market:

Result 7. *When there is common knowledge that there are no manipulators in the market, policy makers trust the market, voting according to the observed transaction prices, and extracting most of the potential gains. Lack of common knowledge has a dramatic effect on trust, with many votes cast for the status quo, leading to sub optimal policy implementation and considerable loss in efficiency.*

The effect of manipulators on voting behavior we saw in the previous section is clearly evident in the middle column of Figure 4. The figure reveals that there are two separate effects in play. First, as can be expected based on the analysis of market prices, the amount of information in the market is substantially diminished with manipulators. Perhaps surprisingly, we see that there is nonetheless still much to gain from trusting the market. In the *CK-Man* treatment, the payoffs for high enough trust ($0.5 \leq \alpha \leq 1$) are in the range of 180 to 230 ECU. The situation is considerably worse when the other traders are not explicitly informed of the existence of manipulators in the market. The corresponding payoffs in the *NCK-Man* treatment are in the lower range of 130 to 175 ECU, though still substantially above the status quo payoff of 100 ECU.

The second effect is observed when comparing the actual-vote payoffs, which are below the status quo payoff in both the *CK-Man* and *NCK-Man* treatments . That is, policy makers not only forgo the potential gains from *trusting* the market, but are even doing *worse* than they would by always *ignoring* the market. This implies that knowledge or suspicion of manipulation leads policy makers to oppose the market, even though high price differentiation typically indicates that the market was successful in reflecting the true state of the world. For example, when $P_2/P_1 \geq 0.25$ (i.e., P_1 is at least four times larger than P_2), S_1 is the True security in 56 of 60 (93.3%) markets in *CK-Man* and in 44 of 52 (84.6%) markets in *NCK-Man*.

Indeed, whereas 89.3% of the votes in the *CK-NoMan* treatment, and 75.6% of the votes in *NCK-NoMan* treatments go to the policy corresponding to the S_1 security, these shares drop to 58.9% and 58.0% in the *CK-Man* and *NCK-Man* treatments, respectively ($p < 0.001$ for the separate and combined comparisons). The share of status quo votes significantly increases from 2.4% in the *CK-NoMan* treatment to 10.7% in the *CK-Man* treatment ($p = 0.004$) and from 13.1% in the *NCK-NoMan* treatment to 20.5% in *NCK-Man* treatment ($p < 0.07$). The share of policy makers opposing the market also increases significantly, from 8.3% in the *CK-NoMan* treatment to 30.4% in the *CK-Man* treatment ($p < 0.001$) and from 11.3% in the *NCK-NoMan* treatment to 21.4% in the *NCK-Man* treatment ($p = 0.008$).

Finally, the considerable difference in voting strategies between the *NCK-NoMan* and *NCK-Man* treatments shows that the market activity provides enough information for policy makers to figure out (to a large extent) whether there are manipulators in the market, and respond to the price ratio accordingly. Nonetheless, we can ask how a simple heuristic that conditions only on the price ratio and ignores all other market information fares. The bottom right panel in Figure 4 shows that such a heuristic can yield substantial gains, with full trust in the market yielding a payoff of 242 ECU. Actual behavior, however, reveals very low trust in the market, potentially yielding a payoff comparable to that obtained with $\alpha = 0.05$.

Result 8. *Manipulators affect policy making via two channels. First, they manipulate market activity sufficiently to obscure the information reflected in market prices, though not sufficiently to completely eliminate the advantage in following the market. Second, policy makers who know or suspect manipulation tend to ignore or even vote against the market, and are therefore unsuccessful in utilizing the information conveyed by the market prices.*

5. Conclusion

Motivated by advancements in the study of information aggregation in markets over the last few decades, many researchers and policy makers advocate the use of markets in guiding policy decisions. This raises the necessity of better understanding how invested parties may be able to misuse the market in order to distort information and influence policy.

Prediction markets proved successful in predicting real-world events (Wolfers and Zitzewitz, 2004). Accordingly, we constructed experimental markets where prices fully reveal the aggregate information in the market. Our design provides a new understanding of the potential effect of manipulators and the roll of uncertainty in information aggregation, and—perhaps more importantly—in fostering trust or mistrust in the market.

The results highlight the importance of trust in the market. When the market designer cannot guarantee that the market is free of manipulators, trade volatility increases, however the market prices still provide ample information for policy makers to reach close to optimal decisions. Nonetheless, policy makers lose trust in the market, leading to substantial loss of potential gains.

This result is reflected in the manipulation markets. While manipulators are able to manipulate the trading activity considerably, the market prices still reflect sufficient information to support socially beneficial policy making. Policy makers who are aware of the manipulation, however, mistrust the market, and fail to utilize the information conveyed in the prices. Uncertainty has two effects. On the one hand, as traders are less certain in how to interpret market activity, manipulators are better able to obscure information in the market. On the other hand, policy makers mistrust the market less. The combined effects lead to similar suboptimal policy making in the face of manipulation with and without common knowledge of the manipulation.

These findings bear important implications for the design of prediction markets and for policy making based on natural observations in financial markets. We find that even markets that have the capability of efficiently aggregating disperse information into prices are susceptible to manipulation.

Thus, if decision makers wish to use markets to guide policy, there is a need to set up dedicated private prediction markets where participation is regulated as a precaution to prevent manipulation, Such precautions, however, are not sufficient, as fostering trust in the market—for example, by increasing familiarity and experience with markets (Maciejovsky and Budescu, in press)—emerges as a necessary condition for market-based policy making to be successful.

References

Arrow, Kenneth J. et al. (2008). The promise of prediction markets. *Science* 320.(5878), pp. 877–878.

Camerer, Colin F. (1998). Can asset markets be manipulated? a field experiment with racetrack betting. *Journal of Political Economy* 106.(3), pp. 457–482.

Chen, Kay-Yut and Charles R. Plott (2002). Information aggregation mechanisms: concept, design and implementation for a sales forecasting problem. Social Science Working Paper 1131. California Institute of Technology.

Choo, Lawrence, Todd R. Kaplan, and Ro'i Zultan (2019). Information aggregation in arrow–debreu markets: an experiment. *Experimental Economics* 22, pp. 625–652.

Corgnet, Brice, Mark DeSantis, and David Porter (2015). Revisiting information aggregation in asset markets: reflective learning & market efficiency. *mimeo*.

Corgnet, Brice et al. (2018). Information (non) aggregation in markets with costly signal acquisition. *Journal of Economic Behavior & Organization*.

Cowgill, Bo and Eric Zitzewitz (2015). Corporate prediction markets: evidence from google, ford, and firm x. *The Review of Economic Studies* 82.(4), pp. 1309–1341.

Deck, Cary, Shengle Lin, and David Porter (2013). Affecting policy by manipulating prediction markets: experimental evidence. *Journal of Economic Behavior & Organization* 85, pp. 48–62.

Deck, Cary and David Porter (2013). Prediction markets in the laboratory. *Journal of Economic Surveys* 27.(3), pp. 589–603.

Dubey, Pradeep, John D. Geanakoplos, and Martin Shubik (1987). The revelation of information in strategic market games: a critique of rational expectations equilibrium. *Journal of Mathematical Economics* 16.(2), pp. 105–137.

Fama, Eugene F. (1970). Efficient capital markets: a review of theory and empirical work. *The Journal of Finance* 25.(2), pp. 383–417.

Fischbacher, Urs (2007). Z-tree: zurich toolbox for ready-made economic experiments. *Experimental Economics* 10.(2), pp. 171–178.

Gandal, Neil et al. (2018). Price manipulation in the bitcoin ecosystem. *Journal of Monetary Economics* 95, pp. 86–96.

Gillen, Benjamin J., Charles R. Plott, and Matthew Shum (2017). A pari-mutuel-like mechanism for information aggregation: a field test inside intel. *Journal of Political Economy* 125.(4), pp. 1075–1099.

Greiner, Ben (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association* 1.(1), pp. 114–125.

Hansen, Jan, Carsten Schmidt, and Martin Strobel (2004). Manipulation in political stock markets–preconditions and evidence. *Applied Economics Letters* 11.(7), pp. 459–463.

Hanson, Robin (2004). Foul play in information markets. In: *Information markets: a new way of making decisions*. Ed. by Robert W Hahn and Paul C Tetlock. Washington, D.C.: AEI-Brookings Press. Chap. 6, pp. 126–141.

Hanson, Robin and Ryan Oprea (2009). A manipulator can aid prediction market accuracy. *Economica* 76.(302), pp. 304–314.

Hanson, Robin, Ryan Oprea, and David Porter (2006). Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior and Organization* 60.(4), pp. 449–459.

Hayek, Friedrich A. (1945). The use of knowledge in society. *American Economic Review* 35.(4), pp. 519–530.

Hellwig, Martin F. (1982). Rational expectations equilibrium with conditioning on past prices: a mean-variance example. *Journal of Economic Theory* 26.(2), pp. 279–312.

Maciejovsky, Boris and David V Budescu (2013). Markets as a structural solution to knowledge-sharing dilemmas. *Organizational Behavior and Human Decision Processes* 120.(2), pp. 154–167.

— (in press). Too much trust in group decisions: uncovering hidden profiles by groups and markets. *Organization Science*.

Muth, John F. (1961). Rational expectations and the theory of price movements. *Econometrica* 29.(3), pp. 315–335.

Ostrovsky, Michael (2012). Information aggregation in dynamic markets with strategic traders. *Econometrica* 80.(6), pp. 2595–2647.

Palan, Stefan, Jürgen Huber, and Larissa Senninger (2019). Aggregation mechanisms for crowd predictions. *Experimental Economics*, pp. 1–27.

Plott, Charles R. (2000). Markets as information gathering tools. *Southern Economic Journal* 67.(1), pp. 1–15.

Plott, Charles R. and Shyam Sunder (1982). Efficiency of experimental security markets with insider information: an application of rational-expectations models. *Journal of Political Economy* 90.(4), pp. 663–698.

— (1988). Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica* 56.(5), pp. 1085–1118.

Radner, Roy (1979). Rational expectations equilibrium: generic existence and the information revealed by prices. *Econometrica* 47.(3), pp. 655–678.

Rhode, Paul W. and Koleman S. Strumpf (2004). Historical presidential betting markets. *Journal of Economic Perspectives* 18.(2), pp. 127–141.

Veiga, Helena and Marc Vorsatz (2009). Price manipulation in an experimental asset market. *European Economic Review* 53.(3), pp. 327–342.

— (2010). Information aggregation in experimental asset markets in the presence of a manipulator. *Experimental Economics* 13.(4), pp. 379–398.

Wolfers, Justin and Eric Zitzewitz (2004). Prediction markets. *Journal of Economic Perspectives* 18.(2), pp. 107–126.

Appendix A. Additional information

A.1. Transaction prices and votes.

We conducted a total of 7 sessions each for the *CK* (sessions 1, 2, 3, 4, 5, 10 and 14) and *NCK* (sessions 6, 7, 8, 9, 11, 12 and 13) treatments. Each session involved a fixed matching group of 12 participants. Figures A1-A14 detail the market transaction prices and votes in each session.

Each Figure is organised by periods along the columns (i.e., periods 1 and 14 in the leftmost and rightmost columns, respectively) and securities along the rows (i.e., the top, middle and bottom rows refer to the True, Fake and Neutral securities, respectively). The black and red lines denote the transaction prices of securities in *NoMan* and *Man* market periods, respectively, across the market duration (in seconds). Finally, the panel headers—the numbers in the grey boxes—detail the number of policy makers who voted for the True (top row), Fake (middle row) and Neutral (bottom row) policies—the omitted votes are those for the status quo.

For example, Figure A1 refers to session 1. Here, we see that the first period is a *Man* market period (i.e., the leftmost column). We see that three policy makers voted for the True policy and one policy maker voted for the status quo.

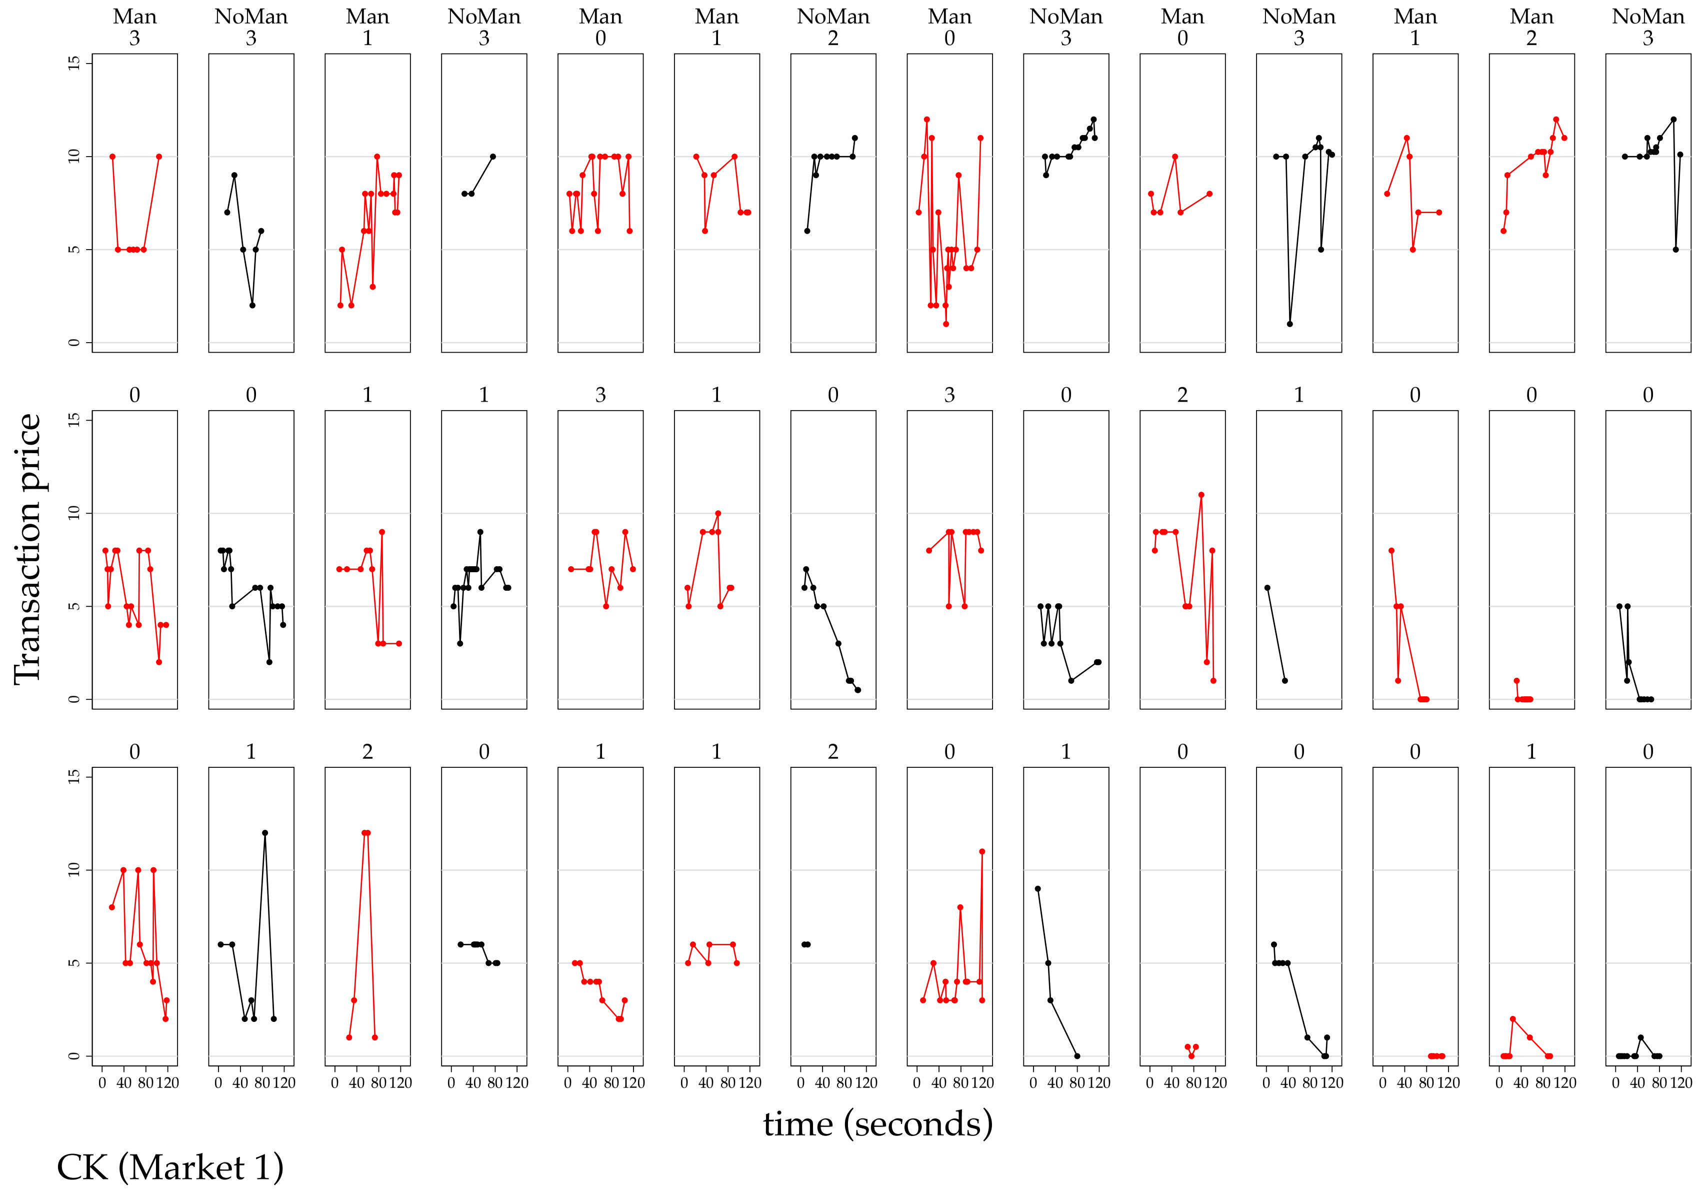

Figure A1: Session 1.

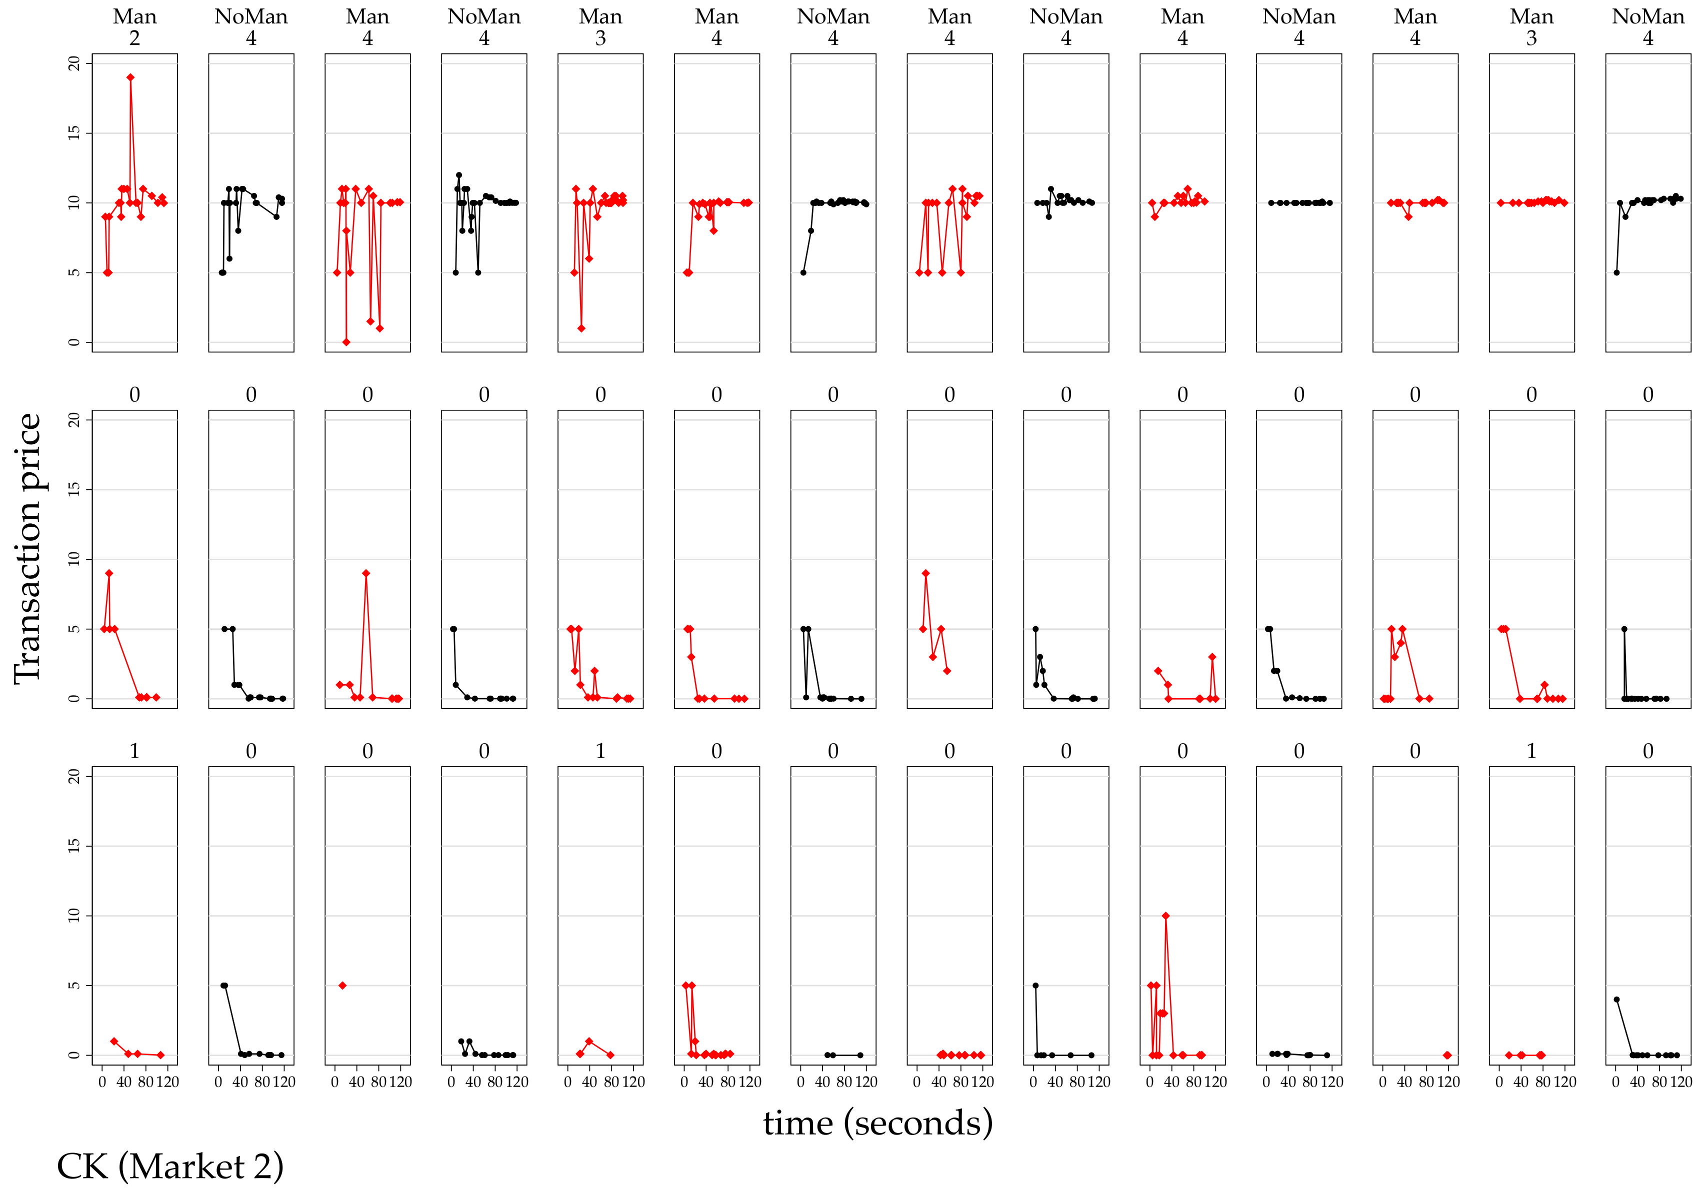

Figure A2: Session 2.

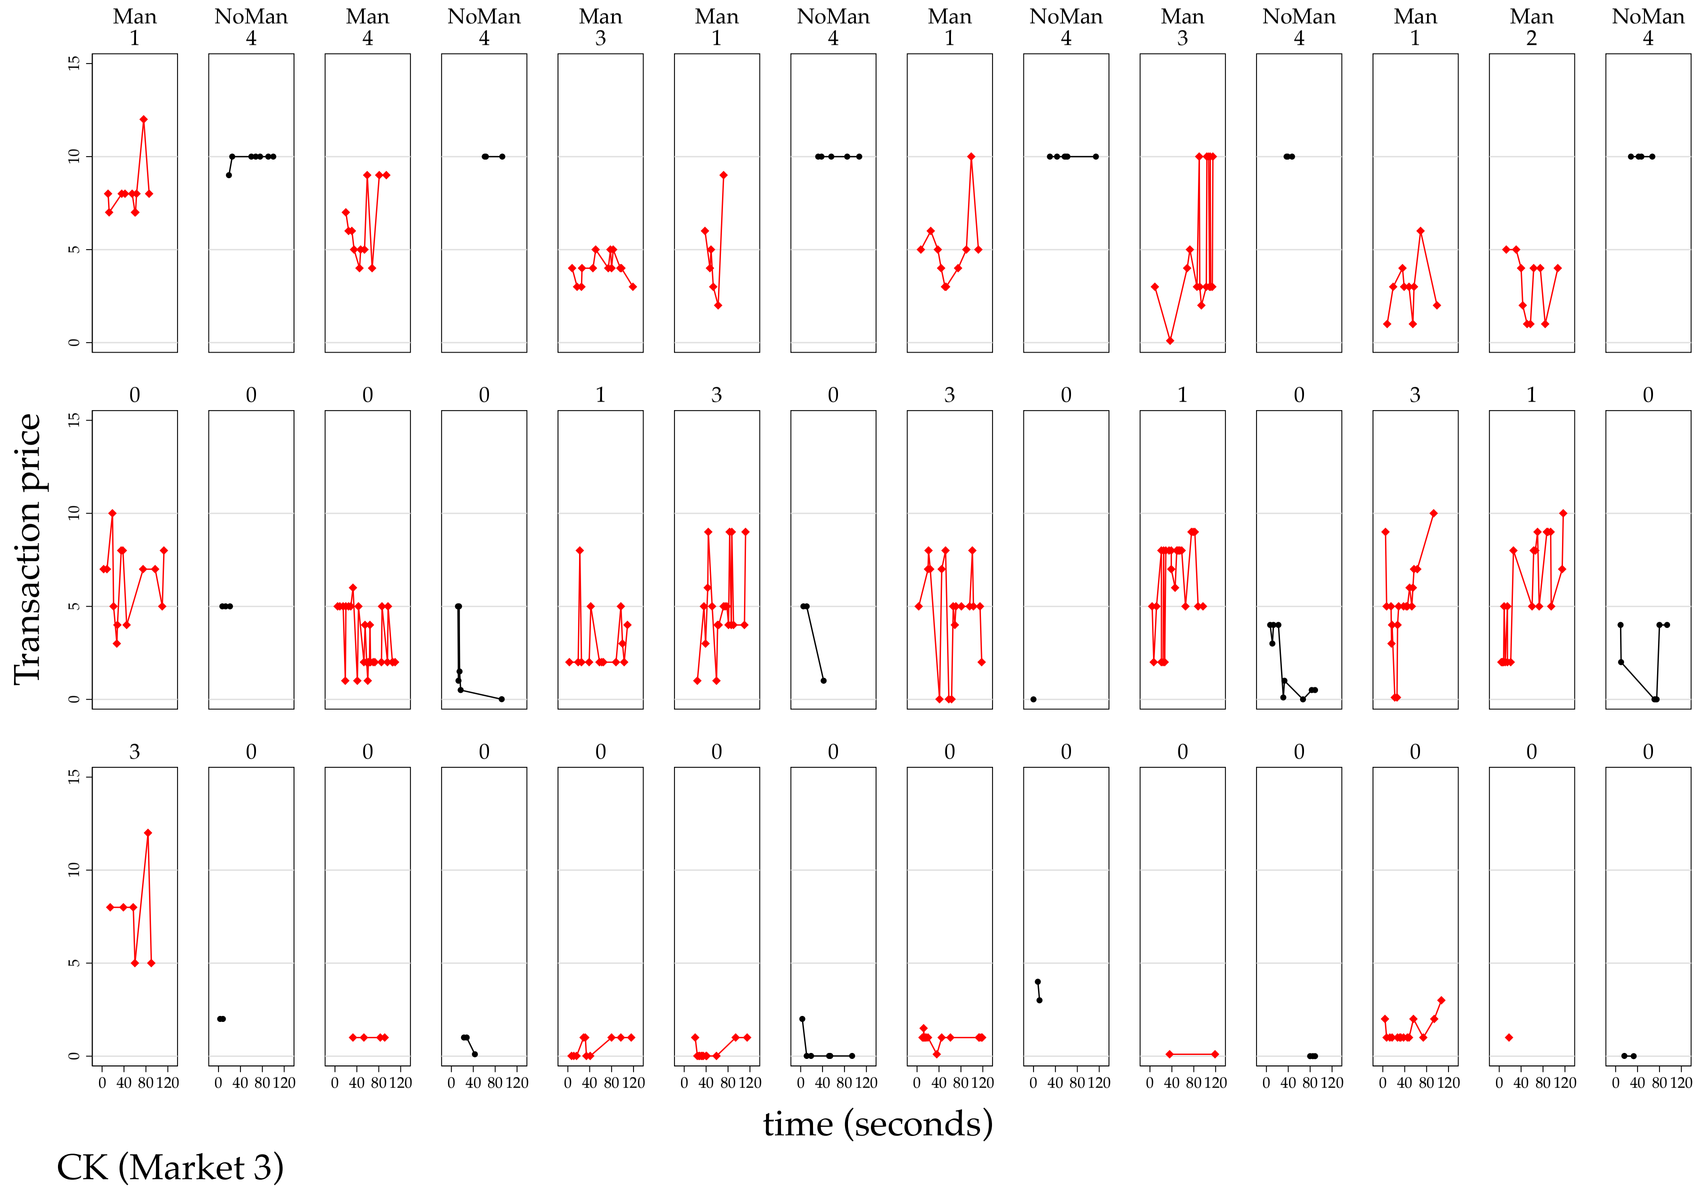

Figure A3: Session 3.

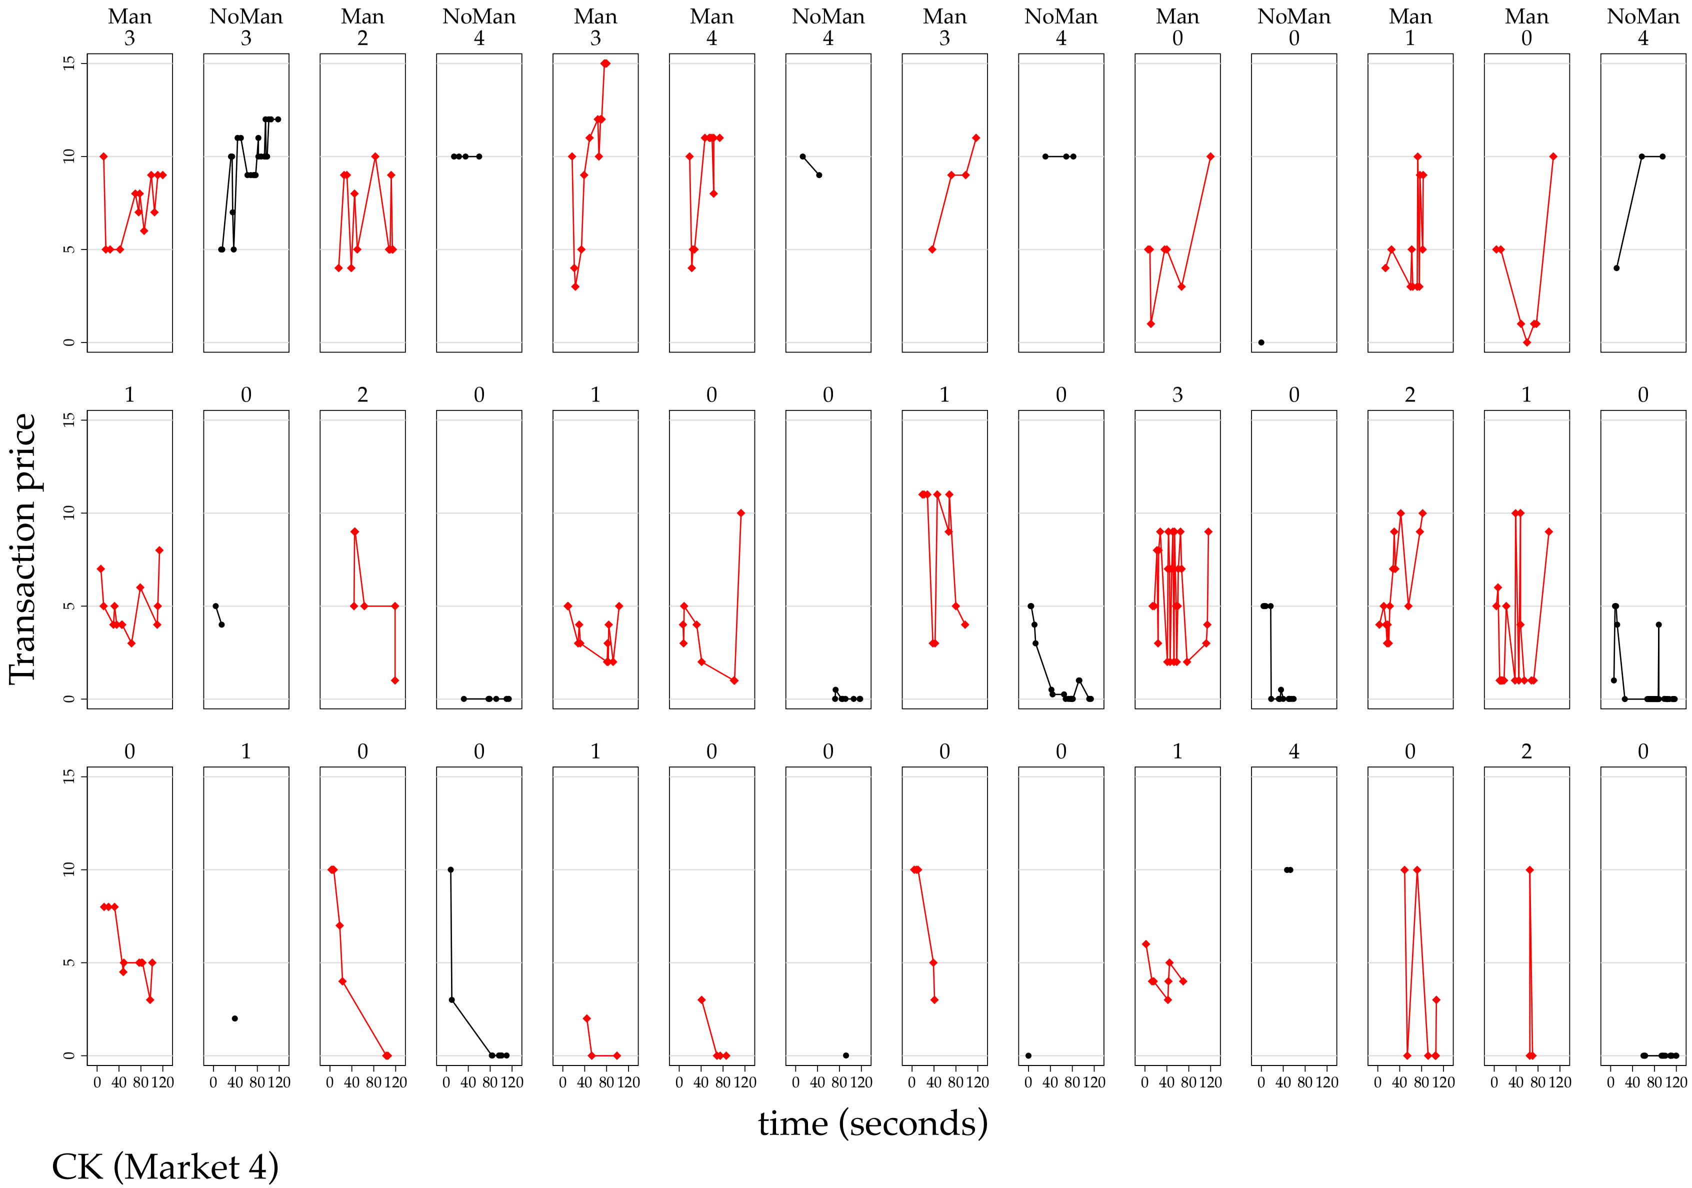

Figure A4: Session 4.

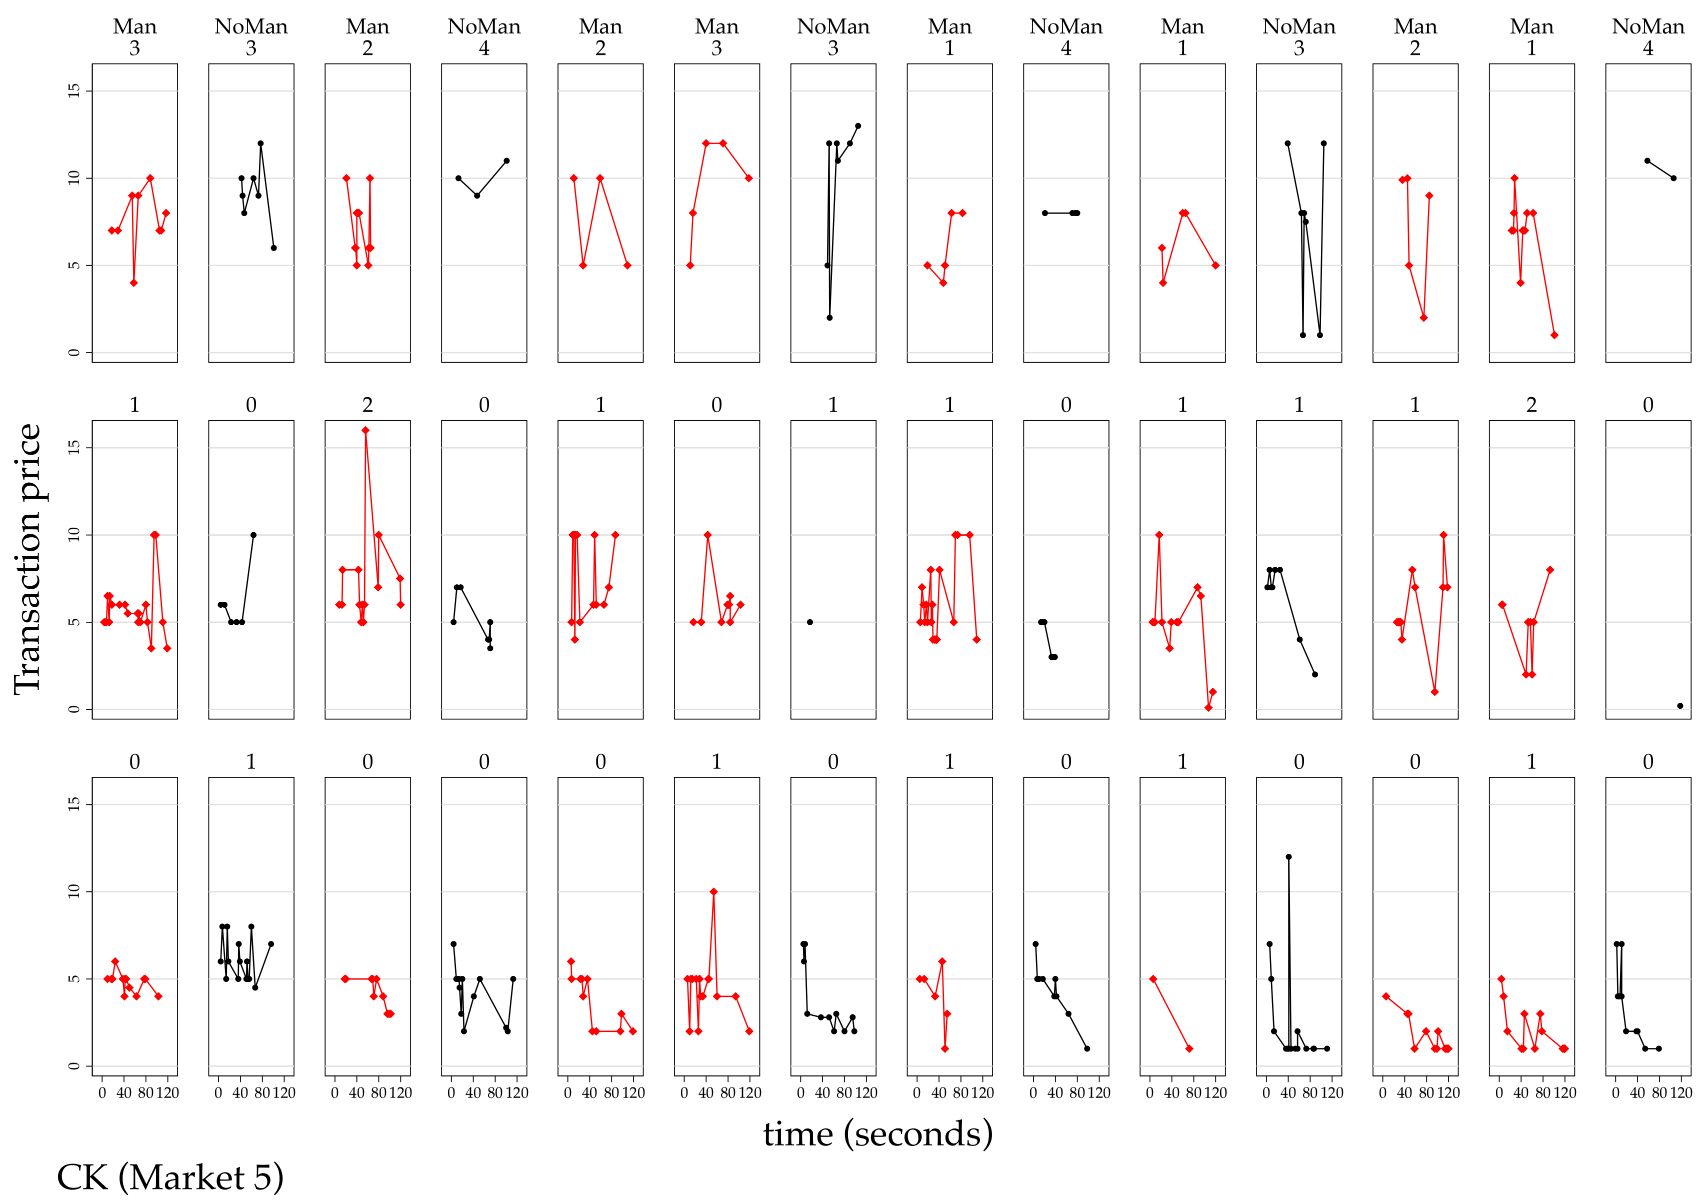

Figure A5: Session 5.

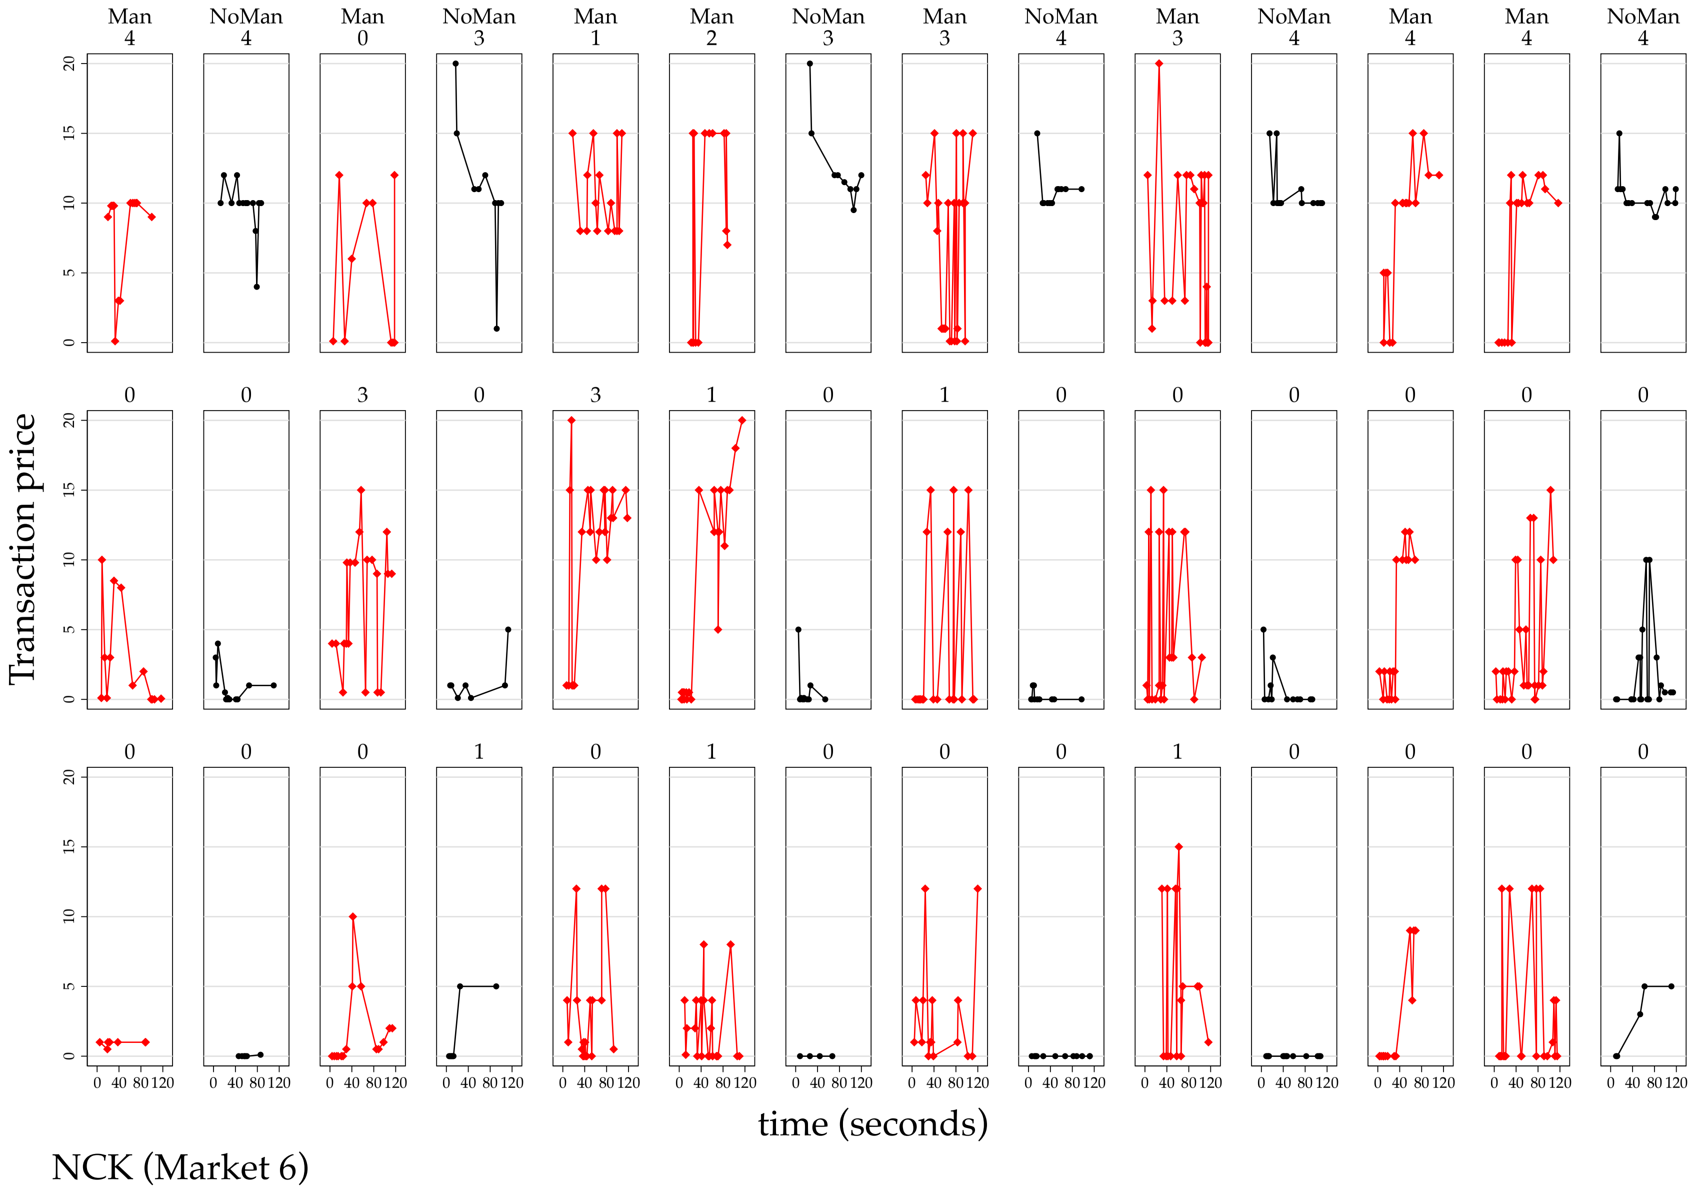

Figure A6: Session 6.

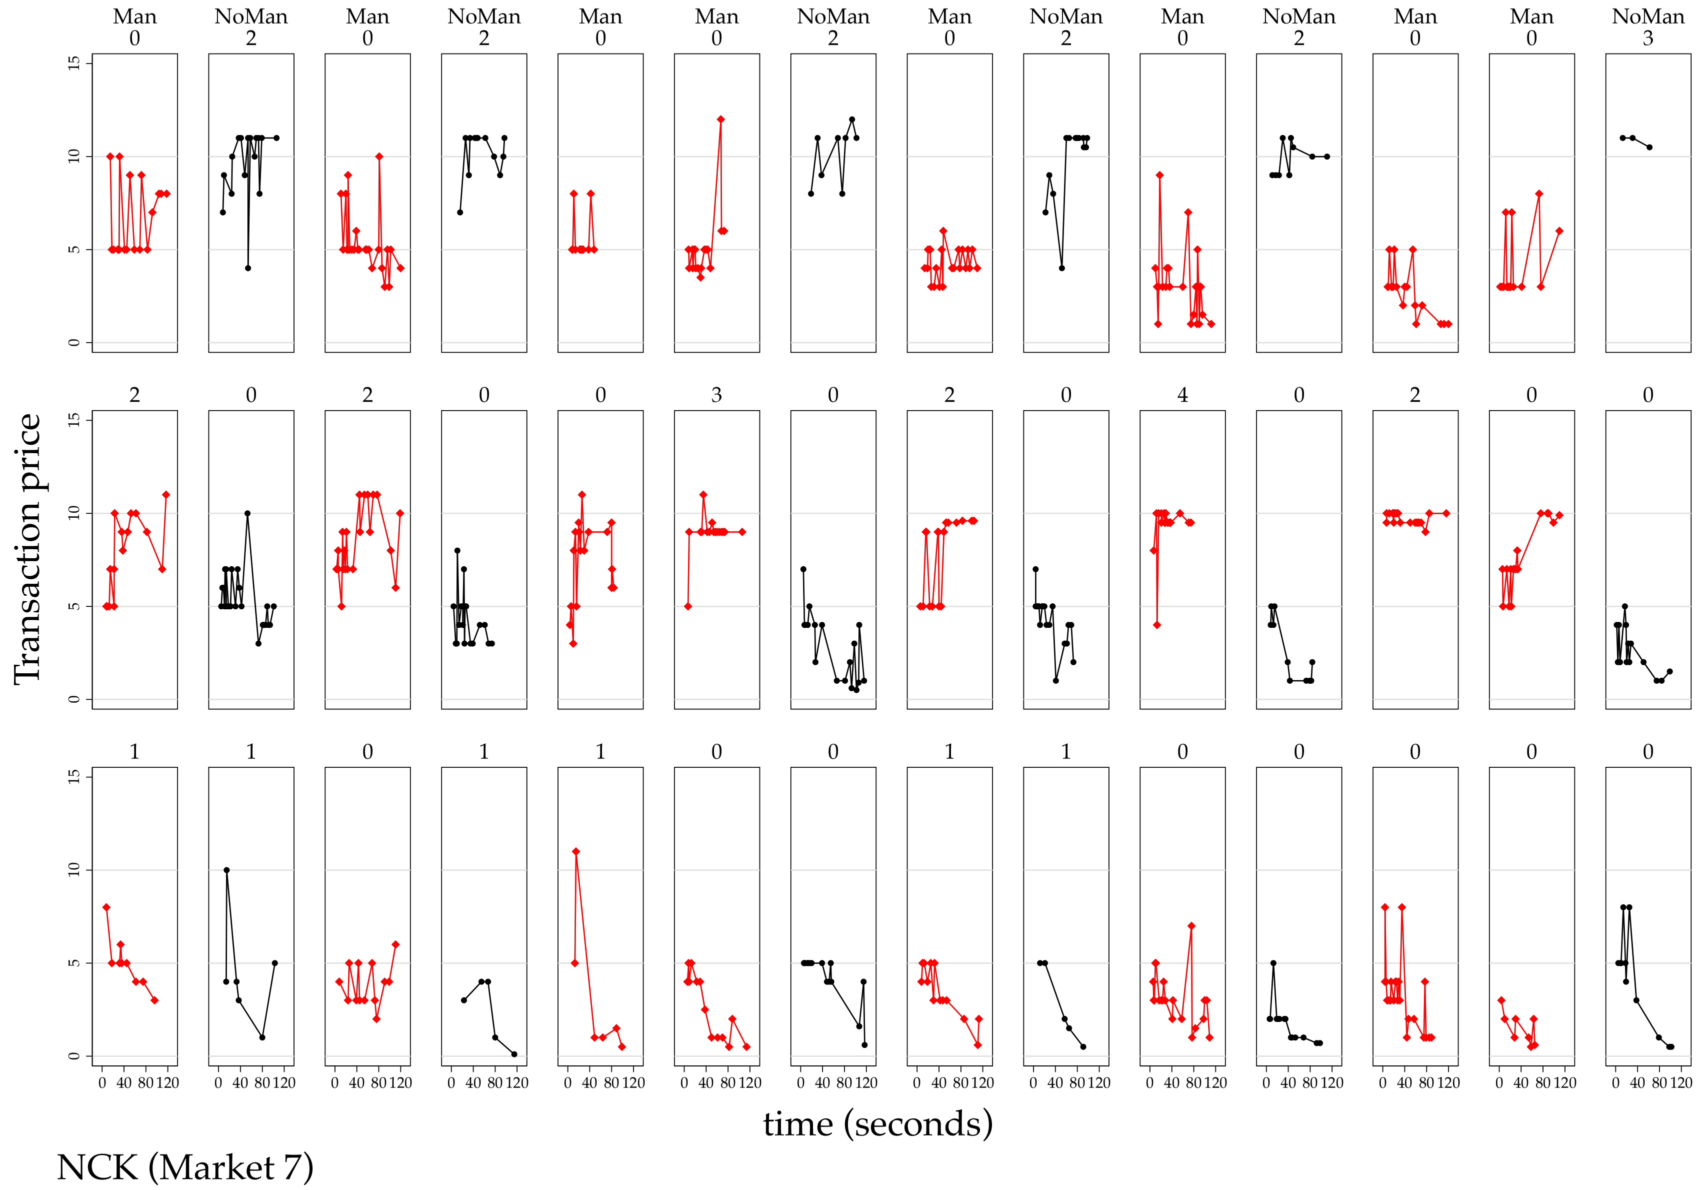

Figure A7: Session 7.

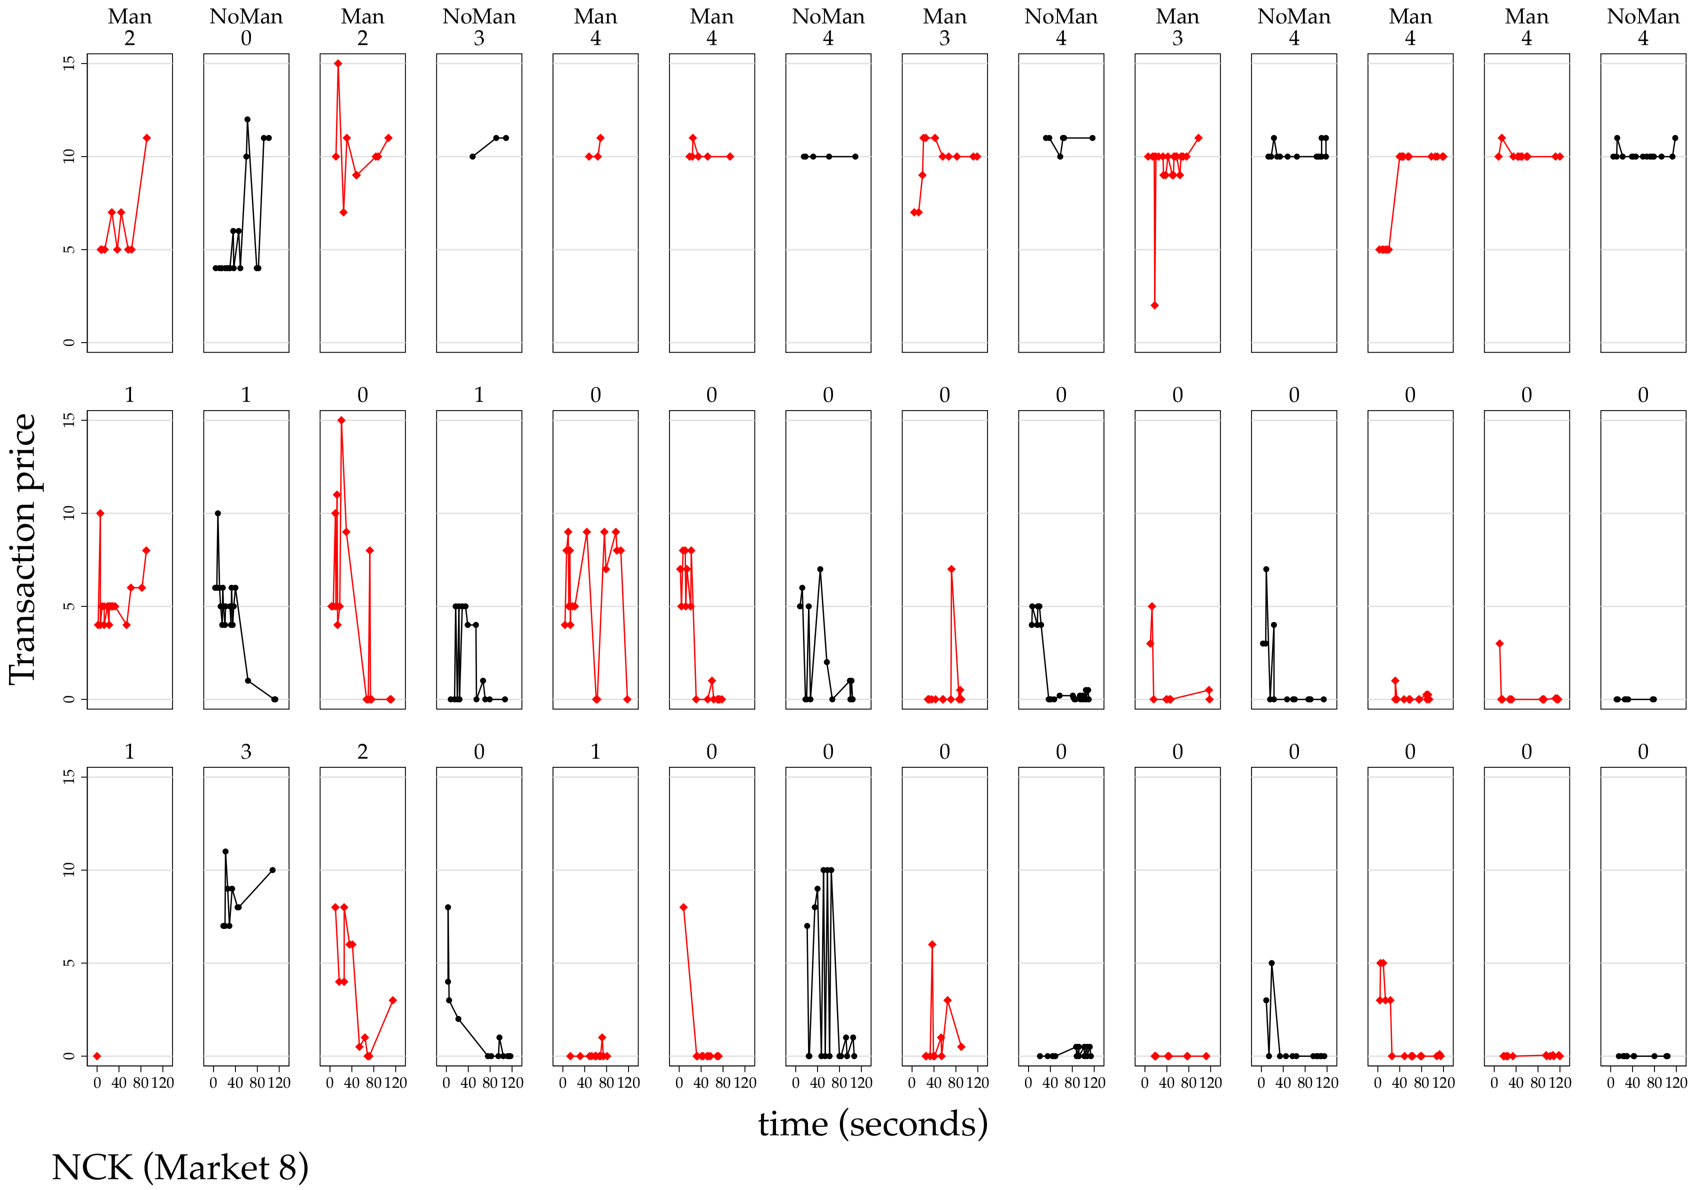

Figure A8: Session 8.

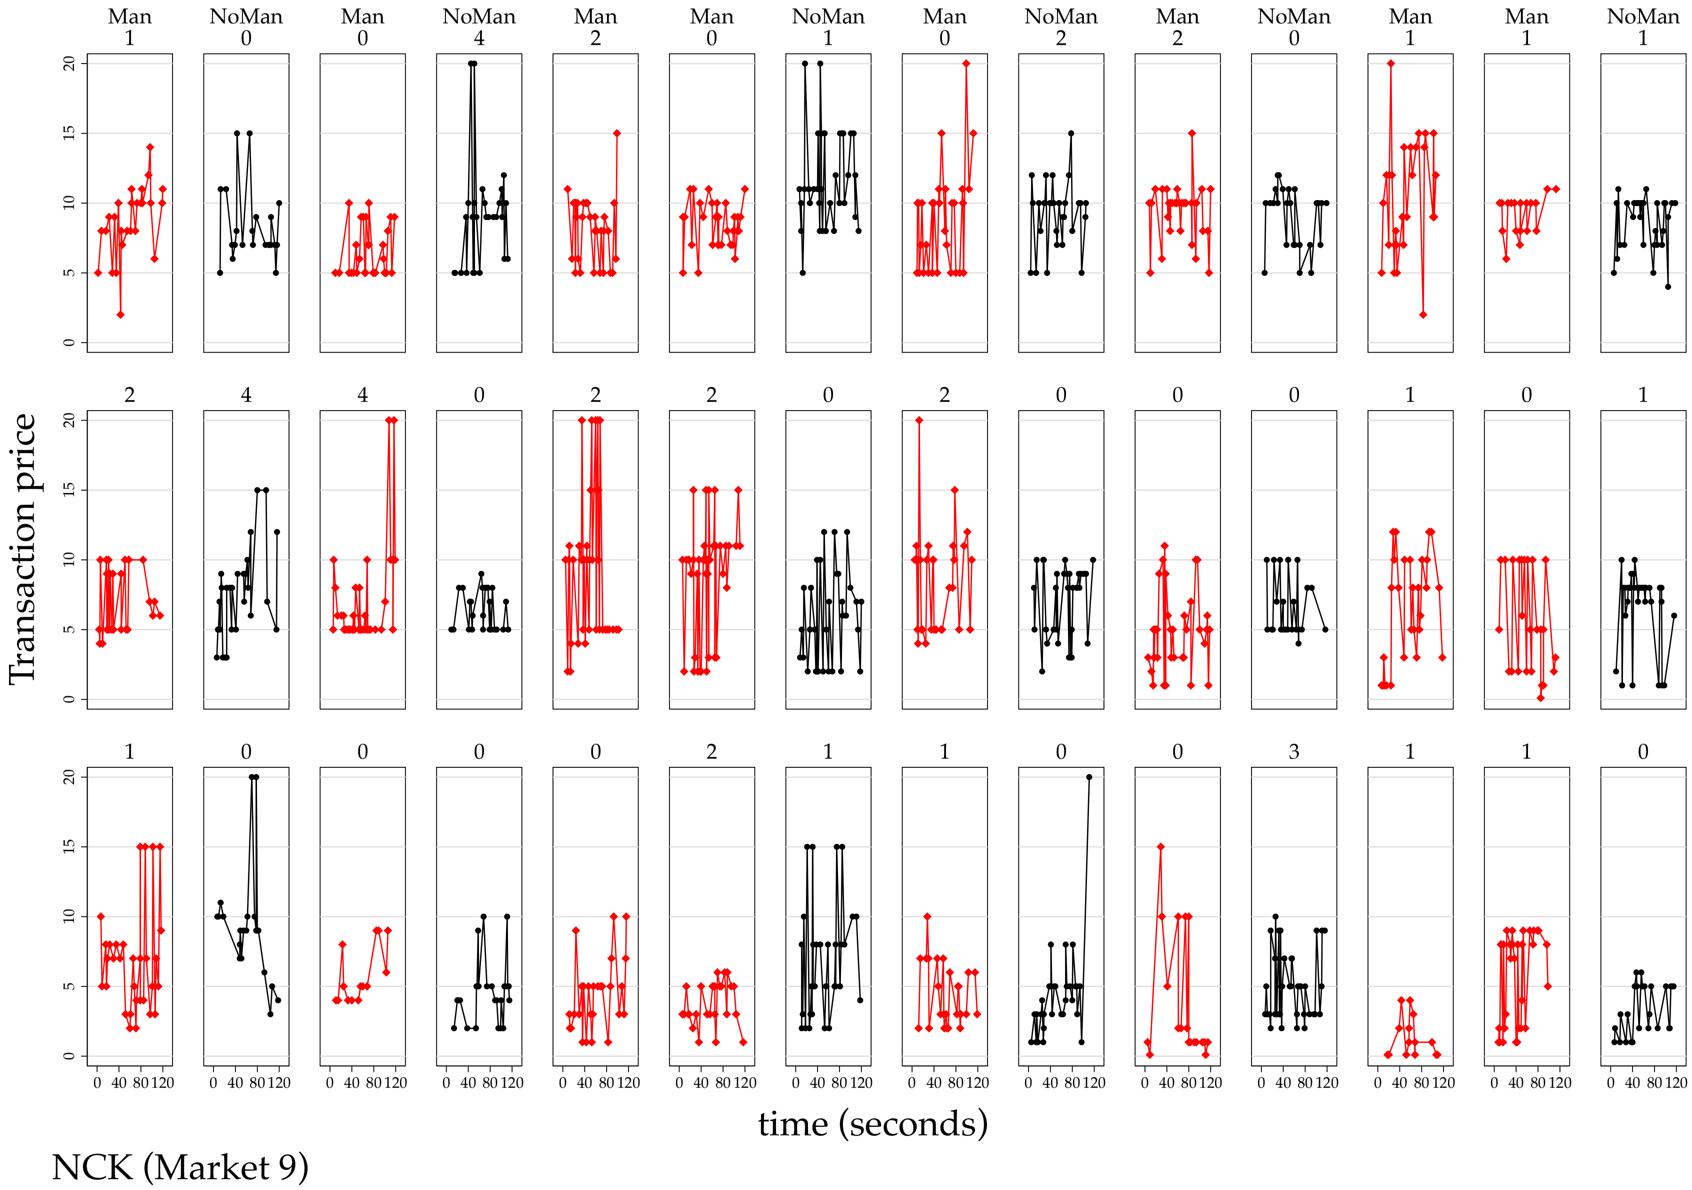

Figure A9: Session 9.

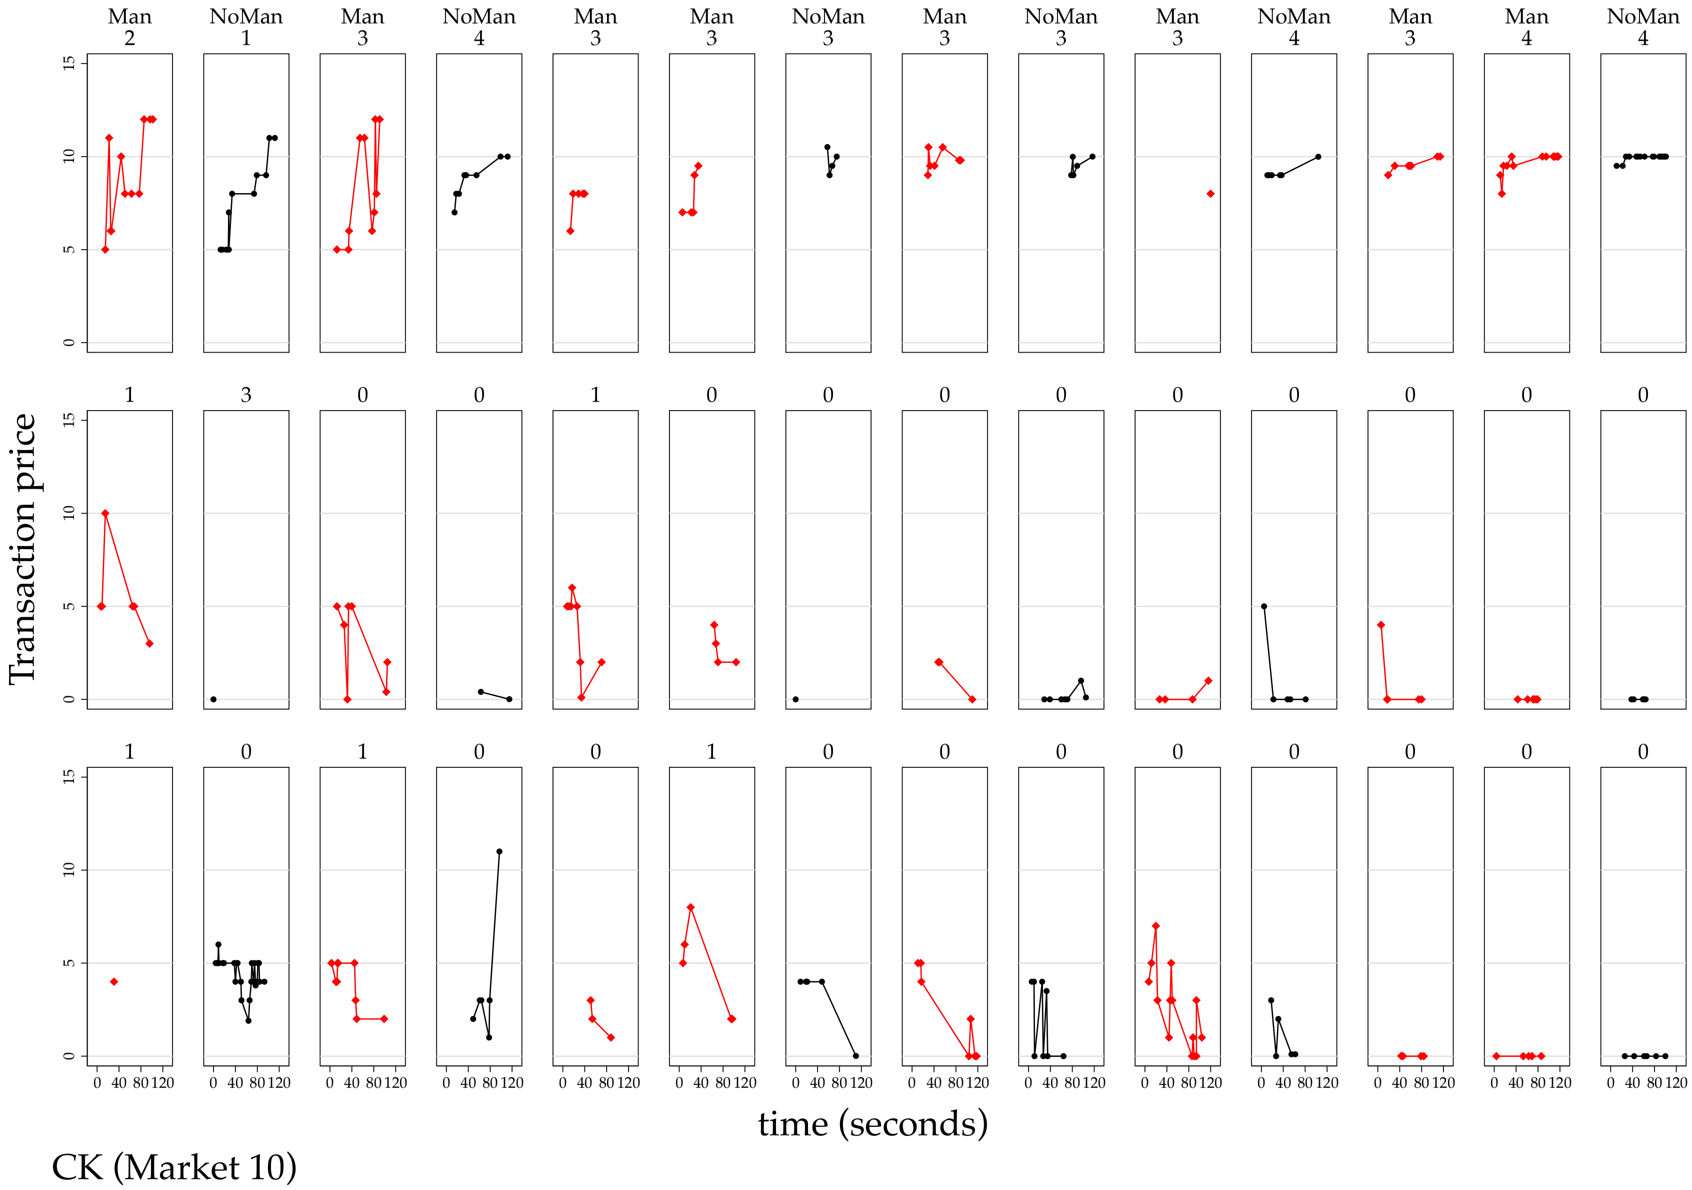

Figure A10: Session 10.

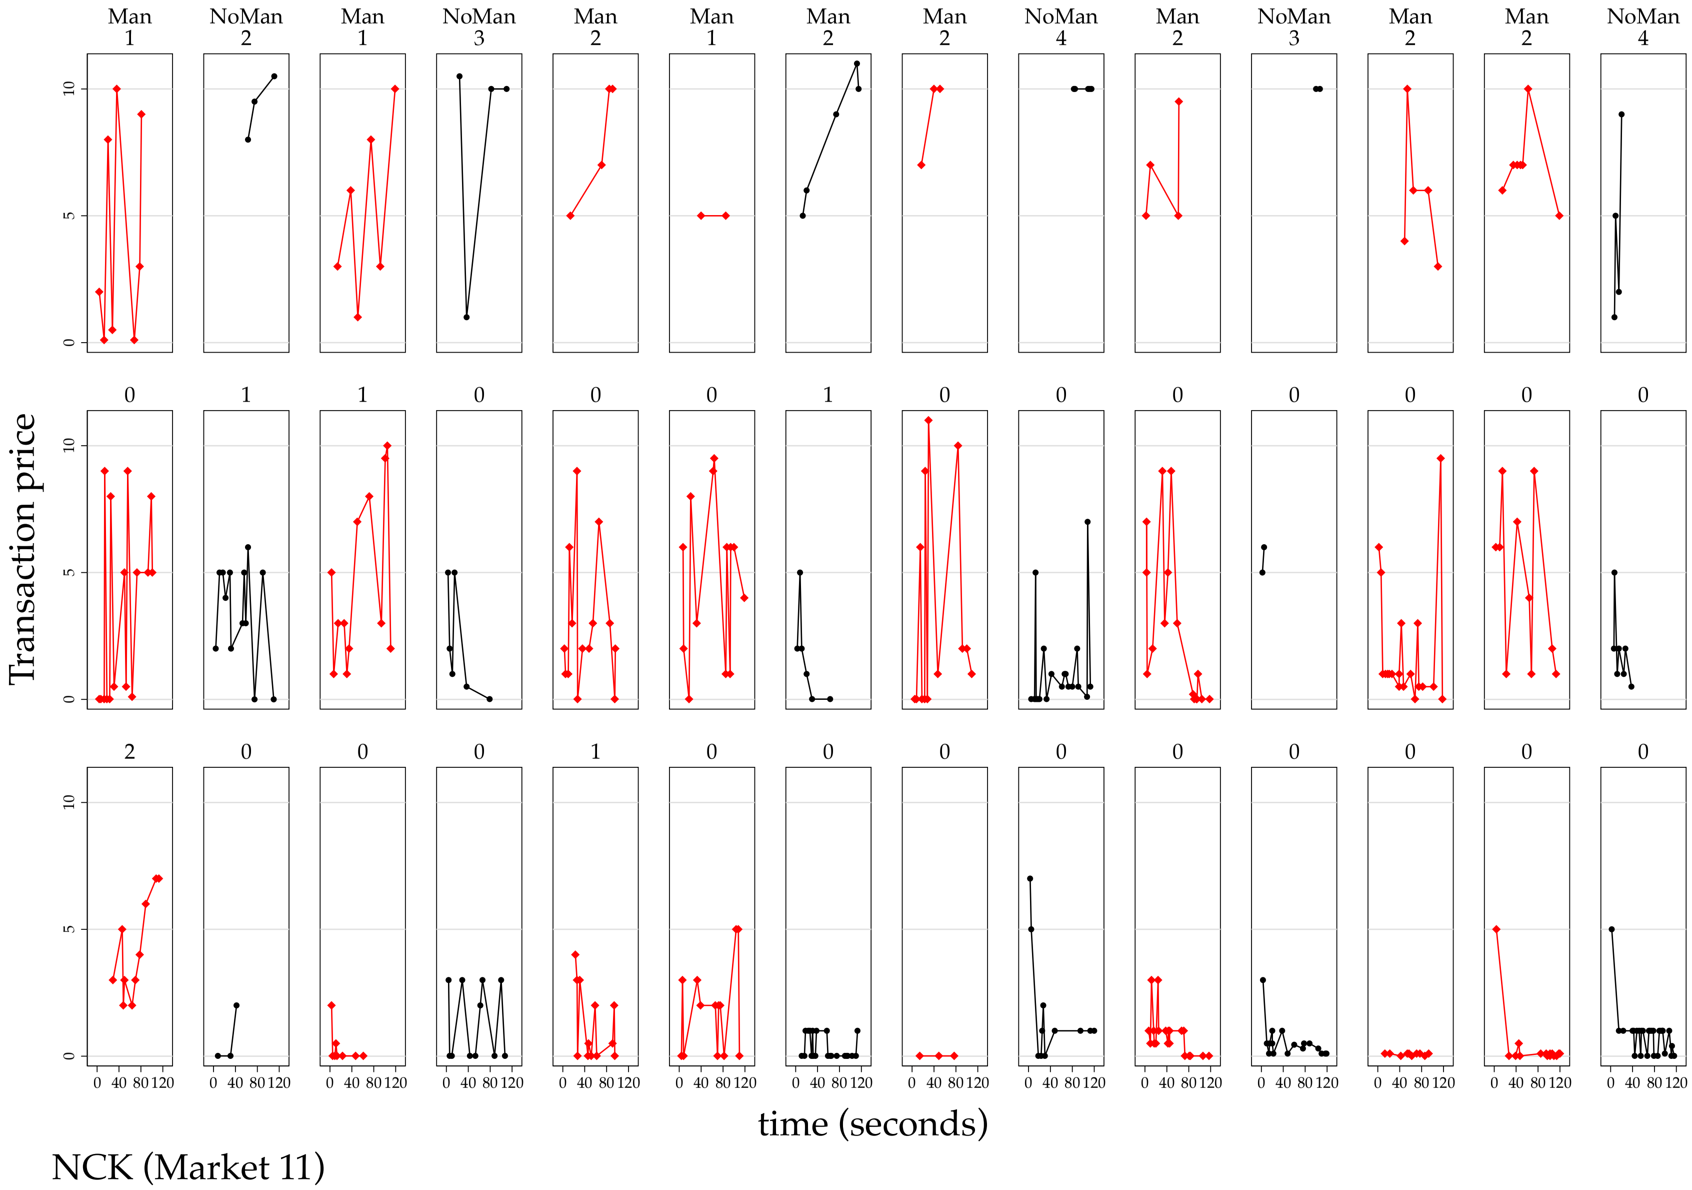

Figure A11: Session 11.

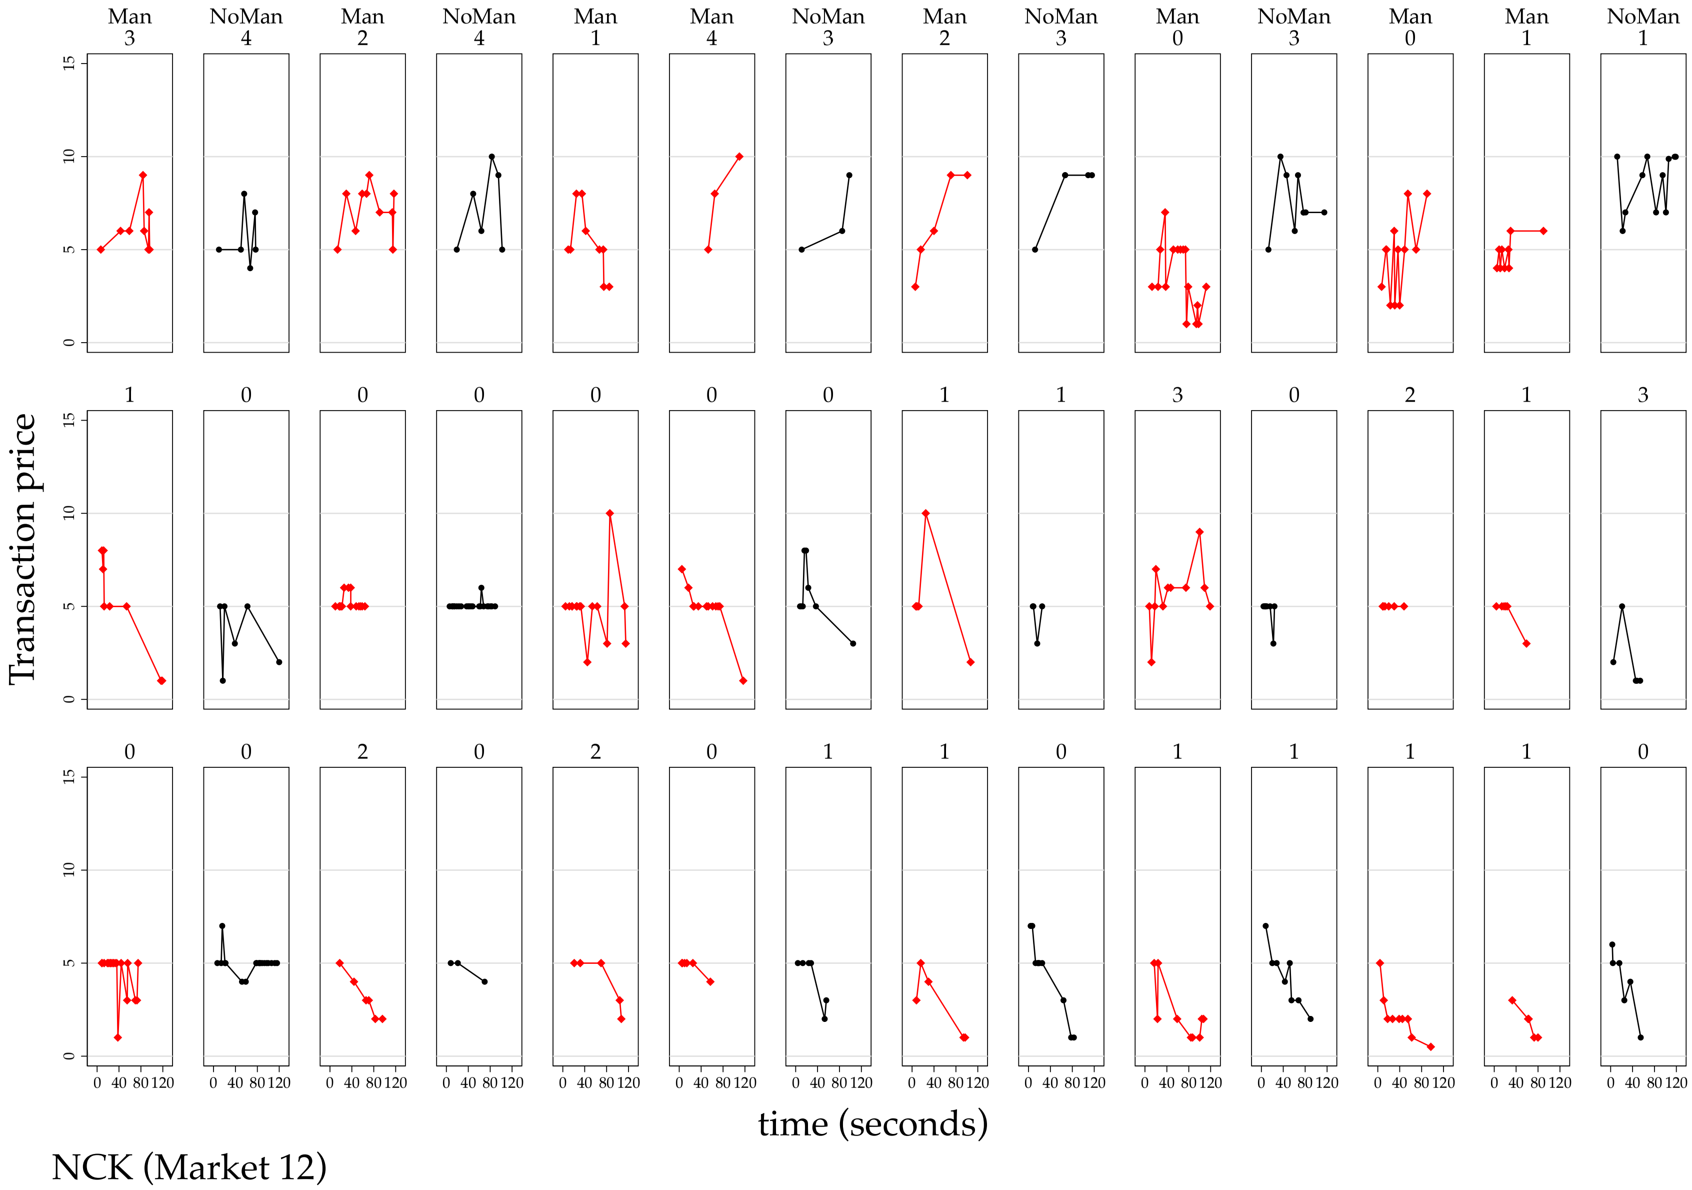

Figure A12: Session 12.

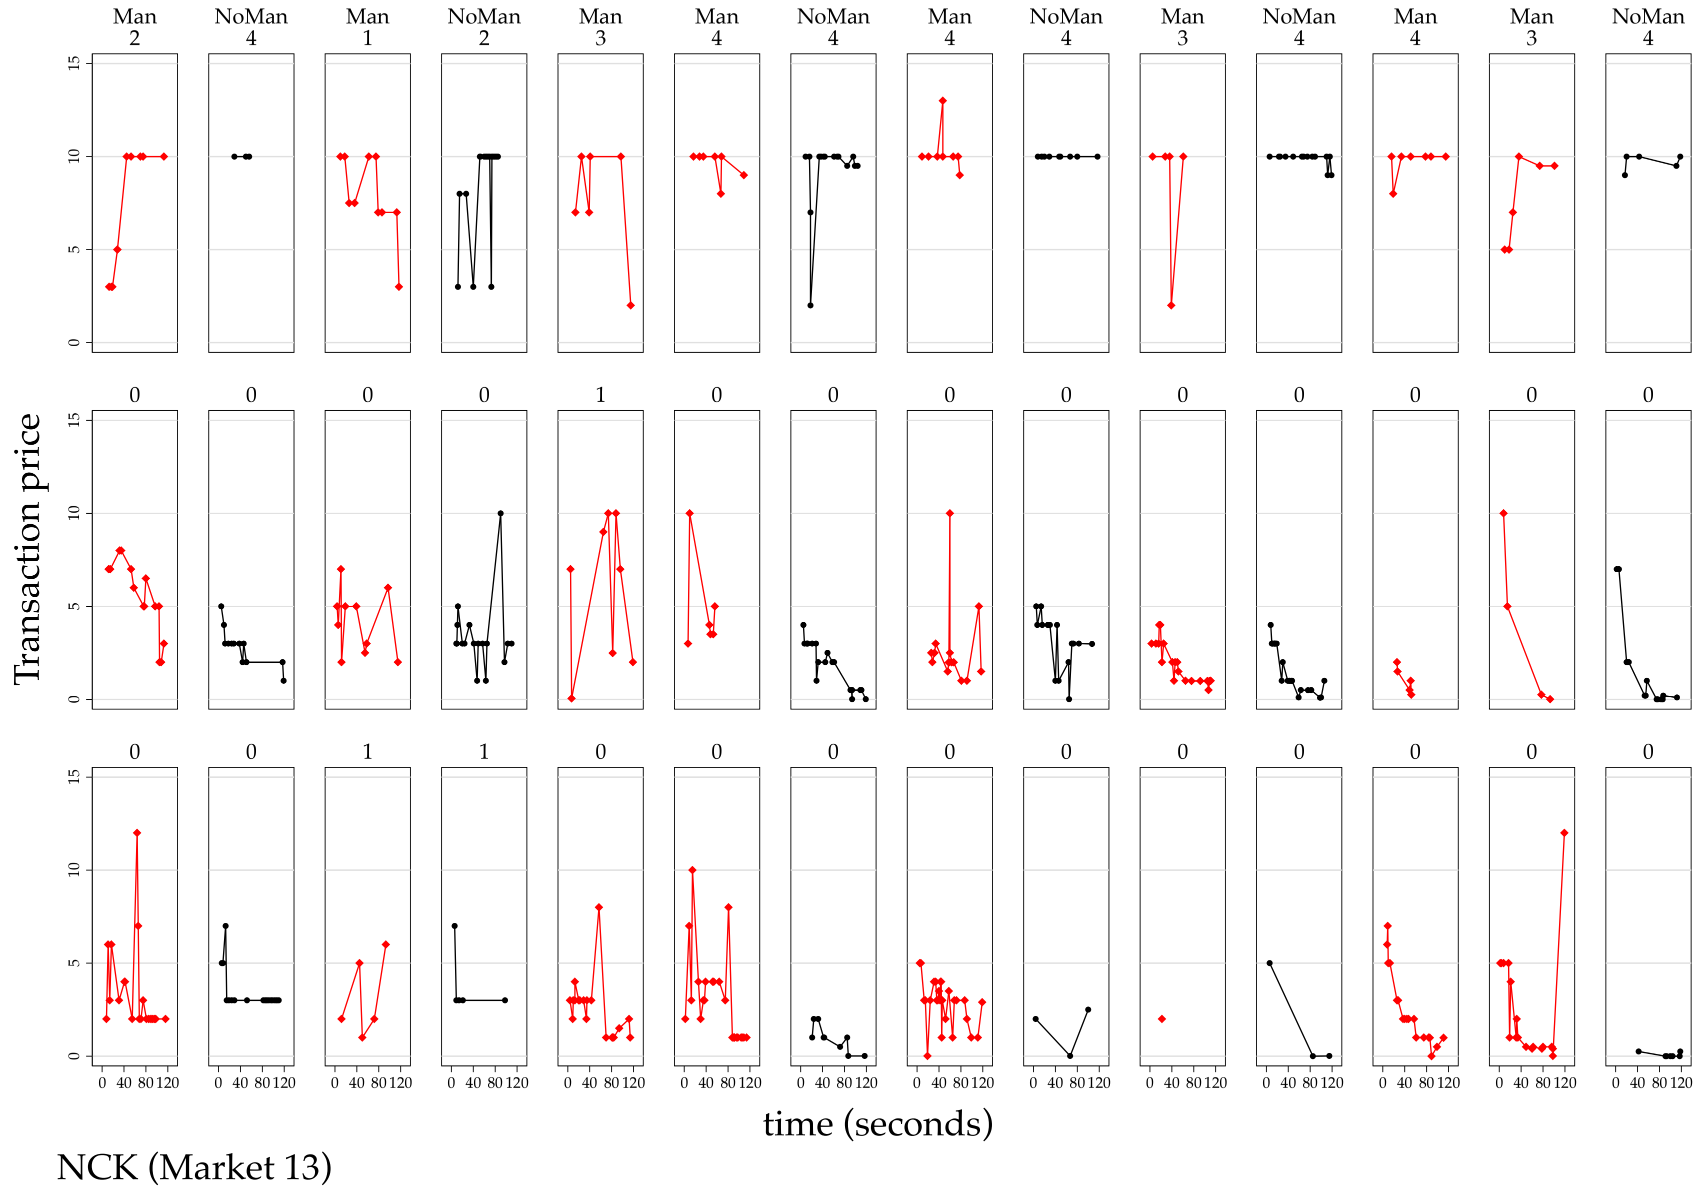

Figure A13: Session 13.

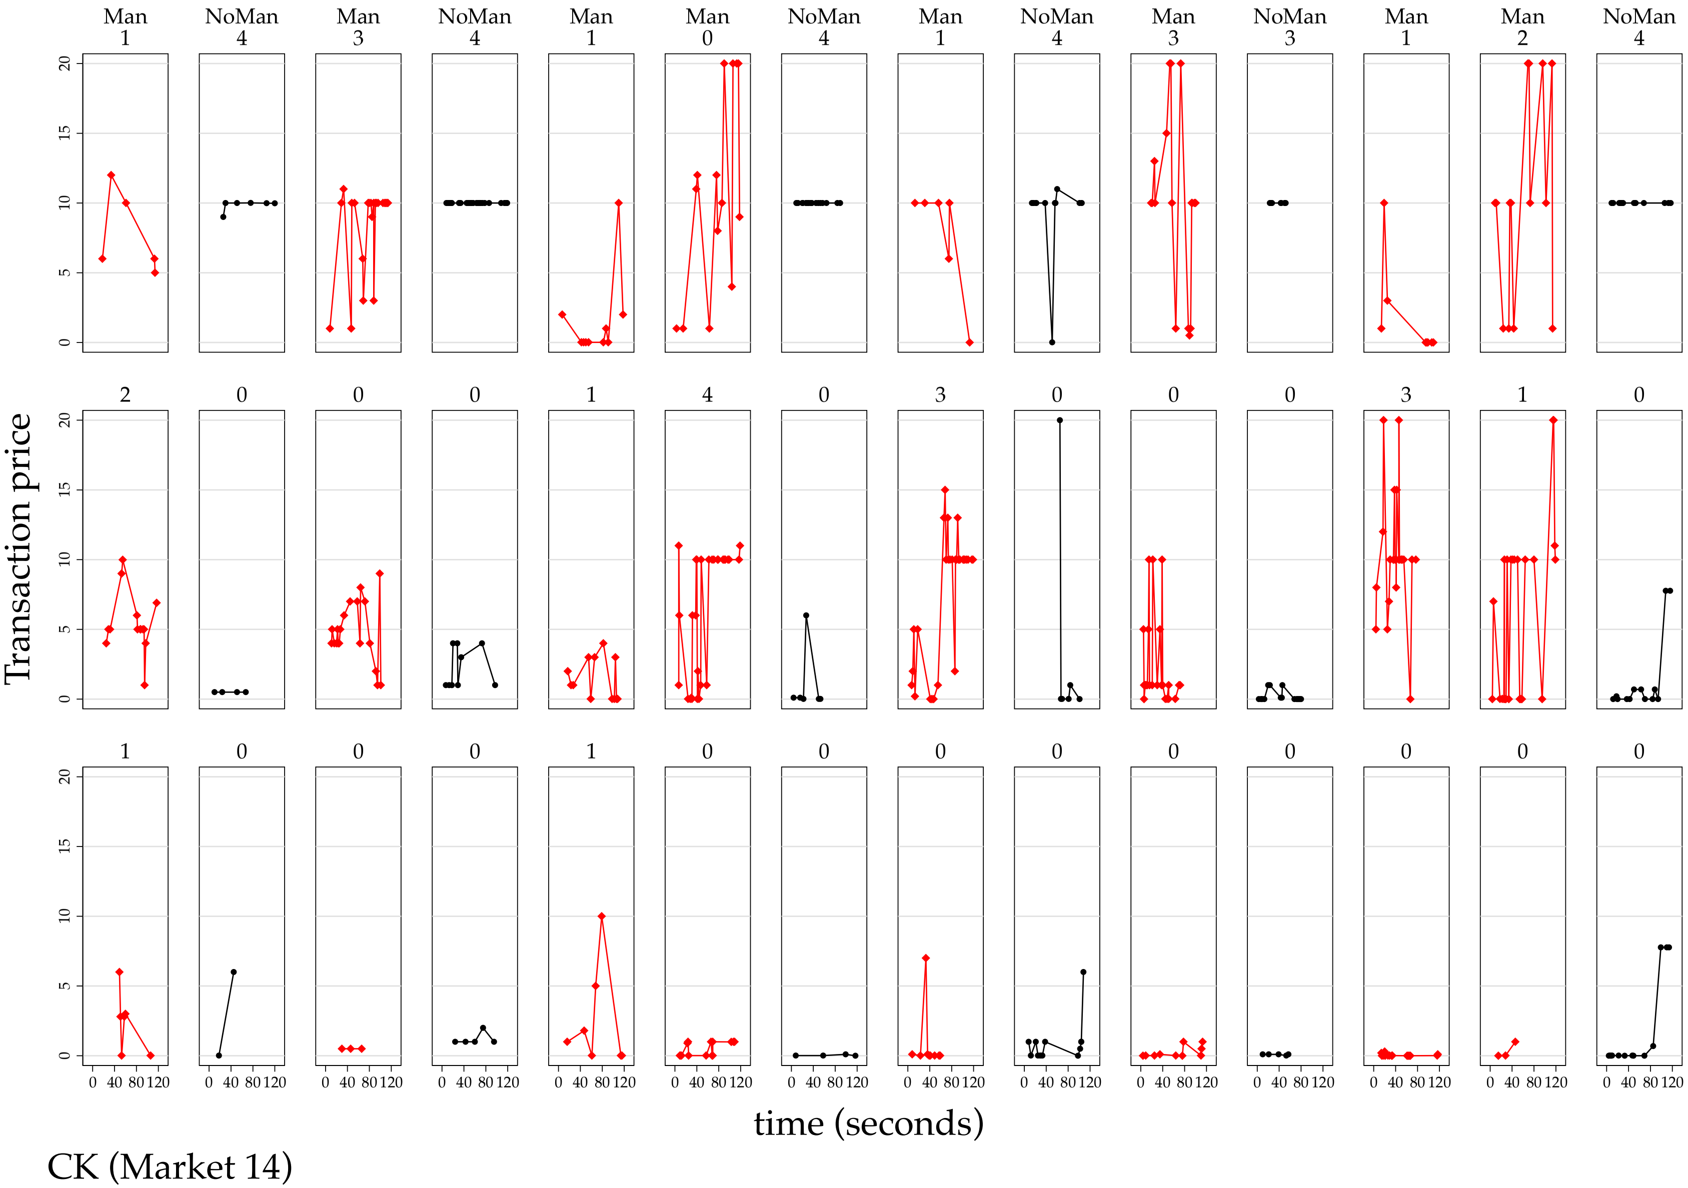

Figure A14: Session 14.

A.2. Trading patterns.

The solid black line of Figure A15 details the mean transaction volume of each security type over 30 second time blocks. The transaction volumes for the Fake and Neutral securities appear to steadily decline over the trading period. In contrast, no such decline is apparent for the True security.

The stacked bars on Figure A15 represent the proportion of transactions in each 30 seconds time block by trader types. For example, the RB^1 area details the the proportion of transactions where a $\mathscr{R}$ trader is the buyer and the $\mathscr{B}^1$ trader is the seller.

The *CK-NoMan* treatment (first row) provides a benchmark description of trading behaviour. In the first 60 seconds of the market duration, we clearly see that traders often condition their demand for securities on their private information about the True state. For example, the Fake security is almost always purchased by the $\mathscr{B}^2$ traders — who assign a positive posterior to the Fake state given their private information — and is often sold by the $\mathscr{R}$ and $\mathscr{B}^1$ traders — who can rule out the Fake state. Likewise, the Neutral security is almost always purchased by the $\mathscr{B}^1$ and R traders and is often sold by the $\mathscr{B}^2$ traders.

In comparison, purchases of the True security — to which all traders assign a positive probability — appear to be evenly distributed amongst all trader types. Interestingly, traders appear to increasingly buy the securities which they know to be worthless in the last 60 seconds of trade. For example in the *NoMan* markets, while the $\mathscr{R}$ and $\mathscr{B}^1$ traders account for only 2% of Fake security purchases in the first 30 seconds, their share of such purchases increases to 21% in the last 30 seconds of trade. This perhaps suggests that some traders may have been trying to engage in speculative trade, creating demand for securities that they know to be worthless.

Turning our attention to the *CK-Man* treatment (second row), we observe that the $\mathscr{B}^1$ and $\mathscr{B}^2$ traders behave similarly to the *CK-NoMan* treatment. In contrast, we now observe that the $\mathscr{R}$ traders account for a substantial proportion of the Fake security purchases. The $\mathscr{R}$ traders are also fairly active in purchasing the Neutral security, possibly in an attempt to obscure the True state from the policy makers by increasing noise trading.

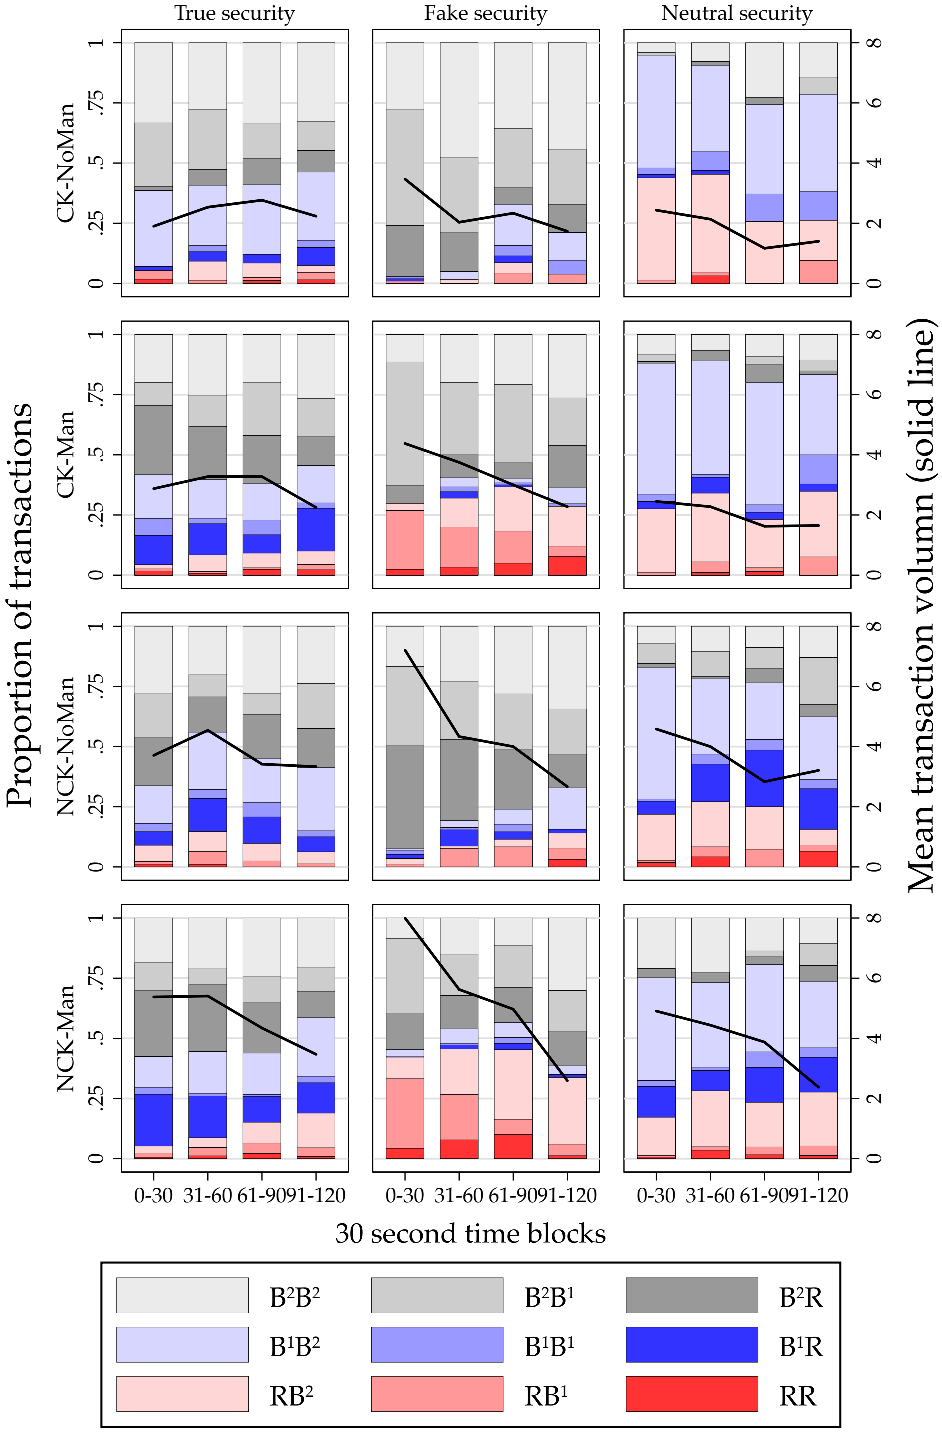

Notes. The panels are organised by treatments (rows) and securities (columns). The solid line (right axis) details the mean transaction volume of each security type for each 30 seconds time block. The stacked bars detail the proportion of transactions by the trader types. For example, the area RB^1 details the proportion of transactions where the buyer is a R trader and the seller is a $\mathscr{B}^1$ trader.

Figure A15: Mean transaction volumes and proportion of transaction by trader types.

Appendix B. Instructions

The experiments were conducted in English at the University of Exeter. Each session consisted of two parts, I and II. Part I is the training phase and Part II is the experimental proper. The *Common knowledge* (*CK*) and *No common knowledge* (*NCK*) treatments only differ in part II. The following details the instructions for both parts. To experimenter also read the instructions to the participants. Proportions of the instructions that are relevant to the *CK* and *NCK* treatments will be labeled "text" and "text", respectively.

B.1. Instructions: Part I

You are participating in an experiment on decision-making. If you follow the instructions and apply them carefully, you can earn some money in additional to the 5 GBP show-up fee which we will give you in any case. From now on you are not allowed to talk to any other participants in the experiment. If you have a question, please raise your hand and one of the instructors will attend to you. The experiment will consist of two parts (Part I and Part II). Your earnings in this experiment will depend on your decisions in Parts I and II.

- In the following, we present the instructions for Part I of the experiment.
- The Part II instructions will be available at the end of Part I.

All payoffs in this experiment will be denoted in Experimental Currency (ECU). Part I of the experiment will consist of one (non-paying) practice round followed by five experimental rounds. At the end of part I, the computer will randomly select one of the five experimental rounds, and your earnings in that round will be paid to you in cash at the exchange rate of 100 ECU to 1 GPB. In this experiment, you will be randomly matched with 11 other participants. You will interact with the same participants for all rounds in part I. We will now describe each experimental round

B.1.1. Design of a round.

There is a box that contains three balls which are labeled X, Y and Z.

- The computer will randomly choose one ball from the box. This means that there is an equal chance for the X, Y and Z balls to be chosen.
- After a ball is chosen, the 12 participants are randomly separated into two groups (6 participants in each group).

Table B1: Chosen ball and information received.

Chosen ball	Information for the first group	Information for the second group
X	The ball is not Y	The ball is not Z
Y	The ball is not Z	The ball is not X
Z	The ball is not Y	The ball is not X

Note. Participants will not know if they belong to the first or second group.

- The computer gives each group some information about the chosen ball. However, the two groups will always receive different information.
- Table B1 details the information that each group will get depending on the chosen ball.

For example, suppose that the X ball is chosen.

- The computer informs one group that the chosen ball is not Y (it can be X or Z).
- The computer informs the other group that the chosen ball is not Z (it can be X or Y).

Notice from the above that the two groups will always receive different information. Also notice that if you receive the message "The ball is not Z", there is a 50% chance the ball is X and a 50% chance the ball is Y. Likewise, for the other two messages.

B.1.2. Market Stage

After all participants have received some information about the chosen ball, the market stage starts. Here, you will have the opportunity to buy and sell 3 classes of certificates labeled cert-x, cert-y and cert-z.

- Each participant begins the market stage with 200 ECU, 5 units of cert-x, 5 units of cert-y and 5 units of cert-z.
- You buy and sell the certificates through a market system that will last for exactly 120 seconds (2 mins).
- After 120 second, the market stage ends and your units of cert-x, cert-y and cert-z certificates will be redeemed at a value that depends on the chosen ball.

Table B2 details how the value of each certificate class depends on the chosen ball — you will only know the value of each certificate at the end of the round.

For example, if the chosen ball is X then only cert-x will be valued at 10. In contrast, cert-y and cert-z will be valued at 0. Similarly, if the chosen ball was Y, then only cert-y

Table B2: Certificate value and chosen ball.

	Chosen ball X	Chosen ball Y	Chosen ball Z
cert-x value	10 ECU	0 ECU	0 ECU
cert-y value	0 ECU	10 ECU	0 ECU
cert-z value	0 ECU	0 ECU	10 ECU

will be valued at 10 ECU and the other certificates at 0 ECU. Your total payoffs will be:

$$\text{Total payoff} = \text{(Money you own)} + \text{(money from certificate value)}$$

For example, suppose that the chosen ball is Y and you ended the market stage with: Money 150 ECU; 2 x cert-x; 10 x cert-y and 3 x cert-z. Your total payoff will be 150 + 2(0) + 10(10) + 3(0) = 250 ECU.

For example, suppose that the chosen ball is X and you ended the market stage with: Money 150 ECU; 2 x cert-x; 10 x cert-y and 3 x cert-z. Your total payoff will be 150 + 2(10) + 10(0) + 3(0) = 170 ECU.

B.1.3. How the market system works

Figure B1 presents an overview of the trading platform. To buy or sell a certificate, first enter the "BID" and "ASK" prices for the desired certificate class.

- ASK Price (between 0 and 20; up to two decimal places) tells the other participants how much you are willing to sell a unit of the certificate for. The software will always show your most recent ASK price.
- Bid Price (between 0 and 20; up to two decimal places) tells the other participants how much you are willing to buy a unit of the certificate for. The software will always show your most recent BID price.

The BID and ASK prices of all participants are listed on the column "Market BID Prices" and "Market ASK Prices", respectively.

- To buy a certificate. Select the price on the "Market ASK Prices" that you wish to buy at and click the "buy button".
- To sell a certificate. Select the price on the "Market BID Prices" that you wish to sell at and click the "sell button".

Please note the following rules:

- You can only buy a certificate if you have sufficient money – you cannot borrow money.

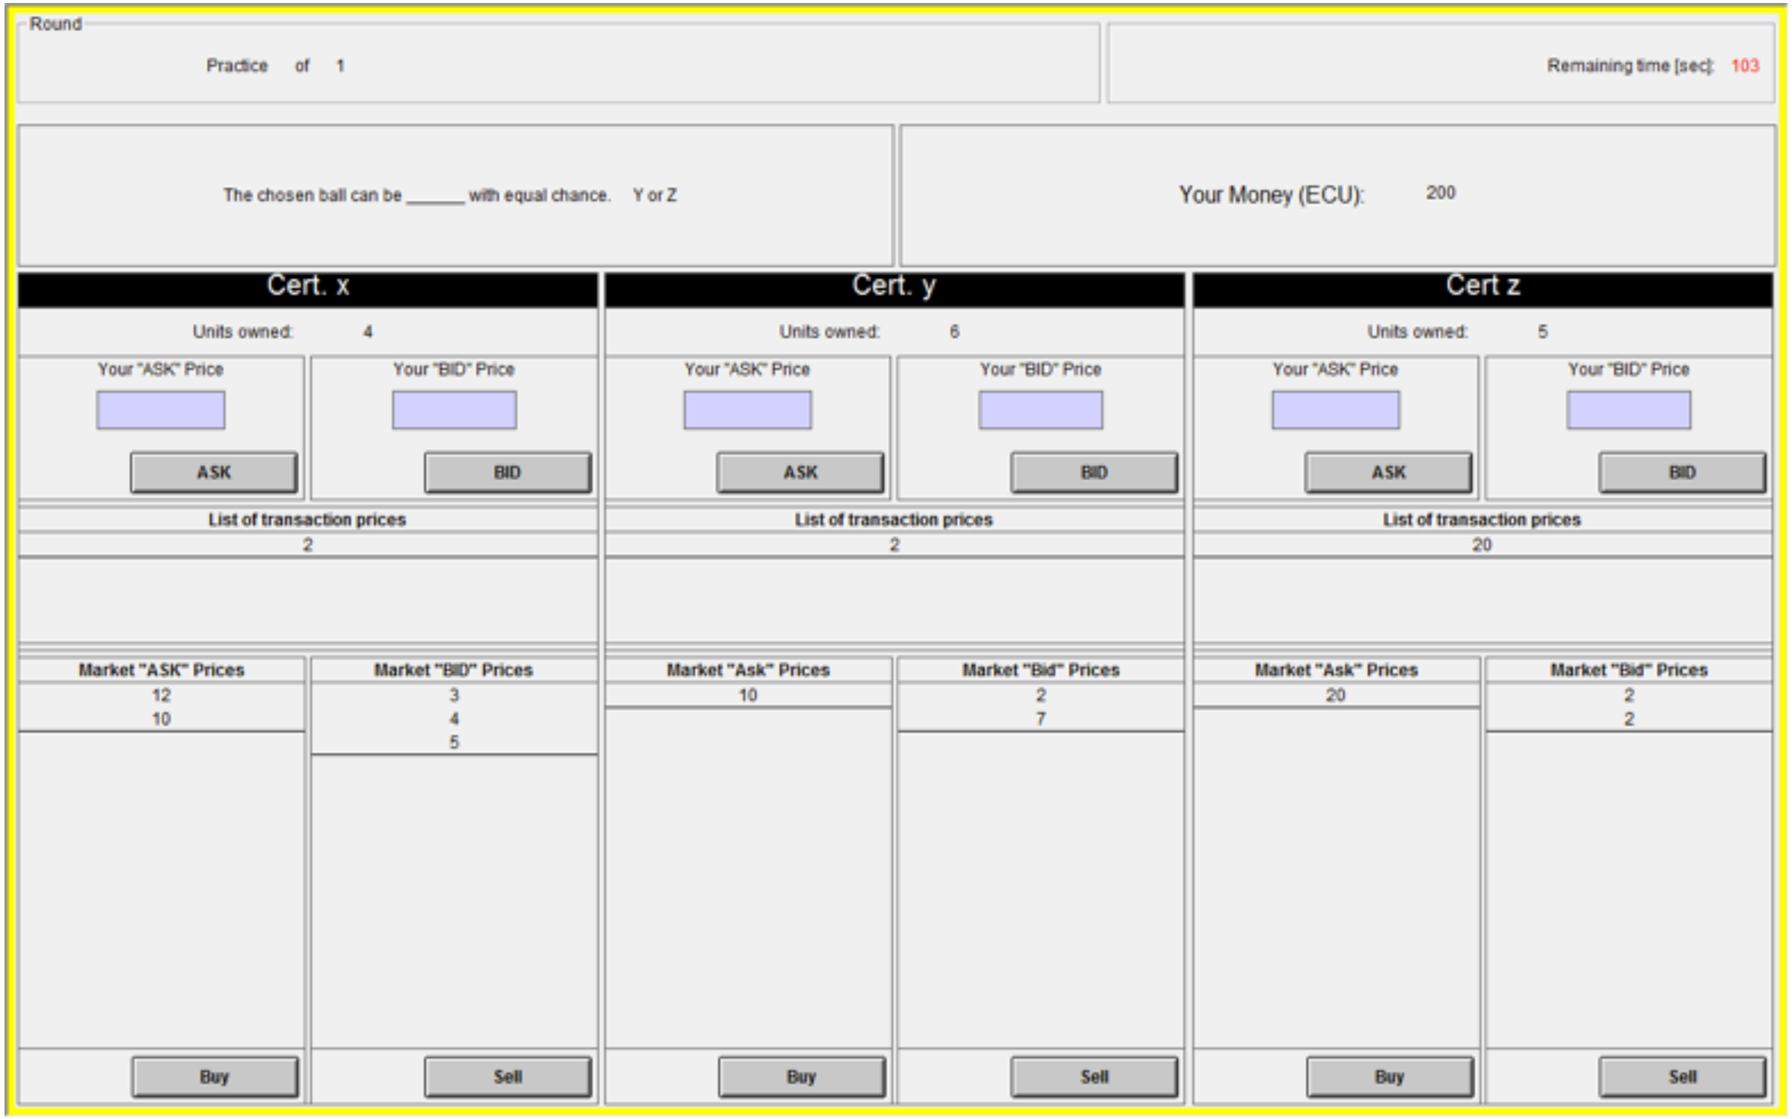

Figure B1: Screenshot of the trading platform.

- You can only sell a certificate if you have at least one unit of that certificate — you cannot borrow certificates.

- You cannot buy and sell with yourself — your BID and ASK prices will be reflected in Blue.

B.1.4. Other information

When the experiment starts, we ask that you manually record the information received and chosen ball at each round. To ensure that everyone understands the experiment design, we included a simple set of control questions. The experiment will only start when everyone completes the control questions. Remember that part I will consist of 5 experimental rounds and you will interact with the same other participants for all rounds. We will distribute the instructions for part II once everyone has completed part I. In the meantime, please feel free to raise your hands if you have any questions and the experimenter will answer you privately.

B.1.5. Control questions.

Please submit your answers to the following questions on the computer.

(a). If you are informed that the chosen ball is not Z, the probability that the chosen ball is Y must be ___ %. (0%; 33.33%; 50%; 67.67%, 100%)

(b). If the chosen ball was Y, then both groups will receive the same information. (TRUE; FALSE)

(c). If the chosen ball is X, cert-x will be valued at ___ ECU.

(d). If the chosen ball is Y, cert-z will be valued at ___ ECU.

(e). Suppose that the chosen ball is X and you ended the market system with: Money 200 ECU; 2 units of cert-x; 3 units of cert-y and 4 units of cert-z. Your total payoff will be ___ ECU.

(f). Suppose that the chosen ball is Z and you ended the market system with: Money 200 ECU; 2 units of cert-x; 3 units of cert-y and 4 units of cert-z. Your total payoff will be ___ ECU.

B.2. Instructions: Part II

Part II of the experiment will consist of one (non-paying) practice round followed by fourteen experimental rounds. At the end of part II, the computer will randomly select two of the fourteen experimental rounds, and your earnings in these two rounds will be paid to you in cash at the exchange rate of 100 ECU to 1 GBP. At each round, you will be matched with the same 11 other participants that you interacted with in part I. Some features of experiment in Part II will be similar to those in Part I. However, to avoid confusion, we will repeat these features in the instructions. We will now describe each experimental round.

B.2.1. Design of a round.

Each of the 12 participants will be randomly allocated to a role, which can either be (a) Voter, (b) Red-Trader or (c) Blue-Trader. In each round, there are exactly,

- 4 x Voters
- 2 x Red-Traders
- 6 x Blue-Traders

All participants will be informed of their role at the start of the round. All participants will assume the same role for the entire experiment. After all participants have been assigned a role, the computer will assign a "type" for the Red-Traders.

- The Red-Traders can be type-A or type-B with equal chance.
- At the start of the round, all participants will know whether the Red-Traders are type-A or type-B.
- At the start of the round, ONLY the Red-Traders will know whether they are type-A or type-B. Voters and Blue-Traders WILL NOT know whether the Red-Traders are type-A or type-B.

B.2.2. Choosing a ball

There is a box that contains three balls which are labeled X, Y and Z.

- The computer will randomly choose one ball from the box. This means that there is an equal chance for the X, Y and Z balls to be chosen.
- Voters receive NO information about the chosen ball.
- Red-Traders and Blue-Traders receive some information about the chosen ball. To do this, the computer randomly separates the Red-Traders and Blue- Traders into two groups (4 participants in each group), with the rule that the Red-Traders will always belong to the same group.
- Thereafter, the computer gives each group some information about the chosen ball. However, the two groups will always receive different information — see Table B3.

Table B3: Chosen ball and information received.

Chosen ball	Information for the first group	Information for the second group
X	The ball is not Y	The ball is not Z
Y	The ball is not Z	The ball is not X
Z	The ball is not Y	The ball is not X

Note. Participants will not know if they belong to the first or second group.

Table B4: Certificate value and chosen ball.

	Chosen ball X	Chosen ball Y	Chosen ball Z
cert-x value	10 ECU	0 ECU	0 ECU
cert-y value	0 ECU	10 ECU	0 ECU
cert-z value	0 ECU	0 ECU	10 ECU

For example, suppose that the X ball is chosen.

- Voters receive NO information.
- One group of Traders are informed that the chosen ball is not-Y (it can be X or Z).
- Another group of Traders are informed that the chosen ball is not-Z (it can be X or Y).

Notice again that the two groups will always receive different information. Also, the Blue-Traders will not know if the Red-Traders are in their group or in the other group.

B.2.3. Market stage

Only Red-Traders and Blue-Traders get to participate in the Market stage. At this stage, there is no difference between the Red and Blue traders — the Voters do not participate in the trading, but can observe the actions of the Traders. Here, Red-Traders and Blue-Traders will have the opportunity to buy and sell 3 classes of certificates labeled cert-x, cert-y and cert-z. Each trader begins the market stage with 200 ECU, 5 units of cert-x, 5 units of cert-y and 5 units of cert-z. The traders buy and sell the certificates through a market system that will last for exactly 120 seconds (2 mins). The market system is identical to that of Part I. After 120 second, the market stage ends and the units of cert-x, cert-y and cert-z certificates will be redeemed at a value that depends on the chosen ball. Table B4 details how the value of each certificate class depends on the chosen ball — traders will only know the value of each certificate at the end of the round.

For example, if the chosen ball is X then only cert-x will be valued at 10 ECU. In contrast, cert-y and cert-z will be valued at 0 ECU. Similarly, if the chosen ball was Y, then only cert-y will be valued at 10 ECU and the other certificates at 0 ECU.

B.2.4. Voting stage

Only Voters get to participate in the voting stage. The Traders do not participate in this stage, but their payoffs are determined by the choices of the Voters. Voters will see the transaction prices of all certificates. Voters will choose between four possible projects:

- Project-X
- Project-Y
- Project-Z
- Project-Q

Each voter can only vote for one project. After all voters have voted, the project with the most votes will be implemented. If there is a tie (i.e., two or more projects have the same number of votes), then Project-Q will be implemented by default.

For example, suppose that: Project-X (2 votes); Project-Y (1 vote); Project-Z (0 vote); and Project-Q (1 vote). Project-X will be implemented as it has the most votes

Another example. suppose that: Project-X (0 vote); Project-Y (0 vote); Project-Z (1 vote); and Project-Q (3 votes). Project-Q will be implemented as it has the most votes

Final Example, suppose that: Project-X (2 votes); Project-Y (2 votes); Project-Z (0 vote); and Project-Q (0 vote). There is a tie between Project-X and Project-Y. In this case, Project-Q will be implemented.

B.2.5. Payoffs

Traders get a sure payoff of 400 ECU. Voters get a sure payoff of 650 ECU. The additional payoffs they earn in the round will depend on their roles, type, market decisions and implemented project.

- The Total payoffs for Voters are:

  $$\text{Total Payoff} = 650 + \text{Project-Earnings}$$

  Here, the Project-Earnings will depend on the chosen ball and voting stage implemented project — we will elaborate on this later.

- The Total payoffs for Red-Traders are:

  $$\text{Total Payoff} = 400 + [\text{Money} + \text{Certificate-values}] + \text{Project-Earnings}$$

  Here, Money refers to the money that Red-Traders own at the end of the market stage. Certificates-Values refer to the values of the certificate owned by Red-Traders at the end of the market stage. The Project-Earnings will depend on the

Table B5: Project-Earnings for Voters, Blue-Traders and Red-Trader (type-A).

	Chosen ball X	Chosen ball Y	Chosen ball Z
Project-X (Implemented)	400 ECU	-400 ECU	-400 ECU
Project-Y (Implemented)	-400 ECU	400 ECU	-400 ECU
Project-Z (Implemented)	-400 ECU	-400 ECU	400 ECU
Project-Q (Implemented)	100 ECU	100 ECU	100 ECU

chosen ball, Red-Traders' type and the voting stage implemented project — we will elaborate on this later.

- The Total payoffs for Blue-Traders are:

$$\text{Total Payoff} = 400 + [\text{Money} + \text{Certificate-values}] + \text{Project-Earnings}$$

Here, Money refers to the money that Blue-Traders own at the end of the market stage. Certificates-Values refer to the values of the certificate owned by Blue-Traders at the end of the market stage. The Project-Earnings will depend on the chosen ball and the voting stage implemented project — we will elaborate on this later.

Notice from above that the Project-Earning for Voters and Blue-Traders depend on the chosen ball and implement project. Project-Earning for Red-Traders depend on the chosen ball, type and implement project. The difference is because, Red-Traders can be type-A or type-B with equal chance. This is explained in the next section.

B.2.6. Project-Earnings

The Project-Earning for the voters and the Blue-Traders are detailed on Table B5.
For example, if the chosen ball is X and project-X is implemented

- Voter: Project-Earnings = 400 ECU
- Blue-Trader: Project-Earnings = 400 ECU

If the chosen ball is X and project-Q is implemented

- Voter: Project-Earnings = 100 ECU
- Blue-Trader: Project-Earnings = 100 ECU

Finally, if the chosen ball is X and project-Y is implemented

- Voter: Project-Earnings = -400 ECU
- Blue-Trader: Project-Earnings = -400 ECU

Table B6: Project-Earnings for Red-Trader (type-B).

	Information not-X	Information not-Y	Information not-Z
Project-X (Implemented)	1000 ECU	-400 ECU	-400 ECU
Project-Y (Implemented)	-400 ECU	1000 ECU	-400 ECU
Project-Z (Implemented)	-400 ECU	-400 ECU	1000 ECU
Project-Q (Implemented)	100 ECU	100 ECU	100 ECU

Notice from Table B5 that voters and Blue-Traders obtain the highest payoff if the implemented project corresponds to the chosen ball. The Project-Earning for the Red-Traders will depend on whether they are type-A or type-B.

- If the Red-Traders are type-A, their Project-Earnings will be identical to those of voters and Blue-Traders — See Table B5. Here, the Red-Traders' Project-Earnings will depend on the implemented project and chosen ball.
- If the Red-Traders are type-B, their Project-Earnings will be different to those of voters and Blue-Traders — See Table B6. Here, the Red-Traders' Project-Earning will depend on the information received (i.e., Not-X, Not-Y, Not-Z) and the implemented project.

For example, if the chosen ball is X, the Red-Traders receive information not-Y and project-X is implemented,

- Type-A Red-Trader: Project-Earnings = 400 ECU
- Type-B Red-Trader: Project-Earnings = -400 ECU

If the chosen ball is X, the Red-Traders receive information not-Y and project-Y is implemented,

- Type-A Red-Trader: Project-Earnings = -400 ECU
- Type-B Red-Trader: Project-Earnings = 1000 ECU

Notice from the Tables B5 and B6 that a type-A Red-Trader obtains the highest payoff if the implemented project corresponds to the chosen ball. In contrast, a type-B Red-Trader obtains the highest payoff if the implemented project is guaranteed to be different from the chosen ball.

B.2.7. Some worked examples.

To help you better understand the payoff design, we included some worked examples. In the following examples, assume that the

- Red-Trader: Has 200 ECU money, 3 units cert-x, 4 units cert-y and 5 units cert-z.
- Blue-Trader: Has 200 ECU money, 3 units cert-x, 4 units cert-y and 5 units cert-z.

Also recall that traders get a sure payoff of 400 ECU. Voters get a sure payoff of 650 ECU.

Example 1. Chosen ball is X, Red-Traders observe not-Y and project-X is implemented.

- Voter: Project-Earning = 400 ECU; Total Payoff = 650 + 400 = 1050 ECU.
- Blue-Trader. Project-Earning = 400 ECU; Market stage earning=200+3(10) + 4(0) + 5(0) = 230 ECU; Total Payoff = 400 + 230 + 400 = 1030 ECU.
- Red-Trader (Type-A). Project-Earning = 400 ECU; Market stage earning = 200 + 3(10) + 4(0) + 5(0) = 230 ECU; Total Payoff = 400 + 230 + 400 = 1030 ECU.
- Red-Trader (Type-B). Project-Earning = -400 ECU; Market stage earning = 200 + 3(10) + 4(0) + 5(0) = 230 ECU; Total Payoff = 400 + 230 - 400 = 230 ECU.

Example 2. Chosen ball is X, Red-Traders observe not-Y and project-Y is implemented.

- Voter: Project-Earning = -400 ECU; Total Payoff = 650 - 400 = 250 ECU.
- Blue-Trader. Project-Earning = -400 ECU; Market stage earning=200+3(10) + 4(0) + 5(0) = 230 ECU; Total Payoff = 400 + 230 - 400 = 230 ECU.
- Red-Trader (Type-A). Project-Earning = -400 ECU; Market stage earning = 200 + 3(10) + 4(0) + 5(0) = 230 ECU; Total Payoff = 400 + 230 - 400 = 230 ECU.
- Red-Trader (Type-B). Project-Earning = 1000 ECU; Market stage earning = 200 + 3(10) + 4(0) + 5(0) = 230 ECU; Total Payoff = 400 + 230 + 1000 = 1630 ECU.

Notice that voters may be able to learn about the chosen ball from the observe prices in the market stage. As such, Voters' choices in the voting stage may depend on the market stage prices of cert-x, cert-y and cert-z. Traders could take this relationship into consideration when deciding upon their market behaviors.

B.2.8. Other Information.

To ensure that everyone understands the experiment design, we included a simple set of control questions. The experiment will only start when everyone completes the control questions. When the experiment starts, we ask that you manually record the information received and chosen ball at each round. Remember that part II will consist of

14 experimental rounds and you will interact with the same other participants for all rounds. In the meantime, please feel free to raise your hands if you have any questions and the experimenter will answer you privately

B.2.9. Control questions.

(A). You will keep the same role for all rounds in Part II. (TRUE; FALSE)

(B). Voters and Blue-Traders will always be informed about the Red-Traders' type. (TRUE; FALSE)

(C). The Red-Traders will always be in the same group. (TRUE; FALSE)

(D). The implemented project for all participants will be determined by the decisions of the voters. (TRUE; FALSE)

(E). If Project-X (2 votes), Project-Y (1 vote), Project-Z (1 vote) and Project-Q (0 vote), then project ___ will be implemented for all participants.

(F). If Project-X (2 votes), Project-Y (2 vote), Project-Z (0 vote) and Project-Q (0 vote), then project ___ will be implemented for all participants.

(G). Voters' and Blue-Traders' Project-Earnings are highest when the implemented project corresponds to the chosen ball. (TRUE; FALSE)

(H). Red-Traders' (type-A) Project-Earnings are highest when the implemented project corresponds to the chosen ball. (TRUE; FALSE)

(I). Red-Traders' (type-B) Project-Earnings are highest when the implemented project corresponds to the chosen ball. (TRUE; FALSE)

(J). The chosen ball is Y and Project-Y is implemented. Red-Traders receive information that the ball is not-X.

1. Project-Earning for voters, Blue-Traders & Red-Traders(type-A): ___ (ECU).
2. Project-Earning for Type-B Red-Traders: ___ (ECU)

(K). The chosen ball is Y and Project-X is implemented. Red-Traders receive information that the ball is not-X.

1. Project-Earning for voters, Blue-Traders & Red-Traders(type-A) : ___ (ECU).
2. Project-Earning for Type-B Red-Traders: ___ (ECU)

(L). The chosen ball is Y and project-X is implemented. They total payoffs for a Type-B Red-Trader who observes Not-X and has (Money: 200 ECU), (2xCert-x), (3xCert-y) and (4xCert-z) will be ___ ECU. Recall that traders also receive a sure payment of 400 ECU.