Endogenous Preferences and Conformity: Evidence from a Laboratory Experiment

Sergio Beraldo*
Department of Economics and Statistics
University of Naples and CSEF

Valerio Filoso
Department of Law
University of Naples

Marco Stimolo
Department of Economics and Statistics
University of Naples and University of Salerno

March 2, 2015

Abstract

We elicited individuals' Willingness to Accept evaluations through an experimental Vickrey auction to investigate whether information on irrelevant signals (the market price, extreme asks) affect the elicited values. We find that individuals' evaluations strongly conform to the asks either at the bottom or the top of the distribution by a factor of 44—66%; this tendency to conform remarkably reduces the shaping effect of the market price. Our results are identified by means of a novel methodology distinguishing between the dynamic and the asymptotic features of preference formation. The overall evidence suggests a decisive role for market interaction in preference formation.

Keywords: Endogenous preferences
Shaping effect
Conformity
Experimental Vickrey auction

JEL codes: C91, C92, D44.

*Corresponding author. E-mail: `s.beraldo@unina.it`

1 Introduction

The existence of anomalies that are at odds with standard theories of preferences is well known (Allais, 1953; Knetsch and Sinden, 1984; Kahneman and Tversky, 1979), as it is well known that such anomalies tend to decay in experimental repeated markets where subjects have the chance of acquiring experience through repetition (Binmore, 1999; Cox and Grether, 1996; Shogren et al., 1994). According to Loomes et al. (2003), this tendency of the anomalies to decay may be explained by grounding either in the hypothesis of exogenous, context independent preferences (Plott, 1996) or in the hypothesis that preferences are endogenous to the institutions through which they are elicited. Whereas the former hypothesis would imply that preferences are revealed in markets, but markets do not have the power to change preferences, according to the latter hypothesis preferences are significantly conditioned by market institutions. On this account, a substantial body of evidence indeed shows that incentive-compatible elicitation mechanisms produce responses in which market signals act as cues affecting the elicited values. In particular, much evidence has been collected about the so-called *shaping effect*, that is, the tendency of both willingness to pay (WTP) and willingness to accept (WTA) valuations to move towards the market price (Ariely et al., 2006; Isoni, 2011; Isoni et al., 2011; Lichtenstein and Slovic, 2006; Loomes et al., 2003, 2010; Thaler, 1985; Tufano, 2010).

A recent attempt to explain the shaping effect in terms of a restricted set of plausible behavioral assumptions is the one carried out by Isoni (2011); Isoni et al. (2011) and it is grounded in a mild version of the theory of reference-dependent preferences (Tversky and Kahneman, 1991; Munro and Sugden, 2003; Kőszegi and Rabin, 2006). The basic idea is that individuals' satisfaction depends not only on the absolute level of consumption, given the reference-independent valuation for the good on trade, but also by the feelings of pleasure or pain associated with the awareness of having performed a good or a bad deal.

Under the hypotheses that the goodness of a deal basically depends on the difference between the actual and the expected market price (price sensitivity) and that the pain associated with bad deals is greater than the pleasure derived from same-sized good deals (bad deal aversion), a plausible account of the experimental evidence showing a tendency of the individuals' evaluations to converge towards the market price is offered[1]. The explanation of the shaping effect offered by Isoni (2011) and Isoni et al. (2011) raises some questions which are worth discussing: is the shaping effect an obvious consequence of the way the individuals understand the setting used to investigate the phenomenon itself? Does the shaping effect persist whenever information other than the market price is made available to traders? What is the effect of con-

[1]Given their price expectations, sellers involved in median price auctions with a low than median evaluation will raise their ask in such a way as to avoid bad deals, whereas sellers with high than median evaluation will lower their ask to try to make good deals. This process ensures that the distribution of asks collapse towards the median ask, which corresponds to the market price.

flicting information on individuals' preferences? Is the explanation of the shaping effect offered so far sufficiently general to cope with cases in which market feedbacks are not limited to the market price?

To investigate these questions, we repeatedly elicited individuals' WTA valuations in a simulated median price auction with the following simple characteristics: in each trading period subjects were required to submit an ask corresponding to their WTA an auctioned bad; the period specific market price was set in correspondence of the median ask and publicly displayed at the end of the trading period; subjects whose asks were lower than the median had to consume the bad, receiving a monetary payment equal to the market price. At the end of every trading period all the subjects were provided with information about the actual market price. Some subjects were further provided with information on the actual asks of individuals at either the bottom or the top of the distribution. These treatments, novel in the literature, proved insightful to investigate the process of preference formation in response to seemingly irrelevant cues accruing from markets.

Our analysis provide two main findings. First, individuals' WTA evaluations are remarkably driven by a strong tendency to conform to others, that is, by a consistent inclination of every individual to adjust her own asks coherently with what the others are expected to do, given the signals accruing from the market and transmitting information on what the others are actually doing. Second, although a clear shaping effect emerges when no other information is being provided, the provision of information about the actual behavior of well identified (groups of) individuals strikingly reduces its magnitude, suggesting that this effect may also qualify as a peculiar case of conformity.

Our results are identified by means of a methodology not yet employed in the field, which helps in distinguishing between the dynamic and the asymptotic features of preference formation - hence between short- and long-run parameters - in the presence of either adaptive or rational expectations. We rely on two main estimators: (1) a linear-dynamic panel estimator of the Arellano and Bond (1991)- type, employed under the assumption that the marginal response to treatments is constant across individuals, and (2) a random-coefficient method inspired Hsiao and Pesaran (2008) in which individual, subject-specific, linear-dynamic regressions are estimated separately for each subject.

A key result of our study, holding under different underlying models of expectation formation, is that providing the subjects with information on the asks at either the top or the bottom of the distribution roughly halves the strength of the shaping effect. Indeed, such information exerts a greater force of attraction on individuals' asks than that exercised by the observation of the market price. In particular, results from the Arellano-Bond estimators show that the reduction in the distance between one's ask ad the market price - when no other information is provided to subjects - is in the range 38—47%. This translates into a long-run conditioning factor falling in a 56—70% range.

Results from Arellano-Bond estimations also show that providing subjects with information on the asks at either the bottom or the top of the distribution noticeably changes the picture. The estimated parameter capturing the effect of such extreme information on individuals's asks show that subjects adjust a significant fraction (24—33%, depending on the estimations) of their asks towards (a central measure of) the asks they can observe. Structural parameters display an effect ranging from 44% to 66%. In addition, the provision of such information halves the size of the parameter capturing the shaping effect.

To test the robustness of our results, we provide a number of alternative empirical estimates and address individuals' heterogeneity through the use of the random coefficients method. We also employ meta-regression techniques to quantify uncertainty about alternative specifications.

The paper is organized as follows. Section 2 describes the experiment design and implementation. Section 3 describes the empirical model and the estimation strategy. Section 4 provides a discussion of the results. Section 5 concludes.

2 Experimental design

The experiment was conducted at the laboratory of experimental economics (LEE) of the University of Prague from the 4th to the 6th of March 2014. The experiment was programmed with the `z-tree` software (Fischbacher, 2007) and consisted of 6 experimental sessions. At the beginning of each session, 27 subjects were randomly assigned to three independent markets of 9 persons each. Individuals participating in the first 6 markets (sessions 1 and 2) constituted the control group. All the other individuals were involved in either a low-info treatment (sessions 3 and 4) or a high-info treatment (sessions 5 and 6).

Subjects were involved in a variant of the Vickrey (1961) multi-unit selling auction, repeated for 11 rounds and with the following characteristics. In each round subjects were asked to state their willingness to accept an auctioned bad, an harmless but disgusting mixture of Fanta, vinegar and salt (a slightly modified version of the mixture used in Ariely et al. (2003) and Tufano (2010)). Subjects whose asks were lower than the median ask (corresponding to the market price), were required to drink 60 ml of the liquid with the promise of a payment in cash equal to the market price at the end of the session.

The Vickrey median price auction is an incentive-compatible mechanism to elicit individuals' preferences, for truthful revelation is the (weakly) dominant strategy: any attempt of strategically manipulating the market price so as to increase earnings is vain. This feature of the auction was strongly emphasized in the instructions given to subjects and in public explanation.[2]

[2]All the materials inherent to the experiment (instructions, screenshots, questionnaires) are available upon request.

Auctioning an unusual item has two some desirable consequences. First, subjects could not rely on their knowledge of real world market prices in stating their WTA, because there is no real world substitute for the item on auction. Second, as consumption on the spot makes re-sale impossible, subjects were forced to focus on their true underlying valuation.

As a necessary condition for participation to the experimental auction, subjects were required to drink a 30 ml sample of the mixture. Tasting liquid allowed subjects to get all the information on its organoleptic properties that are relevant for reflecting about their private evaluation (Plott, 1996, p. 227).

At the beginning of any round the prevailing market price of the previous round was displayed on the right-hand side of the computer screen of the relevant individuals. This was the only information provided to subjects belonging to the control group.

Subjects belonging to the treatment groups were given additional information. In particular, subjects belonging to treatment 1 (low info treatment) could observe the three lowest asks of their market submitted in the previous round (the three lowest askers were exclusively informed on the market price). Analogously, subjects belonging to treatment 2 (high info treatment) could observe the three highest asks; again, the three highest askers could observe just the market price.

In both treatments, the identification of the three lowest (highest) asks was updated at each round. Display of additional information started from the third round, for the distribution of the individuals' asks was supposed to be more stable at that stage of the experiment. Given the complete anonymity of the game, subjects observing the three lowest (highest) asks could not infer who made them, neither were they informed that only information of a particular type was given to them. Indeed, the sample of extreme asks were simply displayed on the left-hand side of their screen with the label "the lowest (highest) asks are…" followed by the list of the three asks in no particular order.

At the end the auction, subjects were required to answer to a socio-economic questionnaire also concerning the perceived influence of both the median price and the sample of observed reservation prices over their asking behavior. The questionnaire controlled for possible dissonances between subjects' actual behavior and their perception of it.

Implementation

A sample of 162 undergraduate and master students from the University of Economics (Vysoka Skola Ekonomicka) and the Charles University of Prague were randomly selected from the database managed by the Laboratory of Experimental Economics (LEE) and formally recruited via e-mail. The recruitment procedure granted gender balancedness and heterogeneity in education, as students' degrees regarded different academic disciplines. Recruited subjects were not over-experienced in economic experiments.

Participant were granted a show-up fee of 100 CZK (€3.7). Subjects were asked to stay in their slots without talking to each other, so to guarantee complete anonymity during the experiment. Once assigned to their slot position, subjects found on their desk the instructions to be read and three plastic glasses, one of them filled with 30 ml of the liquid (the small sample to be tasted) and the other two with 60 ml. This procedure allowed consumption at the assigned position so to preserve complete anonymity during the experiment. Lab assistants refilled the glasses of the subjects who actually drank.

Subjects were first asked to carefully read the instructions, which were also publicly illustrated, then they were required to answer to three questions aimed at testing their comprehension of the rules of the experiment (Lusk and Shogren, 2007).

The program elicited subjects' asks through the question: "Would you be willing to accept x CZK to drink the liquid?", where the value of x ranged from a minimum of 2 CZK (€0.072) in the first question to a maximum of 100 CZK (€3.7) in the last.

Subjects could reply either "yes" or "no" to any of the 30 questions they were confronted with.[3] As soon as a subject answered affirmatively to a given question, the program summarized the implications of her choice by showing her the lowest accepted price (the minimum price at which she was willing to drink the liquid) and the highest rejected price (the maximum price at which she was not willing to drink the liquid). The subject was then asked if she was satisfied with her choice. In case of a positive answer, the subject's highest rejected price was recorded as her ask. Conversely, the elicitation procedure re-started.

It was established from the outset that only rounds 2, 6 and 9 were payoff relevant. This procedure allows for full comparability amongst the experimental groups. Subjects were told that at each round a random mechanism would have determined whether that round were payoff relevant. There was no possibility to infer that payoff relevant rounds were selected deterministically.

3 Empirical estimation and results

This section describes the assumptions, the methods, and the results obtained from the study of the data generated by the experiment.

Our experimental sample is made up of three groups, each including 54 subjects, for a total of 162 subjects who participated in all the experimental sessions. Each round was played 11 times. As it is customary, the first round was played just for illustrative purposes, so the resulting dataset has 10 time periods. The time structure of the experiment is such that players cannot observe the signals m_t (median ask) and e_t (information on extreme asks) as they emerge during the

[3]Starting from $x = 2$ CZK, the value of x was increased by 2 CZK moving from any of the first 12 questions to the subsequent one. It was then increased by 3 CZK (€0.11) until the 23th questions and by 6 CZK (€0.22) thereafter.

round they are playing. Rather, given that the relevant variables are revealed only *after* individual choices have taken place, subjects only observe the median ask and the extreme signal of the *previous* round: this drives naturally to the inclusion of expectational variables. We assume that expectations are formed adaptively (Tufano, 2010; Loomes et al., 2003): subjects form their expectations using a weighted average of all past observed values of the signals and correct their forecast errors only gradually.

The nature of our data allows for the treatment of individual heterogeneity because we have either a cross-section component of variation as well as a temporal component. This provides us with the opportunity of modeling individual differences and to take into account individual differences in treatment response. To estimate the effects of the market price and the extreme information we use two main estimators: (1) a linear-dynamic panel estimator of the Arellano and Bond (1991)-type in which it is assumed that the marginal response to treatments is constant across individuals and the observed systematic behavioral heterogeneity is due to differences in a subject-specific error term, and (2) a random-coefficient method inspired by Hsiao and Pesaran (2008) in which individual, subject-specific, linear-dynamic regressions are estimated separately for each player, giving rise to a whole distribution of individual effects. We also test whether our results are sensitive to changes in the expectations formation process.

As a wealth of alternative estimates emerges from the models, we employ a meta-analysis of our results to provide point estimates of the parameters of interest and capture uncertainty over model specification.

3.1 The empirical model

The choice of the current ask depends upon the expected value of the signals according to the following equation

$$b_t = \alpha + \beta m^e_{t+1} + \delta e^e_{t+1} + \varepsilon_t \tag{1}$$

where m^e_{t+1} and e^e_{t+1} are the expected values of the median ask and the extreme signal forecast at time t. Next, we assume that agents employ an adaptive model of expectation formation of the type

$$x^e_{t+1} = \sum_{i=0}^{\infty} \phi_i x_{t-i} \tag{2}$$

in which $x_t = m_t, e_t$ and $\phi = \phi_0 \gamma^i$. Under these assumptions, subjects forecast the signals using all the observed values of the same signals. After some algebraic manipulation, the expectation equation reduces to

$$x^e_{t+1} - x^e_t = (1 - \gamma)(x_t - x^e_t) \tag{3}$$

which can be interpreted in the following way: the revision in expectation (the term on the left-hand side of the equation) is a fraction $1-\gamma$ of the last round's forecast error. Expectations for both m_t and e_t are modeled adaptively and give raise to the following equations:

$$m_{t+1}^e - m_t^e = (1-\gamma)(m_t - m_t^e) \tag{4}$$

$$e_{t+1}^e - e_t^e = (1-\gamma)(e_t - e_t^e). \tag{5}$$

These equations suggest that the current revision of expectation (left-hand side of the equation) is a fraction $1-\gamma$ of the current forecast error. Subtracting the expression for γb_{t-1} from b_t we obtain an empirically estimable model in terms of only observed variables

$$b_t = \mu_0 + \mu_1 b_{t-1} + \mu_2 m_t + \mu_3 e_t + \eta_t \tag{6}$$

in which $\mu_0 = \alpha(1-\gamma)$, $\mu_1 = \gamma$, $\mu_2 = \beta(1-\gamma)$, $\mu_3 = \delta(1-\gamma)$, and $\eta_t = \varepsilon_t - \gamma\varepsilon_{t-1}$ is a first-order autoregressive error term. The structural parameters can be retrieved from the reduced form in the following way: $\alpha = \mu_0/(1-\mu_1)$, $\beta = \mu_2/(1-\mu_1)$, and $\delta = \mu_3/(1-\mu_1)$.[4]

3.2 Exploratory analysis

The main descriptors of our experiments are gathered in tab. 1. The mean ask for the control group (47.5, σ=24.3) is remarkably lower than the corresponding mean ask for the low (65.6, σ=25.1) and the high (60.7, σ=24.8) information groups, with the median ask also following the same pattern. On average, the low information group faced a signal equal to 47.4 (σ=20.7) whereas the high information group received a signal equal to 84.0 (σ=11.5). Unconditional

[4]The model employed by Tufano (2010), in place of the adaptive expectations assumption outlined above, uses the partial adjustment approach to modelling the change of the asks across experiment's rounds. Adopting his framework our equation would become

$$b_t - b_{t-1} = \nu_0 + \mu_2 (m_t - b_{t-1}) + \mu_3 (e_t - b_{t-1}) + \chi_t \tag{7}$$

where ν_0 is a constant and χ_t is an error term. Seen from this perspective, the parameters μ_2 and μ_3 govern the stability of the dynamic behavior of asks. With $0 < \mu_2, \mu_3 < 1$, the observed change in the value of the asks $(b_t - b_{t-1})$ mirrors the corresponding divergences between (1) the median ask and the ask at $t-1$ and (2) the extreme signal and the ask at $t-1$. While reflecting different assumptions about the underlying model, the interpretations of the parameters μ_2 and μ_3 from eq. 7 and from eq. 6 are similar: accordingly, positive values lower than one in eq. 6 for μ_2 and μ_3 are evidence for convergence of asks toward the median and the extreme signal.

The steady-state solution of eq. 7 is

$$b_t^* = \tau + \frac{\mu_2}{\mu_2+\mu_3} m_t + \frac{\mu_3}{\mu_2+\mu_3} e_t + \xi_t \tag{8}$$

where $\tau = \nu_o/(\mu_2+\mu_3)$ and $\xi_t = \chi_t/(\mu_2+\mu_3)$. In the extreme case in which $\nu_o = 0$, the equilibrium ask is located between m_t and e_t when $0 < \mu_2, \mu_3 < 1$. The lower the ratio μ_2/μ_3, the closer the equilibrium ask will be to the extreme info signal.

correlations between the variables are reported in tab. 2 for all experimental groups pooled together. To make these figures interpretable, we have transformed the data by subtracting the subject-specific mean from each variable. The statistics suggest a certain degree of autocorrelation until the second lag and a moderate association (0.334) between the ask at t and the value of extreme information.

A preliminary analysis of our regression model developed in the next sections can be found in tab. 3 in which we display a number of partial correlations between the individual ask at t and the regressors selected by the theory. Also in this case we have mean-demeaned the data to obtain a fixed-effects-like estimator. Though these estimates cannot take into account the possible endogeneity and error term autoregression issues (both of which are to be addressed shortly), they are nonetheless suggestive of the presence of linear effects of regressors which are worth investigating.

The main statistic to focus on in the table is the *semipartial coefficient of correlation*, a statistic which allows for a comparison of the relative explaining strength of different regressors, for it quantifies the variance of b_t explained by that regressor and *not* shared with other regressors. In other words, it is a measure of the idiosyncratic contribution of a given variable to the variance explained by the linear model. In particular, the squared semipartial coefficient is the gain to the R^2 due to the inclusion of a given variable, weighted by the portion of unexplained variance (Filoso, 2013). Taken together, the independent variables account for a 27.5% of the observed variance of the asks of which 52.7% (14.5/27.5) is due to the specific contributions of the individual regressors. Apart from the first lag of b, the main individual contributor is the value of extreme info. This suggests that the informative content of this variable is a carrier of specific information not already embedded into other regressors.

Finding 1 (Salience of extreme info). *The idiosyncratic contribution of the extreme information signal accounts for 3.9% of the total variance of b_t, while the idiosyncratic contribution of the median ask accounts for 1.9%.*

The autoregressive part of the model accounts for the remaining 8.7%, a figure which points to a significant factor of dynamic adjustment.

3.3 Panel data

Assuming constant marginal effect across individuals and a panel-data structure allows rewriting the reduced-form eq. (6) derived above as

$$b_{i,t} = \mu_0 + \mu_1 b_{i,t-1} + \mu_2 m_{i,t} + \mu_3 e_{i,t} + \omega_i + \eta_{i,t} : \tag{9}$$

now the variable i indexes the subjects and a new term, ω_i, captures time-invariant and subject-specific differences which account for individual factors, possibly correlated with the error term. To provide consistent estimates of our parameters of interest we turn to the literature on dynamic panel data which has been developed for data structures with large N and small T, which is precisely our case. This class of estimators (Blundell and Bond, 1998; Arellano and Bover, 1995; Arellano and Bond, 1991; Holtz-Eakin et al., 1988; Anderson and Hsiao, 1981) explicitly focuses on linear structures with endogenous regressors and autoregressive errors as it happens to be the case for the data generating processes of our experiments.[5]

As we have three experimental groups, we combine them to obtain four main estimation tables using four different sets of data: (1) only the control group; (2) the low information and the high information group; (3) the low information group; (4) the high information group. The estimation results are shown in tables 4–7.

3.3.1 Results

The results for the model including only the control group are displayed in tab. 4. The first step we take is checking the general dynamic properties of b_t: to this extent, we start with a simple OLS model augmented by cluster-corrected standard errors including either one or two lags of the dependent variable. The corresponding models are then labeled as OLS1 and OLS2. We find a strong degree of autocorrelation until the second lag: in the second OLS model the sum of the two autoregressive terms is close to one, which might warn for the presence of a unit root. Nonetheless, these value decrease considerably when a more appropriate fixed-effect estimator is employed (column FE): this suggests that much of the apparent autoregression is actually captured by the individual error term.[6]

The column ABond1 reports the results from an Arellano-Bond estimator with one-step difference GMM and robust standard errors. The following column (ABond2) reports the estimates for the same type of model this time estimated with a two-step difference GMM. The remaining columns use forward orthogonal deviations (Arellano and Bover, 1995), with one-step difference GMM (ABond3) and two-step difference GMM (ABond4). These variations on the Arellano-Bond estimator are reported to account for slight violations in the assumptions about the data structure which can be a source of bias. All two-step estimates show the finite-sample correction of variance due to Windmeijer (2005). As implied by the Arellano-Bond

[5] Compared to the case of adaptive expectations, the estimation of the model with rational expectations is much more challenging. There is no accepted way to deal systematically with this kind of expectations, because parameters are not generally separately identifiable. In the words of Nerlove et al. (2008): "The bottom line is that, if one believes in rational expectations, one is in deep trouble dealing with panel data." Accordingly, we refrain from the temptation of providing estimated figures that are potentially misleading and estimate the rational expectations model with an alternative methodology described in section 3.4.

[6] Moreover, the Levin-Lin-Chu and the Breitung unit root tests, not reported here but available on demand, provide evidence of no unit roots in the panel.

approach, the instruments are the lagged values of the dependent variable (since they are predetermined at t) and the lagged value of the median ask. As an exogenous instrument, we use temporal dummies for the round being played.[7] To limit instruments proliferation, we restrict their number to be below the number of groups included in the estimation (Roodman, 2009): we include no more than three periods lags. Tables 5–7 have exactly the same structure, except for the experimental groups included and for the addition of lagged values of the extreme signal as instruments to control for endogeneity of regressors.

As expected, the autoregressive terms have values lying between the corresponding OLS2 and FE estimates, as suggested by Roodman (2009). This is reassuring since, in normal conditions, OLS estimates tend to be biased upward, while FE are biased downward; also, the sum of the two autoregressive terms is well below one, which suggests dynamic stability. The statistics displayed at the bottom of the table display that residuals are generally not correlated at the second lag (ABond AR(1) and AR(2) tests) and the overidentifying restrictions on instruments are valid (Hansen test).

Finding 2 (Short-run shaping effect)**.** *The results from the Arellano-Bond estimators show that the market price is an attractor for subjects' choices, with this effect in the reduced-form estimates being between 38% and 47%.*

In practice, subjects adjust a significant fraction of their ask toward the median. The corresponding structural parameters can be read in the second main section of the table reporting the values for β. These are obtained using the algebra of the section 3.1, with the distinction that γ now is equal to one minus the *sum* of the two autoregressive terms. Since the structural parameters β and δ are obtained through nonlinear combinations of the reduced-form parameters, the corresponding standard errors are obtained by the delta method (Oehlert, 1992).[8] With this derivation we find that

Finding 3 (Long-run shaping effect)**.** *In the long run, asks depend on the market price by a factor between 56% and 70%.* [9]

The results for the estimations which include the variable e_t (the average value of the extreme information as displayed on subjects' screen) can be seen mainly in the tab. 5 where we pool the low and the high information groups together. The results indicate a strong degree of dependence from extreme information. The corresponding estimated parameter in the reduced

[7]The exclusion of this instrument changes estimated values only marginally. The corresponding tables, not included here, are also available upon request.

[8]Given the substantial equivalence between the adaptive expectations and the partial adjustment model, in what follows we limit our considerations almost exclusively to the long-run estimates of parameters. The same considerations extend immediately to the short-run parameters.

[9]These values are also sensibly distant from their corresponding FE estimates: using $(\beta_{FE} - \beta_{AB})/\beta_{AB}$ as an indicator of estimation bias due to using FE instead of ABond, we find that this value ranges from 12% to 40%.

form varies from 24% to 33%. The figures for the structural parameter δ range from 44% to 66%. Moreover, this is a striking result since the estimated values for δ are uniformly larger than the corresponding values for β (18—24%). To sum up the results for βs from tab. 4 to those from tab. 5 we note that

Finding 4 (Extreme info and shaping effect #1). *The inclusion of extreme information more than halves the effect of the median ask.*

This corroborates our theoretical prior that feeding subjects with normatively irrelevant information can hijack their equilibrium asks significantly, also severely reducing the strength of the market price.

The following two tables estimate the same model as above, now focusing only on the low information group and on the high information group, respectively. These separate estimations are carried out upon the interest about possible asymmetric effects of the low and the high information treatments. The extreme signal was explicitly displayed on the screen as information derived from the highest and the lowest asks. Cognitively, this signal conveys two main messages, one numerical, the other positional: (1) the numerical value of three asks, of which we consider salient their average and (2) the information that this average is representative of one tail of the asks' distribution. The observed effect, δ, is a mixture of these two components.

Finding 5 (Extreme info and shaping effect #2). *Since the estimated values of δ are positive for the high and low info treatment groups, the extreme info e_t is an attractor for asks.*

Finding 6 (Extreme info and shaping effect #3). *Equilibrium asks lie between m_t and e_t, although closer to e_t.*[10]

The parameters estimated for the low information treatment have a ratio δ/β ranging from 1.9 to 2.4, while the corresponding numbers for the high information treatment lie between 1.0 and 1.2. This suggests that the relative strength of the extreme information treatment is more pronounced when subjects are given information on the lowest asks.

3.3.2 Meta-Regression

Our estimates of β and δ exhibit a certain degree of variation across different experimental groups and estimators: to provide a robust estimation of an overall effect and an analysis of sensitivity of our results to alternative specifications we resort to the method of meta-regression (Stanley and Jarrell, 1989). Usually, this method is employed to estimated effects from a body of existing literature, but nothing prevents its use when an intra-study aggregation procedure is needed (Leamer, 1983) to account for model uncertainty. This allows us to average the estimated

[10]See footnote 4.

effects over several specification and to separate the predictive signal of competing assumptions from the inevitable noise.

Meta-regression's outcomes are point estimates for β and γ which take into account both *between-subjects* and *between-specifications* variation: possibly, additional variables reflecting the features of the estimates can be included. In practice, we estimate the model

$$p_j = \bar{p} + \sum_i \pi_i c_i + u_j + v_j$$

where p_j is the parameter of interest (which, in our case, is either β or δ) as it comes out from the j-th specification, $\bar{p}$ is the estimated value of the parameter for the default specification, c_i is a dummy for the i-th feature of the j-th specification, u_j is an error term reflecting within-subjects variance, v_j is another error term reflecting between-specifications variance. The two dummies introduced to capture estimation-specific features are coded in the following way: *Extreme low info* = 1 for the sample including just the first treatment, 0 = otherwise; *Extreme high info* = 1 for the sample including just the second treatment, 0 = otherwise. The estimates are obtained using a variance-weighted least squares method in which every parameter β and δ is weighted according to its original standard error.[11]

Finding 7 (Meta-regression for δ)**.** *Our best estimate of δ is 0.526 (s.e.=0.090, p=0.000).*

This value is obtained by the joint inclusion of the first and the second treatment samples. The estimates carried out over the two experimental treatment groups separately do not seem to add new relevant information. To some extent, the value of δ increases by a factor of 0.116 for the high info treatment, though the marginal effect is noisily estimated.

Finding 8 (Meta-regression for β)**.** *The estimated value of β for the control group is 0.638 (s.e.=0.077, p=0.000).*

This is a result partly consistent with previously obtained results (Tufano, 2010). Here the baseline value is calculated using only the control group to make the figure comparable to the values estimated in the literature on the shaping effect. The inclusion of the estimates from the extreme low and high info groups drives β to a value of 0.415 (0.638—0.223): here we find confirmation that

Finding 9 (Meta-regression, decrease of β due to e_t)**.** *Providing the subjects with information on extreme asks significantly decreases the preference-shaping effect of the market price.*

Moreover, this effect is strongly at work in the case of the provision of extreme low information: here the values of β drops to 0.283 and it is quite precisely estimated (p=0.006). Finally,

[11] The estimation of the parameters is implemented through a random-effects estimator using the Stata command `metareg`. Standard error estimation follows Knapp and Hartung (2003).

also the provision of extreme high information has a moderating effect on β of -0.136, but the high standard error of its estimate (0.121) suggests caution in interpreting it as a precise cue.

The parameter γ governing the speed of adjustment in expectations estimated for the control group is 0.682 (s.e.=0.052, p-value=0.000), which suggests that subjects adjust their expectation by a factor of 31.8% of the previous round's forecast error. We also find that the provision of extreme information decreases γ by a factor equal to 0.168 (s.e.=.072, p-value=0.058), driving it down to 0.515, resulting in an increased speed of expectations adjustment.

3.4 Addressing individual heterogeneity: random coefficients

Next, we tackle the issue of individual heterogeneity. Instead of putting a severely straight jacket on our data under the assumption of a single effect common to all subjects, we consider individual specific parameters as random draws from an unrestricted distribution of causal effects. With this assumption, we can rewrite eq. 9 in the following way:

$$b_{i,t} = \mu_{0,i} + \mu_{1,i} b_{i,t-1} + \mu_{2,i} m_{i,t} + \mu_{3,i} e_{i,t} + \omega_i + \eta_{i,t} \tag{10}$$

where now the parameters μ are allowed to vary over each i-th subject in the spirit of Hsiao and Pesaran (2008). Then, to aggregate this information into a single point estimate, we employ a random coefficients estimator in which the marginal effects of the independent variables are to vary across subjects.

The reduced form of the adaptive expectations model (1) contains an error term, η_t which follows an AR(1) process. In this case, an OLS estimator produces biased results and we must resort to an IV estimator. Since, with regard to time t, the variables m_{t-1} and e_{t-1} are predetermined and uncorrelated to η_t but correlated to b_{t-1}, we use them as instruments for b_{t-1}. The structure of lags is also of some concern, since the Arellano-Bond estimator suggest a two lags specification, while we cannot rule out the possibility that some subjects may be characterized by just one lag. We choose the optimal lag structure on an individual basis, using the Akaike criterion corrected for small sample bias.

The problem with estimating the rational expectations model (12) mostly lies in the necessity of obtaining valid estimates of m_t^e and e_t^e: we need an estimate which is (1) uncorrelated with the level of the same signal and (2) constructed taking all available information into account. In our case, this can be attained by an IV estimator using m_{t-1} and e_{t-1} as instruments for m_t and e_t. Finally, to obtain the maximum informational efficiency, we employ a GMM estimator for both IV models. *We estimate the empirical model separately for each subject in our experimental session.* This provides us with the opportunity of studying the whole distribution of the effects of interest.

Once the previously described estimations have taken place, we are left with a vector of parameters $[\beta_i, \delta_i, \gamma_i]$ for each subject and with another vector $[\sigma^2_{\beta_i}, \sigma^2_{\delta_i}, \sigma^2_{\gamma_i}]$ of the corresponding estimated variances for the same parameters. To get an estimate of the average value of the parameter for the whole population, we envision the observed effects $y_i \in \{\beta_i, \delta_i, \gamma_i\}$ as made up of the sum of the *true* effect θ plus two error terms: of these, the first, u_i reflecting the variability of the estimated effects between subjects and the second, v_i, reflecting the uncertainty of the effect within subjects, according to the formula

$$y_i = \theta + u_i + v_i \tag{11}$$

where $u_i \sim \mathcal{N}(0, \tau^2)$, $v_i \sim \mathcal{N}(0, \sigma^2)$, and the term τ^2 represents the between-subjects variance. To obtain a summary statistic for our parameters of interest we employ a variance-weighted least squares estimator. In practice, every structural parameter is weighted according to its standard error.[12]

This flexible estimation procedure allows us with possibility of experimenting with an alternative assumption about expectation formation, namely *rational expectations*. In this approach (Muth, 1961) the prediction error $\psi_t = x_t - x^e_t$ must be unbiased, since otherwise there would be a systematic component in the forecast error which subjects should be able to correct; in other words, we require that $E(\psi_t) = 0$. Moreover, the prediction error must be uncorrelated with the entire information set that is available at time $t-1$, a set we name I_{t-1}, i.e. $x^e_t = E(x_t|I_{t-1}) + \psi_t$, which amounts to $x_t = x^e_t + \psi_t$. To obtain an estimate of x^e_t we take the expectation of x such that the prediction error is uncorrelated to the information set at $t-1$. This method is implemented through an IV estimator. Since m_t and e_t are unobserved at t we can write equation (1) in the following way:

$$b_t = \alpha + \beta m^e_t + \delta e^e_t + \varepsilon_t. \tag{12}$$

in which we make the following substitutions: $m^e_t = m_t + \psi^0_t$ and $e^e_t = e_t + \psi^1_t$. Valid instruments for estimating the expectational variables are m_{t-1} and e_{t-1}.

3.4.1 Results and meta-regression

All our main results from estimation already transformed in the terms of the structural models are collected in tab. 9.[13] The first results we show are for the adaptive expectations model *without* the extreme signal e_t. As expected, we find evidence that the effect of the median ask, through β, is substantially larger than zero and resembling quite closely the corresponding values obtained through the Arellano-Bond estimator. With regard to the introduction of e_t, its

[12]The general formula for the variance-weighted least squares which takes into account both the variance within subjects and the variance between subjects is due to Harbord and Higgins (2008).

[13]Given the short temporal span of the data, a time-series-like estimation can either fail from some subjects or do not provide the whole set of parameters $[\beta_i, \delta_i, \gamma_i]$: in this case, these observations were dropped from the final estimation.

estimated marginal effect follows the same pattern found in the previous panel estimation: the value of β drops significantly and its predictive power looks much absorbed by δ. More specifically, the effect of extreme high information now plays a major impact on the magnitude of δ rather than β: while β is approximately constant across both experimental groups, the parameter δ looks stronger in the extreme high group (+12%). On the whole,

Finding 10 (Adaptive expectations, random coefficients). *The estimates obtained through the random coefficients method corroborate the results obtained by the Arellano-Bond method.*

Compared to the adaptive expectations approach, the results for the rational expectations model show uniformly lower parameter estimates. To check whether this difference is just a noisy feature of the data or a genuine one, we perform a meta-regression of our results whose results are collected in tab. 10. From this exercise we find that

Finding 11 (Meta-regression, rational expectations). *The values of β are not significantly affected by the assumption of rational expectations.*[14]

As in the case of the Arellano-Bond estimation,

Finding 12 (Meta-regression, rational expectations). *The provision of additional extreme information does decrease markedly the value of β. Furthermore, contrary to the Arellano-Bond, the decrease appears very similar when either extreme high or low information is being fed to the subjects.*

The value of δ is very similar to the value estimated in tab. 8 and the value does not seem to change across the extreme low and high groups, while there is some (imprecise) evidence that a rational expectations modeling does impact the value of extreme information negatively, thereby lowering the estimated δ.

4 Discussion

As emphasized in the previous section, providing the subjects with information on *extreme* asks more than halves the strength of the shaping effect; moreover, such information exercises a greater force of attraction on individuals' asks than that exercised by the observation of the market price. These results are consistent with alternative underlying models of expectation formation. Overall they suggest a strong role for conformity in individual's decision making.

[14] This is a nice illustration of the phenomenon emphasized by Gelman and Hill (2007, p. 22—23). Though *prima facie* evidence would suggest that some difference between the rational and the adaptive expectations models *does* exist, the application of a correct statistical testing procedure can often produce counterintuitive results. Gelman and Hill's warning is actually the main motivation behind our extensive use of meta-regression in this article.

As reported above, over the whole population, Arellano-Bond estimators catch a strong dependence of individuals' asks from the observation of others' asks, with the associated structural parameter indicating a high degree of convergence: estimates point to a $44 - 66\%$ reduction in the distance between one's asks and (a central measure of) the observed asks (Table 5). Besides emerging from the experimental evidence on actual individuals' decisions, a role for conformity is also supported by the individuals' subjective perception of the way the information they receive condition what they do, even if individual perceptions is less pronounced than what emerges from their behavior.[15]

Such a role has to find an explanation. The remaining parts of this Section are indeed devoted to clarify why conformity plays a role in a setting in which there should be no role for social learning.

4.1 Why do individuals conform?

As far as conformity is concerned, a key distinction is between informational and normative motivation to conform to a group norm, depending on whether the individual's aim is that of being correct or that of getting a positive appraisal from the others.

As for all the other sciences (Morgan and Laland, 2012), also the hypotheses put forward by the economic literature to explain conformity basically echo this distinction. Conformity would therefore be either the result of a rational desire of avoiding the disutility connected with punishment and social exclusion (Bernheim, 1994), or a way to economize on the costs of acquiring information (Bikhchandani et al., 1992).

In the experimental auctions on which the present study is based, there is no reason to conform grounded in the desire for a positive social appraisal, for the experiment is held under conditions of anonymity. The motivation behind conformity can be rather seen as one of being correct, as one of providing the right answer, whatever this can signify within the boundaries of our experimental setting.

In a recent study on the evolutionary basis of human social learning (Morgan and Laland, 2012), subjects were involved in computer based binary choice tasks - such as, for example, the 'mental rotation task', in which they had to decide whether two shapes seen from different angles were the same shape or different shapes - with the possibility of resorting to social information to guide decision making. Social information simply presented subjects with the choices of a number of other individuals (demonstrators). Conditional on factors such as the degree of consensus among demonstrators, the task's difficulty or the cost of asocial learning, a strong evidence in

[15]Data from the questionnaires show that only 27% of subjects does not acknowledge the influence of the market price over the asking behavior. Interestingly, subjects are less ready to acknowledge the influence of identified individuals: only 34% of subjects in the high and the low info groups acknowledge the influence of observing their fellows' asks.

favour of a wide variety of social learning strategies was offered.

A key distinction between Morgan and Laland (2012)'s study and ours is that, while subjects involved in their experimental tasks were provided with adequate material incentives to give the correct answer to the puzzle they were confronted with (and this drove behavior in such a way as to use social information adaptively), in our study there is not an objectively right answer to be given, so it is legitimate to wonder why individuals should be subject to the same strong tendency to conform. This attitude to conform, as argued by (Samuelson, 2004), may be the result of an evolutionary process which shaped preferences in such a way as to compensate the individuals for incomplete environmental information[16]. Grounding on similar arguments, studies on non-human animal behavior also point to a strong role for conformity in social learning; evidence is for example provided that conformity shapes primates' foraging decisions (van de Waal et al., 2013). Our study identifies a strong role for conformity even in cases in which conforming does not provide any clear advantage to the individual adopting such a behavioral strategy. This suggests that the *copy-when-uncertain rule* has much of an hard-wired adaptive response whenever individuals are subject to pervasive uncertainty.

4.2 Imprecision as a source of conformity

As it is usually the case when an unfamiliar good is being traded, our subjects did not presumably have any clear idea of which ask corresponded to their true underlying preference for the item being sold. In stating their WTA evaluations they were likely subject to imprecision, at least in the first rounds of the experiment. Our claim is that such imprecision never got completely fixed, and constituted a key condition for conformity to emerge. In this perspective, conformity would be a response to the informational problem individuals face given the inability to determine their true preferences with certainty. In what follows we give an interpretation of our results grounding in these ideas.

The notion of preference imprecision, elaborated by Graham Loomes and colleagues (Loomes, 2005; Butler and Loomes, 2007) recognizes the existence of an (individual specific) imprecision interval from which any subject draws her asks. Such imprecision interval must contain all the alternatives which are not unanimously preferred: preferred according to any of the preference relations candidates to produce the individuals' true preference ordering over the

[16] In Samuelson's model, different consumption levels give rise to different survival probabilities, depending on the state of the environment and some agent-specific characteristics. Conjecturing path dependency in the realized state of the environment - i.e. a good state of the environment (rewarding high consumption) is more likely than a bad state (rewarding low consumption) after a good state is realized – makes the availability of sample information on the consumption levels of survival individuals crucial in forming a reliable expectation on the state of the environment which is more likely to come about. Conforming to the most frequent consumption pattern is henceforth an efficient strategy to enhance one's own survival probability, for it increases the probability of matching the choice of the consumption level with the right state of the environment. In Samuelson's theory, relative consumption's effects are a direct consequence of a copy-when-uncertain rule.

alternatives, what we term active preference relations[17].

Supporters of the *discovered preferences hypothesis*, or of some of its milder versions, such as, for example, the *market discipline hypothesis*, argue that if the individuals were provided with adequate incentives to take the feedbacks coming from the market properly into account, they would eventually discover what their true underlying preferences are (Plott, 1996; Binmore, 1999). In other words, they argue that at some finite time t, the set of active preference relations is singleton, or that, for any couple of still active preference relations, the material consequences of choosing according to any of them are the same for the relevant individual. Thus, it is as if the individual decided according to her true preference relation.

The interpretation we give of our experimental results, supporting the competing conjecture that preferences are endogenously determined by the market process itself, is that a better understanding of the characteristics of the good being sold (and possibly of the rules governing market interaction) helps the individuals to squeeze their imprecision intervals only up to a certain point; some uncertainty, in any case, remains[18]. On this account, in the course of the market process individuals discover that portion of their preferences which is reference-independent (Munro and Sugden, 2003; Kőszegi and Rabin, 2006; Isoni, 2011), on which the size and the location (the boundaries) of their imprecision interval depends. Their actual choices within the imprecision intervals are, to a large extent, determined by the market process itself.

Loosely speaking, experienced subjects can confidently state that they would be willing to drink for a given amount of money, let us say 3, or that they would not be willing to drink for, let us say, 1€; they would nevertheless be very embarrassed if asked whether their true WTA evaluation is 2€, 2.01€, or something in between. Hence, their reference-independent preferences are a guide for choice, but only limited so; only up to a certain point; once the boundaries of the imprecision intervals are determined, individuals need a different criterion than looking inside themselves to determine their WTA evaluations.

In this light, our experimental results (and the results reported in the related literature on shaping) can also be read as a confirmation that the connection between prices and WTA evaluations is even stricter than what is usually reported. Indeed, not only prices depend on WTA evaluations given supply side conditions; WTA evaluations also depend on prices; in particular, prices work as catalysts: it is with respect to them that individuals elaborate what their WTA evaluations are. Without such a reference point, individuals would be fairly disoriented, and the

[17]We conceive a preference relation as a binary relation endowed with the usual properties of completeness and transitivity.

[18]As far as the signals (feedbacks) received by the individuals during the auction are concerned, it is possible to distinguish between normatively relevant signals, conveying information on the desirability of consuming the item being sold, from normatively irrelevant signals, not conveying such information.

In our setting, information about the median ask (which corresponds to the market price), the "extreme asks", or about one's own position with respect to them, are all normatively irrelevant signals. n a reference independent model of preferences, such information should be however regarded as silent as far as the problem of determining one's own preferences is considered.

idea of returning to themselves to find out what their true WTA evaluations are (something that resembles Augustine of Hippo's motto, *in our interior the truth resides*) is a vacuous one.

On this account, we believe that a reasonable explanation of our experimental results must not neglect that while the market price conveys a synthetic statistic of the whole distribution of asks, the sample of *extreme* asks provides direct information on other subjects' WTA evaluation. When individuals observe just the market price, they can compare their WTA evaluations with the moderate preferences of the marginal subject and assume the median ask as the standard of value for the auctioned bad; put differently, the median ask is perceived as the prototypical valuation of the auctioned bad (Niedrich et al., 2001). Information on extremely low (high) asks allows the individual to infer her position relative to a reference group at the tail of the distribution. More importantly, the individual gets direct information on others' reservation prices, which evidently do not synthesize the whole distribution of asks but can be perceived as exemplars (i.e., instances) of the distribution of asks and not as a prototype, like the median ask. What our results show is that direct observation of the individual's actual behavior is more salient to the individual and exercises a greater force of attraction.

As a last step, we believe that it is necessary to illustrate why the theory offered by Isoni (2011); Isoni et al. (2011) to explain the shaping effect does not give reason of conformity as it emerges from our experiment. To this end, consider for simplicity our low info treatment. From our understanding, what their theory predicts is a tendency of the asks to increase. The argument may go as follows: any of the three individuals whose asks are at the bottom of the distribution is obtaining a surplus, given by the difference between the market price and her WTA. In an attempt to avoid a bad deal, i.e. in an attempt to avoid exchanges which would imply a surplus not comparable to that obtaining by these individuals, any other subject would increase her ask. As the previous Section instead showed, this is however, the opposite of what the experimental evidence tells us.

5 Conclusions

In this paper we reported results from an experiment testing the differential impact of presumably irrelevant signals over WTA valuations for an auctioned bad. We found evidence supporting the shaping hypothesis: when only information about the market price is provided, asks are conditioned by a factor between 38% and 47%; in the long run this translates in a conditioning factor falling in the 56—70% range.

We however found strong evidence supporting the hypothesis of conformity: individuals' WTA evaluations are remarkably driven by a strong although unexpected tendency to conform to the asks either at the bottom or the top of the distribution by a factor of 44-66%, moreover the provision of extreme information more than halves the influence of the market price on individual's asks.

We interpret these results by relying on a notion of conformity as an adaptive response to a problem of preference imprecision in market contexts where uncertainty is pervasive. More precisely, we claim that whenever individuals are uncertain about their own preferences, they rely on signals not conveying any information on their private utility. The adjustment of individuals' asks towards the market price is hence to be interpreted as a process of adaptation of individuals' WTA valuations to the moderate preferences of the marginal subject, in a context in which direct information on others' asks deliver more salient cues. In this light, conforming to direct information on others' asking behavior rather than to the market price is to be seen as a more effective heuristics to reduce individuals' uncertainty about their own preferences.

6 Tables of estimates

6.1 Descriptive statistics

Table 1: Descriptive statistics

		ask		Median		Extreme info	
Experimental group	N	Mean	σ	Mean	σ	Mean	σ
Control	540	47.47	24.28	43.33	14.96		
Low info treatment	540	65.63	25.14	62.62	18.77	47.42	20.65
High info treatment	540	60.71	24.78	56.75	16.48	84.03	11.46
Total	1620	57.94	25.88	54.23	18.64	65.72	24.78

Table 2: Cross correlations

	b_t	b_{t-1}	b_{t-2}	Median
b_{t-1}	0.417			
	[0.000]			
b_{t-2}	0.298	0.422		
	[0.000]	[0.000]		
Median	0.282	0.173	0.095	
	[0.000]	[0.000]	[0.005]	
Extreme info	0.326	0.160	0.131	0.334
	[0.000]	[0.000]	[0.000]	[0.000]

Note — The table reports unconditional correlation coefficients. p-values in brackets. All experimental groups are included.

Table 3: Regression anatomy

	Partial correlations				
	Simple	Semipartial	$(\text{Simple})^2$	$(\text{Semipartial})^2$	p-value
b_{t-1}	0.304	0.272	0.093	0.074	0.000
b_{t-2}	0.134	0.115	0.018	0.013	0.000
Median	0.159	0.137	0.025	0.019	0.000
Extreme info	0.226	0.197	0.052	0.039	0.000
Variance explained by idiosyncratic contributions				0.145	
Variance common to all independent variables				0.130	
Variance explained by the linear model (R^2)				0.275	

Note — The dependent variables is b_t. The groups included in the estimation are the high and the low info. The table displays an analysis of model's explained variance as the sum of the shared and the idiosyncratic contribution of each independent variable. The variance explained by the idiosyncratic contribution of the j-th independent variable (correlated with b_t but uncorrelated with other regressors) is measured by its semipartial squared correlation coefficient.

6.2 Panel data estimation

Table 4: Panel estimation — Control group

	OLS1	**OLS2**	**FE**	**ABond1**	**ABond2**	**ABond3**	**ABond4**
b_{t-1}	0.816	0.672	0.237	0.263	0.302	0.275	0.284
	(0.054)	(0.061)	(0.070)	(0.074)	(0.076)	(0.066)	(0.066)
	[0.000]	[0.000]	[0.001]	[0.001]	[0.000]	[0.000]	[0.000]
b_{t-2}		0.208	0.013	0.021	0.030	0.046	0.050
		(0.056)	(0.068)	(0.062)	(0.060)	(0.074)	(0.076)
		[0.000]	[0.849]	[0.734]	[0.612]	[0.532]	[0.513]
Median	0.216	0.165	0.525	0.453	0.375	0.468	0.418
	(0.066)	(0.054)	(0.080)	(0.125)	(0.128)	(0.082)	(0.086)
	[0.002]	[0.004]	[0.000]	[0.001]	[0.005]	[0.000]	[0.000]
β		1.368	0.700	0.632	0.562	0.690	0.627
		[0.000]	[0.000]	[0.001]	[0.002]	[0.000]	[0.000]
γ		0.121	0.750	0.716	0.667	0.679	0.666
		[0.008]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Observations	486	432	432	378	378	378	378
Subjects			54	54	54	54	54
Hansen test p-value				0.302	0.302	0.312	0.312
ABond AR(1) test				0.002	0.001	0.001	0.001
ABond AR(2) test				0.454	0.646	0.273	0.554

Note — The dependent variable is b_t. Standard errors in parentheses, p-values in brackets. The estimation is carried out over the entire sample using the Stata command `xtabond2`. The model ABond1 is an Arellano-Bond estimator obtained using lagged values (from $t-1$ to $t-3$) of ask, median ask, and extreme signal as instruments for the current ask. The model also includes the time variable as an additional instrument; the estimator is a one-step difference GMM. ABond2 employs the same specification of variables, but uses a two-step estimator. ABond3 uses the forward orthogonal deviations transform instead of differencing and a one-step estimator. Finally, ABond4 uses the forward orthogonal deviations transform instead of differencing and a two-step estimator. All models include dummy variables for each round of the experiment.

Table 5: Panel estimation — Low and high information treatments

	OLS1	**OLS2**	**FE**	**ABond1**	**ABond2**	**ABond3**	**ABond4**
b_{t-1}	0.882	0.622	0.299	0.333	0.321	0.346	0.319
	(0.026)	(0.071)	(0.072)	(0.114)	(0.103)	(0.106)	(0.091)
	[0.000]	[0.000]	[0.000]	[0.004]	[0.002]	[0.001]	[0.001]
b_{t-2}		0.289	0.129	0.167	0.160	0.158	0.140
		(0.065)	(0.045)	(0.048)	(0.037)	(0.049)	(0.040)
		[0.000]	[0.005]	[0.001]	[0.000]	[0.002]	[0.001]
Median	0.127	0.118	0.221	0.181	0.194	0.205	0.241
	(0.030)	(0.030)	(0.072)	(0.120)	(0.072)	(0.088)	(0.060)
	[0.000]	[0.000]	[0.003]	[0.136]	[0.008]	[0.021]	[0.000]
Extreme info	0.009	0.013	0.354	0.320	0.258	0.330	0.237
	(0.015)	(0.014)	(0.099)	(0.102)	(0.065)	(0.094)	(0.077)
	[0.542]	[0.377]	[0.001]	[0.002]	[0.000]	[0.001]	[0.003]
β		1.318	0.386	0.362	0.374	0.413	0.446
		[0.000]	[0.001]	[0.150]	[0.008]	[0.019]	[0.000]
γ		0.090	0.572	0.500	0.519	0.496	0.541
		[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
δ		0.140	0.619	0.640	0.497	0.664	0.439
		[0.366]	[0.001]	[0.006]	[0.001]	[0.004]	[0.005]
Observations	972	864	864	756	756	756	756
Subjects			108	108	108	108	108
Hansen test p-value				0.918	0.918	0.899	0.899
ABond AR(1) test				0.000	0.005	0.000	0.004
ABond AR(2) test				0.206	0.501	0.190	0.421

Note — The dependent variable is b_t. Standard errors in parentheses, p-values in brackets. The estimation is carried out over the entire sample using the Stata command `xtabond2`. The model ABond1 is an Arellano-Bond estimator obtained using lagged values (from $t-1$ to $t-3$) of ask, median ask, and extreme signal as instruments for the current ask. The model also includes the time variable as an additional instrument; the estimator is a one-step difference GMM. ABond2 employs the same specification of variables, but uses a two-step estimator. ABond3 uses the forward orthogonal deviations transform instead of differencing and a one-step estimator. Finally, ABond4 uses the forward orthogonal deviations transform instead of differencing and a two-step estimator. All models include dummy variables for each round of the experiment.

Table 6: Panel estimation — Low information treatment

	OLS1	**OLS2**	**FE**	**ABond1**	**ABond2**	**ABond3**	**ABond4**
b_{t-1}	0.859	0.609	0.253	0.210	0.222	0.197	0.200
	(0.042)	(0.112)	(0.104)	(0.147)	(0.153)	(0.146)	(0.148)
	[0.000]	[0.000]	[0.018]	[0.159]	[0.153]	[0.181]	[0.182]
b_{t-2}		0.280	0.131	0.180	0.184	0.161	0.168
		(0.102)	(0.063)	(0.063)	(0.067)	(0.069)	(0.072)
		[0.008]	[0.043]	[0.006]	[0.008]	[0.023]	[0.024]
Median	0.066	0.032	0.163	0.171	0.163	0.175	0.190
	(0.069)	(0.066)	(0.088)	(0.089)	(0.091)	(0.093)	(0.092)
	[0.339]	[0.633]	[0.070]	[0.061]	[0.079]	[0.067]	[0.044]
Extreme info	0.082	0.107	0.414	0.410	0.389	0.417	0.375
	(0.064)	(0.066)	(0.150)	(0.161)	(0.160)	(0.162)	(0.152)
	[0.204]	[0.112]	[0.008]	[0.014]	[0.018]	[0.013]	[0.017]
β		0.287	0.264	0.280	0.275	0.272	0.301
		[0.620]	[0.058]	[0.056]	[0.069]	[0.057]	[0.026]
γ		0.111	0.616	0.610	0.594	0.641	0.633
		[0.005]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
δ		0.968	0.672	0.672	0.655	0.651	0.592
		[0.094]	[0.004]	[0.006]	[0.013]	[0.005]	[0.016]
Observations	486	432	432	378	378	378	378
Subjects			54	54	54	54	54
Hansen test p-value				0.400	0.400	0.466	0.466
ABond AR(1) test				0.000	0.053	0.001	0.056
ABond AR(2) test				0.112	0.424	0.112	0.412

Note — The dependent variable is b_t. Standard errors in parentheses, p-values in brackets. The estimation is carried out over the entire sample using the Stata command `xtabond2`. The model ABond1 is an Arellano-Bond estimator obtained using lagged values (from $t-1$ to $t-3$) of ask, median ask, and extreme signal as instruments for the current ask. The model also includes the time variable as an additional instrument; the estimator is a one-step difference GMM. ABond2 employs the same specification of variables, but uses a two-step estimator. ABond3 uses the forward orthogonal deviations transform instead of differencing and a one-step estimator. Finally, ABond4 uses the forward orthogonal deviations transform instead of differencing and a two-step estimator. All models include dummy variables for each round of the experiment.

Table 7: Panel estimation — High information treatment

	OLS1	**OLS2**	**FE**	**ABond1**	**ABond2**	**ABond3**	**ABond4**
b_{t-1}	0.900	0.635	0.347	0.305	0.290	0.290	0.280
	(0.031)	(0.081)	(0.096)	(0.142)	(0.132)	(0.140)	(0.151)
	[0.000]	[0.000]	[0.001]	[0.036]	[0.032]	[0.043]	[0.070]
b_{t-2}		0.289	0.131	0.141	0.143	0.121	0.128
		(0.072)	(0.063)	(0.074)	(0.038)	(0.063)	(0.048)
		[0.000]	[0.043]	[0.062]	[0.000]	[0.059]	[0.011]
Median	0.079	0.087	0.289	0.304	0.260	0.311	0.280
	(0.047)	(0.047)	(0.110)	(0.117)	(0.091)	(0.121)	(0.121)
	[0.097]	[0.069]	[0.011]	[0.012]	[0.006]	[0.013]	[0.024]
Extreme info	0.070	0.048	0.352	0.356	0.256	0.361	0.287
	(0.047)	(0.051)	(0.146)	(0.150)	(0.115)	(0.152)	(0.150)
	[0.145]	[0.352]	[0.020]	[0.021]	[0.030]	[0.021]	[0.061]
β		1.142	0.554	0.549	0.458	0.527	0.474
		[0.013]	[0.004]	[0.005]	[0.011]	[0.003]	[0.016]
γ		0.076	0.522	0.554	0.567	0.590	0.592
		[0.011]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
δ		0.632	0.674	0.642	0.451	0.612	0.485
		[0.339]	[0.028]	[0.045]	[0.048]	[0.046]	[0.115]
Observations	486	432	432	378	378	378	378
Subjects			54	54	54	54	54
Hansen test p-value				0.995	0.995	0.861	0.861
ABond AR(1) test				0.001	0.051	0.002	0.056
ABond AR(2) test				0.729	0.849	0.871	0.892

Note — The dependent variable is b_t. Standard errors in parentheses, p-values in brackets. The estimation is carried out over the entire sample using the Stata command `xtabond2`. The model ABond1 is an Arellano-Bond estimator obtained using lagged values (from $t-1$ to $t-3$) of ask, median ask, and extreme signal as instruments for the current ask. The model also includes the time variable as an additional instrument; the estimator is a one-step difference GMM. ABond2 employs the same specification of variables, but uses a two-step estimator. ABond3 uses the forward orthogonal deviations transform instead of differencing and a one-step estimator. Finally, ABond4 uses the forward orthogonal deviations transform instead of differencing and a two-step estimator. All models include dummy variables for each round of the experiment.

Table 8: Meta-regression of dynamic panel estimation results

	β			δ		
	est.	s.e.	p	est.	s.e.	p
Parameter	0.638	(0.077)	[0.000]	0.526	(0.090)	[0.000]
Extreme low info group	-0.355	(0.105)	[0.006]	0.116	(0.152)	[0.463]
Extreme high info group	-0.136	(0.121)	[0.281]	0.003	(0.167)	[0.986]
Extreme low + high info groups	-0.223	(0.106)	[0.056]			

Note — The dependent variables are the parameters β and δ. Standard errors in parentheses, p-values in brackets. The estimation of β and δ is performed over the various samples using the Stata command `metareg` which implements a random-effects meta-regression.

6.3 Heterogeneity reconsidered

Table 9: Random coefficients estimation

	Expectations							
	Adaptive				Rational			
	Control	Treatments			Control	Treatments		
		Both	Low	High		Both	Low	High
β	0.509	0.208	0.205	0.222	0.716	0.156	0.252	0.127
	(0.091)	(0.100)	(0.140)	(0.143)	(0.116)	(0.137)	(0.168)	(0.232)
	[0.000]	[0.041]	[0.151]	[0.129]	[0.000]	[0.260]	[0.140]	[0.587]
Subjects	50	80	43	37	50	90	45	45
γ	0.712	0.850	0.895	0.761				
	(0.092)	(0.076)	(0.129)	(0.082)				
	[0.000]	[0.000]	[0.000]	[0.000]				
Subjects	50	80	43	37				
δ		0.635	0.595	0.664		0.300	0.374	0.229
		(0.119)	(0.165)	(0.205)		(0.110)	(0.146)	(0.147)
		[0.000]	[0.001]	[0.003]		[0.007]	[0.014]	[0.127]
Subjects		80	43	37		90	45	45

Note — The dependent variable is b_t. Standard errors in parentheses, p-values in brackets. The parameters are estimated using the adaptive and the rational expectations models described in the text. All models employ a two stages estimation. First, we estimate separately the model for each subject using a GMM method with one or two lagged values of endogenous variables as instruments and robust standard errors. Then, we aggregate the estimates for the parameters using a weighted average in which weights are the inverse of standard errors. This second stage is performed using the Stata command `metareg` which performs a random-effects meta-regression on the results obtained in the first stage.

Table 10: Meta-regression of random coefficient estimation results

	β			δ		
	est.	s.e.	p	est.	s.e.	p
Parameter	0.563	(0.096)	[0.000]	0.597	(0.145)	[0.000]
Rational expectations	0.084	(0.119)	[0.478]	-0.292	(0.184)	[0.114]
Extreme low info group	-0.374	(0.130)	[0.004]	0.056	(0.174)	[0.750]
Extreme high info group	-0.413	(0.139)	[0.003]			

Note — The dependent variables are the parameters β and δ. Standard errors in parentheses, p-values in brackets. The estimation of β and δ is performed over the various samples using the Stata command `metareg` which implements a random-effects meta-regression.

References

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole américaine. *Econometrica*, 21(4):503–546.

Anderson, T. W. and Hsiao, C. (1981). Estimation of Dynamic Models with Error Components. *Journal of the American Statistical Association*, 76(375):598—606.

Arellano, M. and Bond, S. (1991). Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations. *The Review of Economic Studies*, 58(2):277—297.

Arellano, M. and Bover, O. (1995). Another Look at the Instrumental Variable Estimation of Error-Components Models. *Journal of Econometrics*, 68(1):29—51.

Ariely, D., Loewenstein, G., and Prelec, D. (2003). "Coherent Arbitrariness": Stable Demand Curves Without Stable Preferences. *The Quarterly Journal of Economics*, 118(1):73–105.

Ariely, D., Loewenstein, G., and Prelec, D. (2006). Tom Sawyer and the construction of value. *Journal of Economic Behavior & Organization*, 60(1):1–10.

Bernheim, B. D. (1994). A theory of conformity. *Journal of Political Economy*, 102(5):841–877.

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026.

Binmore, K. (1999). Why experiment in economics? *The Economic Journal*, 109(453):16–24.

Blundell, R. and Bond, S. (1998). Initial Conditions and Moment Restrictions in Dynamic Panel Data Models. *Journal of Econometrics*, 87(1):115—143.

Butler, D. J. and Loomes, G. C. (2007). Imprecision as an account of the preference reversal phenomenon. *American Economic Review*, 97(1):277–297.

Cox, J. C. and Grether, D. M. (1996). The preference reversal phenomenon: response mode, markets and incentives. *Economic Theory*, 7(3):381–405.

Filoso, V. (2013). Regression Anatomy, Revealed. *The Stata Journal*, 13(1):92—106.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.

Gelman, A. and Hill, J. (2007). *Data Analysis Using Regression And Multilevel/Hierarchical Models*. Cambridge University Press, New York, NY.

Harbord, R. M. and Higgins, J. P. (2008). Meta-regression in Stata. *The Stata Journal*, 8(4):493—519.

Holtz-Eakin, D., Newey, W., and Rosen, H. S. (1988). Estimating Vector Autoregressions With Panel Data. *Econometrica*, 56(6):1371—95.

Hsiao, C. and Pesaran, H. (2008). Random coefficient models. In Mátyás, L. and Sevestre, P., editors, *The Econometrics of Panel Data. Fundamentals and Recent Developments in Theory and Practice*, chapter 9, page 185—213. Springer, Berlin, 3rd edition.

Isoni, A. (2011). The willingness-to-accept/willingness-to-pay disparity in repeated markets: Loss aversion or 'bad-deal'aversion? *Theory and Decision*, 71(3):409–430.

Isoni, A., Brooks, P., Loomes, G., and Sugden, R. (2011). Do markets reveal preferences-or shape them? Technical report, School of Economics, University of East Anglia, Norwich, UK.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291.

Knapp, G. and Hartung, J. (2003). Improved tests for a random effects meta-regression with a single covariate. *Statistics in Medicine*, 22(17):2693—2710.

Knetsch, J. L. and Sinden, J. A. (1984). Willingness to pay and compensation demanded: Experimental evidence of an unexpected disparity in measures of value. *The Quarterly Journal of Economics*, 99(3):507–521.

Kőszegi, B. and Rabin, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165.

Leamer, E. E. (1983). Let's Take the Con Out of Econometrics. *The American Economic Review*, 73(1):31—43.

Lichtenstein, S. and Slovic, P. (2006). *The construction of preference*. Cambridge University Press, Cambridge, MA.

Loomes, G. (2005). Modelling the stochastic component of behaviour in experiments: Some issues for the interpretation of data. *Experimental Economics*, 8(4):301–323.

Loomes, G., Starmer, C., and Sugden, R. (2003). Do Anomalies Disappear in Repeated Markets? *The Economic Journal*, 113(486):C153—C166.

Loomes, G., Starmer, C., and Sugden, R. (2010). Preference reversals and disparities between willingness to pay and willingness to accept in repeated markets. *Journal of Economic Psychology*, 31(3):374–387.

Lusk, J. L. and Shogren, J. F. (2007). *Experimental auctions: Methods and applications in economic and marketing research*. Cambridge University Press, Cambridge, MA.

Morgan, T. J. H. and Laland, K. N. (2012). The biological bases of conformity. *Frontiers in Neuroscience*, 6:87.

Munro, A. and Sugden, R. (2003). On the theory of reference-dependent preferences. *Journal of Economic Behavior & Organization*, 50(4):407–428.

Muth, J. F. (1961). Rational Expectations and the Theory of Price Movements. *Econometrica*, 29(3):315—335.

Nerlove, M., Sevestre, P., and Balestra, P. (2008). Introduction. In Mátyás, L. and Sevestre, P., editors, *The Econometrics of Panel Data. Fundamentals and Recent Developments in Theory and Practice*, chapter 9, page 185—213. Springer, Berlin, 3rd edition.

Niedrich, R. W., Sharma, S., and Wedell, D. H. (2001). Reference price and price perceptions: A comparison of alternative models. *Journal of Consumer Research*, 28(3):339–354.

Oehlert, G. W. (1992). A Note on the Delta Method. *The American Statistician*, 46(1):27—29.

Plott, C. R. (1996). Rational individual behavior in markets and social choice processes: the discovered preference hypothesis. In Arrow, K., Colombatto, E., Pearlman, M., and Schmidt, C., editors, *The Rational Foundations of Economic Behaviour*, pages 225–250. McMillan, London, UK.

Roodman, D. (2009). How to do `xtabond2`: An Introduction to Difference and System GMM in Stata. *The Stata Journal*, 9(1):86.

Samuelson, L. (2004). Information-based relative consumption effects. *Econometrica*, 72(1):93–118.

Shogren, J. F., Shin, S. Y., Hayes, D. J., and Kliebenstein, J. B. (1994). Resolving differences in willingness to pay and willingness to accept. *The American Economic Review*, pages 255–270.

Stanley, T. D. and Jarrell, S. B. (1989). Meta-Regression Analysis: A Quantitative Method of Literature Surveys. *Journal of Economic Surveys*, 3(2):161—170.

Thaler, R. (1985). Mental accounting and consumer choice. *Marketing science*, 4(3):199–214.

Tufano, F. (2010). Are 'true'preferences revealed in repeated markets? an experimental demonstration of context-dependent valuations. *Experimental Economics*, 13(1):1–13.

Tversky, A. and Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *The Quarterly Journal of Economics*, pages 1039–1061.

van de Waal, E., Borgeaud, C., Whiten, A., Inomata, T., Triadan, D., Aoyama, K., Castillo, V., Yonenobu, H., Chen, Y., and Dorn, G. W. (2013). Potent social learning and conformity shape a wild primate's foraging decisions. *Science*, 340(6131):483–485.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37.

Windmeijer, F. (2005). A Finite Sample Correction for the Variance of Linear Efficient Two-Step GMM Estimators. *Journal of Econometrics*, 126(1):25—51.