Rebuilding Macroeconomics Working Paper Series

Cognitive Variability Matches Speculative Price Dynamics

JIAN-QIAO ZHU*

*Department of Psychology
University of Warwick*

JAKE SPICER

*Department of Psychology
University of Warwick*

ADAM N. SANBORN

*Department of Psychology
University of Warwick*

NICK CHATER

*Warwick Business School
University of Warwick*

Working Paper No.43

<https://www.rebuildingmacroeconomics.ac.uk/publications>

National Institute of Economic and Social Research
2 Dean Trench Street, Westminster
London, SW1P 3HE

This Working Paper is issued under the Macroeconomic Instability Hub of the ESRC's Rebuilding Macroeconomics Network. One or more of the authors of the Working Paper has received support from Rebuilding Macroeconomics or been an active participant in its research hubs and events.

Rebuilding Macroeconomics is funded by the UK's Economic and Social Research Council (ESRC). Our aim is to transform macroeconomics back into a policy relevant social science. We bring together scholars of from many disciplines, policy makers, representatives of civil society and interested members of the public. We are particularly interested in interdisciplinary research and bringing new methods into macroeconomic analysis.

Working Papers may include preliminary or incomplete work and they are circulated to encourage discussion and comment. Citations and use of such papers should take this position into account.

*Corresponding author: Jianqiao.Zhu@warwick.ac.uk

Summary Paragraph

Price series in speculative markets show a common set of statistical properties, termed ‘stylised facts’. While some facts support simple efficient markets composed of homogenous rational agents (e.g., the absence of autocorrelation in price increments), others do not (e.g., heavy-tailed distributions of price changes and volatility clustering) (Campbell et al., 1997; Fama, 1970; Mandelbrot, 1966; Mandelbrot, 1963; Cont, 2001). Collectively, these facts have been explained by either more complex markets or markets of heterogeneous agents (Cont 2007; Giardina & Bouchaud, 2003; Hommes, 2006; Barberis & Thaler, 2005), with asset-market experiments validating the latter approach (Hommes 2011; Kirchler & Huber, 2009). However, it is unknown whether markets are necessary to produce these features. Here we show that within-individual variability alone is sufficient to produce many of the stylised facts. In a series of experiments, we increasingly simplified a price prediction task by first removing external information, then removing any interaction between participants. Finally, we removed any resemblance to an asset market by asking participants to simply reproduce temporal intervals. All three experiments produced the main stylised facts. The robustness of the results across tasks suggests a common cognitive-level mechanism underlies these patterns, and we identify a candidate that is a general-purpose approximation to rational behavior. We recommend a stronger focus on individual psychology in macroeconomic theory, and particularly within-individual variability. Combining these insights with existing economic mechanisms could help explain price changes in speculative markets.

Keywords: cognition, finance, behavioural science, stylised facts

Cognitive Variability Matches Speculative Price Dynamics

What causes the price of speculative assets (e.g., stocks, commodities, or currencies) to change from time to time? A host of approaches and assumptions have been proposed to help address this fundamental problem: for example, when human behaviours are assumed to be rational and risk neutral, price series should follow a random walk in which all price changes represent random departures from previous prices (Samuelson, 1965; Mandelbrot, 1966). The idea of a random-walk model of price is closely related to the efficient market hypothesis (Fama, 1970), which supposes markets reflect all known information efficiently; that is, whenever any news, which is naturally unpredictable, is made available, it will be incorporated into the price without delay. Thus, a price series as predicted by the efficient market hypothesis must be a random walk in which price changes are independent of one another. And indeed, one of the most important stylised empirical facts of financial markets is the absence of autocorrelation in returns (i.e., changes in the log price series) for lags greater than a few minutes; this is also demonstrated in the power spectrum, which resembles $1/f^2$ noise, the spectral character of random walks (see Figure 2 S&P500 panels; Cont, 2001).

Further assuming that the log of a price series follows a Gaussian random walk (i.e., the price follows the geometric Brownian motion), the famous Black-Scholes-Merton (BSM) model offers a tractable solution for the dynamics of a financial market (Black & Scholes, 1973; Merton, 1973). However, Gaussian random walks do not describe real markets. First, Gaussian distributions of price changes underestimate extreme movements: in a hypothetical Gaussian market, a loss of greater than 4 standard deviations is only expected to occur once every 126 years (Frain, 2009), but the FTSE-100 index has shown losses of such magnitude 11 times between 1987 and 2008, even if excluding the 2008 global financial crisis. The distribution of returns is thus better described as having heavy tails (see Figure 2 S&P500

panels; Mandelbrot, 1963), which are more pronounced for short lags (e.g., intra-day or 30 minutes).

Besides heavy tails, there exists another critical feature of the financial markets which is not captured by the Gaussian random-walk model: long range dependencies in price series. While the absence of autocorrelation in price changes supports a random-walk model of prices, it does not necessarily imply that price changes are independent. Indeed, another stylised fact, volatility clustering, demonstrates that price changes are dependent: absolute price changes exhibit persistently positive autocorrelations (Mandelbrot, 1963; Campbell, Lo, & McKinlay, 1999). Thus, while the directions of price changes are unpredictable, their magnitudes are not: large price changes are more likely to be followed by large price changes and vice versa.

Two threads of economic models have been developed to accommodate these stylised facts and make useful predictions for asset prices: Dynamic Stochastic General Equilibrium (DSGE) and heterogeneous agent-based models. The former approach typically assumes a homogenous representative agent that is rational in forecasting future economic variables and optimizes economic decisions given that forecast (Kydland & Prescott, 1982). The latter approach instead assumes heterogeneous agents which are fundamentally limited (Farmer & Foley, 2009). Traditionally, DSGE models assume homoscedasticity; hence they are unable to reproduce stylised facts such as volatility clustering, though recent efforts have been made to alter the models to produce more stylised facts (Fernández-Villaverde & Rubio-Ramírez, 2010; Uhlig, 2007). In contrast, the heterogeneous agent-based models more easily produce the stylised facts, and indeed, this has been accomplished in numerous ways. For example, market participants could be assumed to attend to information at various time scales; long-term investors tend to focus on long-term prices behaviours, whereas short-term investors tend to exploit short-term fluctuations (Guillaume et al., 1997). Prices thus are an aggregation

of different information flows due to the diversity in time horizons; this mechanism has been shown to produce persistent volatility clustering in both simulations (Andersen & Bollerslev, 1997; Bouchaud & Potters, 2003) and asset-market experiments (LeBaron, 2001; Hommes 2011; Kirchler & Huber, 2009). Furthermore, models with heterogeneity in agent-level behaviours (e.g., switching between trading strategies) have also been used to produce the stylised facts (Cont, 2007; Kirman, 1993).

These two approaches do however share a common focus on market level explanations, either through market mechanisms or agent interactions. Here we explore the necessity of assuming markets or interacting agents by examining the variability of cognitive fluctuations and the extent to which they match the price dynamics on aggregate. In the spirit of exploring what conditions are sufficient to produce the stylised facts in the laboratory, our first experiment simplified the design of conventional asset-market experiments: participants only need to predict the next price, which will be the average prediction of all group members. In contrast to conventional asset-market experiments, the buying-selling procedure and information about the fundamental value of the asset are removed, making this a pure belief-based pricing task. There was also no heterogeneity of information, as we presented the same price information to all participants simultaneously. In this simplified asset-market experiment, we were able to reproduce the absence of autocorrelation in price changes, heavy tails, and volatility clustering (see Figure 2). This result suggests that a very basic experimental setup is sufficient to generate the stylised facts of asset prices.

But is interaction among agents really necessary for the emergence of stylised facts, or could individual cognition play a role? Experiment 1, in common with other asset market experiments, does not separate these factors; that is, the complexity of existing asset-market experiments prevents a clean account regarding *which* aspect of the task or model is responsible for generating the stylised facts and whether all the ingredients are needed (c.f.

LeBaron, 2000). Manipulations of information and group structures certainly influence the resulting price series, but perhaps the individual psychology of decision makers, in particular their within-individual variability, may itself exhibit the stylised facts. Indeed, there are reasons to think this is possible. Almost any human judgment is noisy, including economic decisions between risky options: in a series of repeated gambles, participants will make inconsistent choices (Mosteller & Nogee, 1951; Kahneman, Sibony, & Sunstein, 2021). The noise in human judgment is also not independent from trial to trial. For example, when participants were asked to repeatedly estimate a fixed quantity (e.g., 1-sec or 1-inch), they produced noisy estimates that showed long-range dependence (Gilden et al., 1995), and long-range dependence occurs in the time taken to make a response in a number of tasks (Gilden, 1997).

To test whether internal cognitive fluctuations alone are sufficient to reproduce the stylised facts, we further simplified the design of Experiment 1 by cutting out group interactions. In Experiment 2, participants were asked to individually predict the next price in a series in which price was not influenced by their predictions; instead, the target price series followed a Gaussian random walk in log space, in essence reflecting an idealized rational price sequence without heavy tails or volatility clustering. The reward-maximising behavior in this task, if the nature of underlying price series is known or has been learnt, is to simply use the preceding price as prediction for the next. However, participants in Experiment 2 deviated from optimal behavior. Moreover, their predictions showed the stylised facts of markets despite the absence of volatility clustering and heavy tails in the target price series (see Figure 2), though the magnitude of the volatility clustering was less than that observed in Experiment 1. This suggests that a common cognitive mechanism is generating the stylised facts.

The framing of the task in terms of predicting asset prices, however, could have misled participants into believing that the target price series had the same properties as real-world price series. Hence, despite presenting a Gaussian random walk as a target, participants' greater experience with real-world markets may have overshadowed the “small” number of training trials in the laboratory, resulting in a strategy of reproducing the statistics of real-world markets. So, in Experiment 3, we further abstracted away from real-world markets by removing any reference to price, financial assets or predictions. Participants were instead simply asked to reproduce temporal intervals (see Figure 1). Rather than predicting the next target, participants were asked to reproduce the most recent target as accurately as possible immediately after its presentation, where the true duration varied from trial to trial following the same Gaussian random-walks as in Experiment 2, but with the digital numbers converted to physical times. As such, the optimal behavior was again to directly replicate the target series (and therefore its properties), though in this case the optimal behavior was the explicit goal of the task. Again, participants' time series of estimates did not behave like the target Gaussian random-walk but showed the stylised facts of financial markets (see Figure 2) though again the magnitude of the volatility clustering was less than that observed in Experiment 1. Another difference was that the high-frequency regions of the power spectrum curved away from the straight $1/f^2$ noise line. This deviation is possibly due to motor noise dominating cognitive noise in short lags; the same effect has also been previously documented in the repeated estimation task of a fixed time interval (Gilden et al., 1995).

Why might these stylised facts appear in the two very different tasks of forecasting and time estimation? One possible explanation is that the responses of any complex system always exhibit these properties. For example, it has been argued that the complex process of air turbulence shows the same heavy tails as price series (Ghashghaie et al., 1996). However, the turbulence data was later criticized for being unable to match the dependence statistics

(e.g., absence of autocorrelation) of price series (Mantegna & Stanley, 1996). Furthermore, the near-ubiquitous dependencies in psychological time series can be made to disappear when people are instructed to simply respond as quickly as possible to an unpredictable target, suggesting that cognitive processes are necessary to produce the stylised facts (Gilden, et al., 1995). In sum, cognition may be the necessary ingredient to produce the empirical observations in our three experiments.

If it is cognition that generates these stylised facts, how are agents able to produce both features that align with rational agents (absence of autocorrelation) and those that do not (heavy tails and volatility clustering)? Instead of investors being entirely rational, it could be that human beings approximate rational solutions with limited cognitive resources (Lieder & Griffiths, 2020). One scheme for implementing resource-rationality is to assume that people are sampling forecasts or time estimates according to their probability rather than always producing the single best response, and use the sampling algorithms that computer scientists and statisticians use (Sanborn & Chater, 2016; Griffiths, Vul, & Sanborn, 2012). Many of these algorithms produce dependent samples, and, along with assuming that people only generate a small number of samples due to the mental cost involved, this approach successfully explains a number of well-known deviations of human behavior from rationality (Zhu, Sanborn, & Chater, 2020; Vul et al., 2014). We tested a sophisticated sampling algorithm as the resource-rational approximation and found it replicated the stylised facts well (see Supplementary). And indeed, assuming that people make stochastic approximations could potentially explain other puzzles of the market that our experiments were not designed to address, such as why markets show more volatility than they should (Shiller, 1981, Le Roy & Porter, 1981): market volatility could be inflated by cognitive volatility.

One potential criticism of this work, often used in past rebuttals of the importance of human deviations from rationality, is that the presence of rational arbitrageurs or the

aggregation of many individuals in a market will reduce the influence of within-individual variability on asset prices to zero. There are however reasons to believe that individual cognition could have an influence. First, unsophisticated noise traders who are subject to these cognitive fluctuations can determine a price series even if rational arbitrageurs are present (De Long et al., 1990). Second, the impact of individual cognition can be amplified by power-law regularities in economics: while market aggregation using equally-weighted averaging may wash out the distinctive properties of individual cognition (due to the central limit theorem), power-law averaging does not (see Supplementary for details). Power-law distributions are ubiquitous in economics and finance, demonstrated in personal wealth, firm sizes, and trading volumes in stock markets (Gabaix, 2009). It is then possible that economic activity could be concentrated around fluctuations in the beliefs of a few individuals/entities, and so affect the wider market.

We have here provided a complementary, cognitive perspective to the properties of fluctuations in asset markets. Cognitive fluctuations are clearly not the only cause of the stylised facts as shown by quantitative mismatches between psychological and financial time series in the magnitude of volatility clustering. Even so, our findings suggest closer attention be paid to the role of cognitive fluctuations in macroeconomic activity and policy intervention. The spontaneous emergence of heavy tails and volatility clustering in cognition is particularly worrisome given its potential to cause financial disasters. However, it does suggest investigating whether interventions that improve the independence or accuracy of repeated estimates, such as slowing down decision making or considering problems from multiple perspectives (Herzog & Hertwig, 2014), could help alleviate these issues, though first investigations are needed on how these interventions influence the stylised facts. We see this endeavor as an interdisciplinary one which will benefit from dialogues between economics, psychology, statistics, computer science and beyond.

Figure 1. Schematic illustrations of a typical trial in group prediction (a), individual prediction (b), and time estimation tasks (c).

Figure 2. Comparing statistical behaviours of financial (daily close of S&P500 for the 32-year period Jan 1986 to Jun 2018) and psychological time series (group price predictions, individual price predictions, and individual time estimations). **(a)** Power spectral density (95% CIs shaded around the solid blue line as mean). Dashed lines denote Brown noise (or $1/f^2$ noise). **(b)** Empirical probability density function in quantile-quantile plot (95% CIs shaded in empirical quantiles). Dashed horizontal lines denote normal distributions. **(c)** Autocorrelation function of logarithm of successive changes (95% CIs shaded around correlational coefficients at each lag). Dashed horizontal lines denote no autocorrelation. **(d)** Autocorrelation function of volatility or

absolute log changes (95% CIs shaded around correlational coefficients at each lag). Dashed horizontal lines denote no autocorrelation.

References

- Andersen, T. G., & Bollerslev, T. (1997). Heterogeneous information arrivals and return volatility dynamics: Uncovering the long-run in high frequency returns. *The Journal of Finance*, 52(3), 975-1005.
- Barberis, N., & Thaler, R. (2005). *A survey of behavioral finance*. Princeton University Press.
- Black, F., & Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*. 81 (3): 637–654.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307-327.
- Bouchaud, J. P., & Potters, M. (2003). *Theory of financial risk and derivative pricing: from statistical physics to risk management*. Cambridge university press.
- Campbell, J. Y., Lo, A. W., & MacKinlay, A. C. (1997). *The econometrics of financial markets*. Princeton University press.
- Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9, 88–97.
- Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2), 223-236.
- Cont, R. (2007). Volatility clustering in financial markets: empirical facts and agent-based models. In *Long memory in economics* (pp. 289-309). Springer, Berlin, Heidelberg.
- De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of Political Economy*, 98(4), 703-738.
- Fama, E. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *Journal of Finance*, 25(2), 383–417.

- Farmer, J. D., & Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460(7256), 685-686.
- Fernández-Villaverde, J., & Rubio-Ramírez, J. (2010). *Macroeconomics and volatility: Data, models, and estimation* (16618). National Bureau of Economic Research.
- Frain, J. C. (2009). *Studies on the Application of the Alpha-stable Distribution in Economics*. Trinity College.
- Gabaix, X. (2009). Power laws in economics and finance. *Annual Review Economics*, 1(1), 255-294.
- Ghashghaie, S., Breymann, W., Peinke, J., Talkner, P., & Dodge, Y. (1996). Turbulent cascades in foreign exchange markets. *Nature*, 381(6585), 767-770.
- Giardina, I., & Bouchaud, J. P. (2003). Bubbles, crashes and intermittency in agent-based market models. *The European Physical Journal B-Condensed Matter and Complex Systems*, 31(3), 421-437.
- Gilden, D. L. (2001). Cognitive emissions of $1/f$ noise. *Psychological Review*, 108(1), 33.
- Gilden, D. L., Thornton, T., & Mallon, M. W. (1995). $1/f$ noise in human cognition. *Science*, 267(5205), 1837-1839.
- Guillaume, D. M., Dacorogna, M. M., Davé, R. R., Müller, U. A., Olsen, R. B., & Pictet, O. V. (1997). From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets. *Finance and Stochastics*, 1(2), 95-129.
- Heemeijer, P., Hommes, C., Sonnemans, J., & Tuinstra, J. (2009). Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economic Dynamics and Control*, 33(5), 1052–1072.
- Herzog, S. M., & Hertwig, R. (2014). Harnessing the wisdom of the inner crowd. *Trends in Cognitive Sciences*, 18(10), 504-506.

- Hommes, C. (2011). The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics and Control*, 35(1), 1-24.
- Hommes, C. H. (2006). Heterogeneous agent models in economics and finance. *Handbook of Computational Economics*, 2, 1109-1186.
- Hommes, C. H., Sonnemans, J., Tuinstra, J., & Van de Velden, H. (2008). Expectations and bubbles in asset pricing experiments. *Journal of Economic Behavior & Organization*, 67(1), 116–133.
- Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A flaw in human judgment*. William Collins, London.
- Kirchler, M., & Huber, J. (2009). An exploration of commonly observed stylized facts with data from experimental asset markets. *Physica A: Statistical Mechanics and its Applications*, 388(8), 1631-1658.
- Kirman, A. (1993). Ants, rationality, and recruitment. *The Quarterly Journal of Economics*, 108(1), 137-156.
- Kydland, F. E., & Prescott, E. C. (1982). Time to build and aggregate fluctuations. *Econometrica: Journal of the Econometric Society*, 1345-1370.
- LeBaron B. (2001). Evolution and time horizons in an agent-based stock market, *Macroeconomic Dynamics*, 5, pp. 225–254.
- LeBaron, B. (2000). Agent-based computational finance: Suggested readings and early research. *Journal of Economic Dynamics and Control*, 24(5-7), 679-702.
- Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 43.
- Mandelbrot, B. B. (1963). The variation of certain speculative prices. In *Fractals and scaling in finance* (pp. 371-418). Springer, New York, NY.

- Mandelbrot, B. B. (1966). Forecasts of future prices, unbiased markets, and "martingale" models. *The Journal of Business*, 39(1), 242-255.
- Mantegna, R. N., & Stanley, H. E. (1996). Turbulence and financial markets. *Nature*, 383(6601), 587-588.
- Merton, Robert C. (1973). Theory of Rational Option Pricing. *Bell Journal of Economics and Management Science*. The RAND Corporation. 4 (1): 141–183.
- Mosteller, F., & Nogee, P. (1951). An experimental measurement of utility. *Journal of Political Economy*, 59(5), 371-404.
- Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., ... & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, 51(1), 195-203.
- Robert J. Shiller. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71:421—436, 1981
- Stephen F. LeRoy and Richard D. Porter. The present value relation: Tests based on implied variance bounds. *Econometrica*, 49:555—574, 1981.
- Uhlig, H. (2007). Explaining asset prices with external habits and wage rigidities in a DSGE model. *American Economic Review*, 97(2), 239-243.
- Zhu, J.-Q., Sanborn, A. N., & Chater, N. (2020). The Bayesian sampler: Generic Bayesian inference causes incoherence in human probability judgments. *Psychological Review*.

Methods

Experiment 1: Group Prediction

Participants

One-hundred-and-fifty participants were recruited through Prolific in return for financial compensation between £6 and £7 based on accuracy of prediction in a randomly selected trial. 120 participants qualified for further analysis (44 females, 73 males, 3 declined to answer, aged between 18 and 80; exclusion criteria explained below).

Procedure

Experiment 1 examined estimates in small experimental markets: groups of participants each made separate predictions of the price of a risky fictional asset based on a common history, with the actual price of that asset then being based on aggregated predictions from all participants in each period. This was based on the ‘learning-to-forecast’ experiments of Hommes and colleagues, most notably Hommes et al. (2008) and Heemeijer et al. (2009), in which participants were asked to predict the next price of a fictional asset based on the price history up to that point. Unlike these experiments, however, we do not define a dividend for the target asset, nor do we note any safe alternative asset with a fixed interest rate, meaning the asset has no defined fundamental price; this was to focus attention on price movements arising from cognitive fluctuations in expectations without any external information. Price in this experiment is therefore purely a function of participant predictions rather than any pre-determined series.

Experiment 1 ran in a series of 1-hour sessions performed online using the software oTree (Chen, Schonger, & Wickens, 2016). Upon entering the experiment, participants were first sorted into groups of 6 by order of arrival, which remained consistent across the session. If a group did not reach 6 members within 5 minutes of the start of the session, the task

proceeded with that incomplete group, with price being calculated according to the average of the present members; incomplete groups were then removed from subsequent analysis.

Participants were each framed as an advisor to a large financial firm deciding whether to invest in a particular asset, making repeated individual predictions of the price of the asset in the next immediate trial, mirroring the one-period ahead predictions of Heemeijer et al. (2009). Firms were stated to base their demand for the asset on the predictions of their advisor: high price predictions lead to greater demand, and vice versa. Price would then be set according to the total demand of all firms within the market, thus reflecting the predictions of all participants in the group for that period. Participants were however informed that they would be rewarded purely based on the accuracy of their predictions, and so should focus only on being as close to the true price as possible. Bonuses were noted to be determined by accuracy on a randomly selected trial to incentivize accuracy on all trials.

After reading the instructions, participants began the main trial block. Each trial requested participants predict the price of the asset for that trial (i.e. today's as-yet undetermined price) by typing their estimate into an on-screen text box. Predictions were required to be positive integers with an upper limit of 10000 to prevent extreme increases in price, as used in Hommes et al. (2008); participants were not made aware of this limit at the start of the task, but would be asked to enter a new estimate if a submitted prediction exceeded this value. In addition, to counter potential typing errors, predictions with an absolute log difference of 0.3 from the most recent price required secondary confirmation by the participant to be accepted; this threshold was determined via pilot testing.

To assist predictions, each trial showed a line plot illustrating the history of both the actual price and that participant's predictions for all past trials. For the first trial in which no previous information was available, participants were instead given a fixed starting value of 200 representing the most recent price of the asset, common to all experimental sessions.

A time limit of 10 seconds was placed on each prediction to assure task progression; if no prediction was entered within this time, the task automatically advanced and that participant had no influence on price for that trial.

Price was then calculated by taking the average prediction across participants who responded for that trial, simulating a basic equilibrium between supply and demand, with the addition of Gaussian noise:

$$p_t = \frac{\sum_{i=1}^{n_t} p_{i,t}^e}{n_t} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0,2)$$

where $p_{i,t}^e$ is the prediction of participant i on trial t , and n_t is the number of participants who made a prediction on that trial. Prices were rounded to the nearest integer to match with responses. This price was then revealed to each group member alongside the potential reward for their prediction if that trial were selected; this feedback remained on-screen for 3 seconds before advancing to the next trial.

The task continued until either a maximum of 300 trials or a total duration of 60 minutes. The experiment then ended by calculating a bonus payment for each participant according to their prediction accuracy on a randomly sampled trial. Reward was calculated using an exponential function of absolute log deviation from true price on the sampled trial restricted to fall between £0 and £1: $R_{i,t} = \exp(-20 \times |\ln p_t - \ln p_{i,t}^e|)$. Participants therefore maximise reward by coordinating their predictions with each other, thus maximising accuracy. This also means that no single prediction strategy necessarily outperforms any other so long as all participants follow a common strategy.

Experiment 2: Individual Prediction

Participants

Seventy-four participants were recruited through the University of Warwick SONA subject pool, and completed the 60-min experiment in exchange for monetary rewards between £2

and £20 based on accuracy of prediction in a randomly selected trial. Participants were excluded from analysis if they reported an interruption during the task, or if less than 350 predictions were completed within the 60-minute session; this removed 33 participants, leaving 41 participants for further analysis (15 females, 26 males, and aged between 19 and 32).

Procedure

Experiment 2 was conducted remotely using the PsychoPy software packages (Peirce et al., 2019), installed on participants' personal computers. Participants were first assigned to one of two potential target time series, with 20 participants viewing series A, and 21 viewing series B. Both target time series were defined using random walk processes with Gaussian increments in log scale:

$$\ln x_t = \ln x_{t-1} + \epsilon_t, \epsilon_t \sim \mathcal{N}(0, \sigma^2)$$

The two series differed in standard deviation of step size ($\sigma_A = 0.25, \sigma_B = 0.2$) and the random seed used to generate each step. Participants viewed the target in absolute scale, which was multiplied by 100 and rounded to the nearest integer for display.

Participants were instructed that they would be predicting the price of a financial asset that moved up and down randomly, and would be rewarded based on their accuracy of predicting the target in the next trial. Each trial presented the current target to the participant on-screen for 2 seconds. After presentation, participants were immediately asked to predict the next target by typing a positive integer. Once the prediction was made, a feedback screen was shown noting the next target and the potential monetary reward that the participant would receive if that trial were selected for payment. Trials were separated by an inter-trial interval of 1 second while displaying a cross fixation at the center of the screen. Trials continued until a maximum duration of 60 minutes.

At the end of the experiment, participants were also asked to report whether they had been interrupted during the task, with any participants indicating that they had being removed from subsequent analysis. Payments were then determined using a nonlinear transform of the participant's log error in their prediction on a randomly selected trial: $R_t = 2 + 18 \times \exp(-9 \times |\ln p_t^e - \ln x_{t+1}|)$.

Experiment 3: Time Estimation

Participants

Fifty-eight participants were recruited through the University of Warwick SONA participant pool in exchange for monetary rewards between £2 and £20 based on accuracy of estimate in a randomly selected trial. Exclusion criteria were identical to Experiment 2, removing 17 participants and leaving 41 for subsequent analysis (22 females, 18 males, 1 non-disclosed gender, and aged between 19 and 45).

Procedure

Experiment 3 was again conducted remotely using the PsychoPy software packages (Peirce et al., 2019), installed on participants' personal computers. Participants were first assigned to one of two potential target time series, with 20 participants viewing series A, and 21 viewing series B. These were the exact same series as in Experiment 2, though the log changes were scaled down by a factor of 4 from the ones used in Experiment 2 so that the time intervals were not too long or too short.

Participants were told that the task involved viewing and replicating a time interval, and they would be rewarded based on their accuracy. Each trial presented the target time interval to the participant as a red circle on-screen for the given duration; that is, the red circle appeared on-screen for x_t seconds in the t -th trial. After presentation, participants were immediately asked to replicate the target duration by holding the space-bar for the same

length of time as the red circle had appeared. To aid response, a matching red circle was displayed as participants held the space-bar. Once the estimate was made, a feedback screen was shown noting the potential bonus that the participant would receive if that trial were selected for payment. Trials were separated by an inter-trial interval of 1 second while a cross fixation was displayed at the center of the screen. Trials continued until a maximum duration of 60 minutes. Four practice trials were given to all participants prior to the main task, following a sequence of target time intervals in the order of 1sec, 2sec, 1sec, and 2sec. Finally, participants were also asked to report whether they had been interrupted during the task, with any participants indicating that they had being removed from subsequent analysis. Payments were then determined by applying the same reward function as in Experiment 2 on a randomly selected trial.

Supplementary Materials

Additional Stylised Facts

Here we explore other stylised facts of price series that are documented in the extensive survey of financial markets conducted by Cont (2001). These are sometimes derivatives of the three stylised facts from the main text or do not occur in all types of financial markets. There are also some features which make reference to specific market aspects which cannot be explored in the current data set: for example, the correlation between trading volumes and volatility requires measures of volumes.

Gain/loss asymmetry

In stock markets, people typically observe greater magnitudes in losses than in gains; indeed, Figure 1b shows the S&P500 index has a heavier left tail than right tail. However, this property is not true for other speculative markets, most notably foreign exchange markets (Cont, 2001). This is likely because an increase in the relative valuation of one currency automatically leads to a decrease in another currency, and vice versa; hence, there is greater symmetry between upward and downward movements for foreign exchange markets. Our group and individual tasks, however, also exhibit more symmetry than the S&P500 index. This might be because the incentives are unbiased.

Aggregational Gaussianity

While return distributions demonstrate heavy tails across shorter time lags, increasing the timescale over which log changes are computed makes the tails of the distribution thinner, approaching a Gaussian distribution. That is, lag-1 changes are expected to have heavier tails than lag-100 changes. Hence, to demonstrate this trend, we calculated the excess kurtosis for changes at lags from 1 to 100 (see Figure S1a), where zero excess kurtosis indicates a Gaussian distribution and larger values indicate heavier tails (we limited our investigation to the range of 1 to 100 lags because some of our psychological time series appear unreliable

after 100 lags because the limited amount of trials tested). It is clear that both the financial and the psychological time series show smaller and smaller excess kurtosis over longer lags.

Conditional heavy tails

Heavy tails are not explained away by statistically correcting for volatility clustering: the residual time series from fitting a GARCH-type model still exhibit heavy tails. GARCH denotes generalized autoregressive conditional heteroskedasticity; it is a statistical model to describe time series that exhibit volatility clustering (Bollerslev, 1986). In particular, the error term of GARCH is assumed to be related to previous error terms (i.e., the error variance follows an autoregressive moving average model). We fitted a GARCH(1,1) model (meaning the lag lengths for the error variance are 1) to all four type of time series and checked the tail properties of the residuals (see Figure S1b). The conditional heavy tails were replicated across both group and individual tasks, though with some notable asymmetry between upward and downward movements.

Slow decay of autocorrelation in absolute returns

The quantitative details of autocorrelation in volatility (or absolute log changes) further suggest that the decay of these autocorrelation functions is slow across time lags. Often the decay can be best described as a power law function. As shown in Figure 1d, psychological time series did not replicate this quantitative property in their autocorrelation function; the autocorrelation of group and individual tasks drops to zero faster than the S&P500 data. This is possibly due to differences in timescale between the financial (ticks of day) and psychological data (ticks of trial): participants in our experiments were free to take as long as they wished to make a response within a trial. As a result, unlike the fixed physical time taken as ticks in analyzing financial data, the physical time taken to complete a trial can vary from person to person.

Figure S1. Comparing financial and psychological time series on additional stylised facts: (a) aggregational Gaussianity and (b) conditional heavy tails. Horizontal dashed lines indicate Gaussian distribution.

Comparing Stylised Facts with Models of Individuals

In this section, we study three models of individuals and compare their statistical properties to the stylised facts. For ease of definition, we focus here on the individual price prediction task of Experiment 2, where price moves according to a Gaussian random walk. First, the optimal model that maximizes the expected payoffs reproduces the present price; the prediction made on trial t for the target that will realize in the next trial should be the target presented on current trial:

$$\ln p_t^e = \ln x_t.$$

Clearly the optimal model will just reproduce the statistics of target time series, which is a Gaussian random walk. As shown in Figure S2a, the statistics of a Gaussian random walk did not reproduce the stylised facts nor the behavioral data in Experiments 2 and 3.

As an alternative, participants may consider only a sample from the distribution of potential movements and make a prediction based on the sample. Sampling-based models of human cognition have provided algorithmic-level accounts that approximate Bayes-optimal solutions (Griffiths, Vul, & Sanborn, 2012; Sanborn & Chater, 2016). They have also successfully explained why behavior is near-optimal on aggregate (Griffiths & Tenenbaum, 2006) yet individual judgments show systematic biases (Dasgupta et al., 2018; Zhu, Sanborn, & Chater, 2020). In our tasks, the correct responses are distributed according to a Gaussian distribution with a fixed standard deviation:

$$\ln p_t^e \sim \mathcal{N}(\ln x_t, \sigma^2).$$

Assuming that people have the correct mental representation to draw samples from (i.e., the Gaussian distribution above) and their predictions only rely on a single sample, we can test whether the dynamics of sampling algorithms match the dynamics of psychological time series documented in our individual tasks. Here, we test two sampling algorithms: a simple direct sampler and a sophisticated MC³ sampler. The direct sampler simply draws an

independent, identically distributed sample from the true Gaussian distribution of price steps. The simulated time series of a direct sampler effectively equates to the target Gaussian random walk with additional Gaussian noise added on each timestep. As a result, the direct sampling model does not predict heavy-tailed changes or volatility clustering (see Figure S2b).

We also considered an alternative sampling model of human cognition: MC³ (algorithmic details can be found below). The MC³ sampler has been previously found to match the statistical properties of long-range dependence and heavy-tailed movements in repeated estimation tasks where targets are typically fixed (Zhu, Sanborn, & Chater, 2018). The stylised facts in the main text were replicated by running the MC³ sampler with a Gaussian random walk as a target (see Figure S2c). The simulated behaviours also match the empirical data in the individual tasks (Experiments 2 and 3). MC³ therefore offers a viable explanation for individual behavior in these tasks, and so is a potential foundation to build more complex models with interacting agents upon.

Algorithmic details of Metropolis-coupled Markov Chain Monte Carlo (MC³)

The MC³ sampler is composed of multiple Markov chains, which are stochastic processes for which future and past states are independent when conditioned on the current state. When the transitions between states are carefully chosen, such as in the Metropolis-Hastings algorithm, the number of iterations the chain spends in a state corresponds to the probability of that state; effectively, simulating the Markov chain acts as if sampling from a probability distribution (Roberts & Casella, 2011). MC³ runs parallel Markov chains each at a different temperature of the target distribution (Geyer, 1991); the more heated chains see more flattened versions of the distribution, and hence can traverse through valleys of probabilistic landscapes (e.g., local minima) more easily. These chains essentially conduct local searches in their heated target distributions and then two randomly selected chains are

chosen to potentially swap positions. The swapping rule is governed by a Metropolis-Hastings acceptance rule with the differences in temperatures between the two chains adjusted appropriately. If there is only one chain, the algorithm reduces to the standard Metropolis-Hastings algorithm.

Simulation details

For each model mentioned above, we simulated 512 (2^9) predictions or time estimations in which the target distribution evolves through time as a Gaussian random walk. This process is repeated 10^4 times with different random seeds. The simulated time series were then analyzed using the same methods in the main text. In simulating the MC³ sampler, we set the number of parallel chains at 8, the proposal distribution as the true Gaussian distribution of price movements as with the direct sampler, and the temperature of the i -th parallel chain as $1/(4i-3)$. The first chain of MC³ is the cold chain which did not distort the target distribution. The location of this cold chain in the mental representation is then taken as the prediction or time estimation.

Figure S2. Simulated model behaviours for agents using (a) the optimal model, (b) direct sampler, and (c) MC³ sampler. The same statistical analyses from the main text were applied to the simulated time series here.

Market aggregation: Power-law averaging

In this section, we examine whether individual-level statistics will be preserved in market aggregation. That is, we assess whether the statistical properties of the group task can be replicated by assuming additional market mechanisms on the individual tasks where participants made estimates independently. While participants in the individual tasks have no chance to communicate with one another, participants who belong to the same condition are predicting future outcomes (Experiment 2) or reproducing time intervals (Experiment 3) with regard to the same underlying target time series. Of course, due to the central limit theorem, simple uniform averaging will wash away non-Gaussian or dependence features that are unique to the individual time series. However, economic and psychological studies have found that the distributions of wealth and income and of social influence have power-law tails (Gabaix, 2009). Overall, this could lead to a power law distribution of individual impacts on the aggregate price: a small number of individuals dominate.

We first randomly sample a subset of 6 participants from a common experimental condition. Their individual time series were aggregated according to a power law weighting with exponent of 3 (Gabaix, 2009; Atkinson & Piketty, 2007) where the weighting order is randomized across participants. The aggregate time series was then analyzed as in the main text. This whole process of subsampling, power law averaging, and analyzing aggregated data is repeated 10^4 times. The results are shown in Figure S3, with all conditions qualitatively reproducing results of the group task except the dependence in volatility of the individual prediction task in Condition A.

Figure S3. Power-law averaging of psychological time series obtained from the individual tasks.

Supplementary References

- Atkinson, A. B., & Piketty, T. (2007). *Top Incomes over the Twentieth Century*. Oxford University Press.
- Dasgupta, I., Schulz, E., Goodman, N. D., & Gershman, S. J. (2018). Remembrance of inferences past: Amortization in human hypothesis generation. *Cognition*, 178, 67-81.
- Geyer, C. (1991). Markov chain Monte Carlo maximum likelihood. In E. M. Keramidas (Ed.), *Proceedings of the 23rd symposium on the interface: Computing science and statistics*. Interface Foundation.
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, 17(9), 767-773.
- Griffiths, T. L., Vul, E., & Sanborn, A. N. (2012). Bridging levels of analysis for probabilistic models of cognition. *Current Directions in Psychological Science*, 21(4), 263-268.
- Robert, C., & Casella, G. (2011). A short history of Markov chain Monte Carlo: Subjective recollections from incomplete data. *Statistical Science*, 102-115.
- Zhu, J.-Q., Sanborn, A. N., & Chater, N. (2018). Mental sampling in Multimodal Representations. In S. Bengio, H. Wallach, H. Larochelle, K. Graumann, N. Cesa-Bianchi, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Montreal, Canada.