

Cognitive Imprecision and Adjustment Dynamics: Some Evidence from the Lab

Luminita Stevens
University of Maryland

August 31, 2021

Overarching Questions

- How can we measure individual-level limitations in perception, information processing, and cognition that are “systematic enough” to affect aggregate outcomes?
- How might complementarities, markets, and policies amplify or dampen the individual-level systematic mistakes made due to these limitations?

Rational Expectations

- The **rational expectations** modeling standard:
 - economic agents are **forward-looking** & **reason** through the consequences of their own actions and those of others
 - they **recall**, **perceive** and **incorporate perfectly** all relevant information
 - they (correctly) believe that **all agents** behave in the same way

Confronting the Evidence

- RE implies **instantaneous complete adjustment** to shocks and regime changes
 - data: **discrete, delayed adjustment** (prices, hiring, investment,...)
 - *response*: keep RE and add adjustment costs, habits
- RE implies **deterministic homogeneous** beliefs and responses
 - data: **huge dispersion** in both beliefs and outcomes
 - data: pervasive **stochasticity of choice**
 - *response*: keep RE, add unbiased idiosyncratic shocks, RUM

Confronting the Evidence

- *Remark 1:* these extensions tend to have limited **implications** for aggregate outcomes and welfare
 - e.g. Golosov and Lucas (2007) in the monetary context
- *Remark 2:* the empirical success has been mixed in identifying **quantitatively meaningful** adjustment costs or randomness in individual preferences

The Cognitive Limitations Hypothesis

- An alternative view: deviations from REE predictions occur because agents economize on the cognitive resources used to
 - acquire the knowledge needed for economic decisions
 - think through the implications of this knowledge
 - predict and incorporate the actions of others
- Extensive survey evidence documenting dynamics and correlates of forecasts and forecast errors
 - heterogeneity, predictability, under/over-reaction
 - strong (but transient) information treatment effects
(e.g., Coibion and Gorodnichenko, 2015; Bordalo, Gennaioli, Ma, Schleifer, 2018; Kohlhas and Walther, 2020; Angeletos, Huo, Sastry, 2020; Coibion, Gorodnichenko, Ropele, 2019;...)

Breaking FIRE

- This evidence has emboldened modelers to unshackle expectations from FIRE's tight grip
- Survey data on expectations is increasingly used directly in the **estimation of DSGE models** (e.g., Aruoba and Schorfheide, 2011; Del Negro and Eusepi, 2011; Carvalho, Eusepi, Moench, Preston, 2021; ...)
- Also increasingly used to guide **model selection** among different non-FIRE or behavioral models (e.g., Coibion and Gorodnichenko, 2015; Angeletos, Huo, Sastry, 2020; ...)

Breaking FIRE

- Importantly, non-FIRE expectations often result in **systematic distortions** at the individual level that are consequential for aggregate outcomes and policy objectives
- *Remark:* standard economic data becomes insufficient for
 - disentangling **preferences and technologies from information and attention**
 - identifying the **mechanisms** behind the deviations from FIRE

Controlled Lab Experiments

- Complementary tool to study
 - **real-time information processing**: the response to news under uncertainty
 - **strategic sophistication**: the extent to which behavior incorporates attempts to predict the forecasts of other subjects in the session
 - **the types of strategies used to problem-solve** : eductive versus inductive reasoning

Why Controlled Lab Experiments?

- Can specify the **objective**
 - participants are rewarded for a specific task
- Can control what **information** is available
 - at least to extent that it is fully processed by subjects
- Can control the true **data-generating process** and what subjects know about it
 - appropriate prior for inference about hidden state

Why Controlled Lab Experiments?

- Can control non-cognitive **adjustment costs** or grounds for **habit formation**
 - relevant for identifying sources of discrete, lumpy adjustment that pervades econ choices
- Can generate **state-dependent stochastic choice data** (Caplin and Dean, 2013)
 - subjects given each of several decision problems multiple times
- Building from huge experimental literature documenting **distorted probabilities** in a wide range of contexts

Illustrations

1. Discrete adjustment
2. Distorted probability estimation
3. Forecasting with strategic complementarities

EXPERIMENT I
APPLICATION TO DISCRETE ADJUSTMENT
Khaw, Stevens & Woodford (2017)

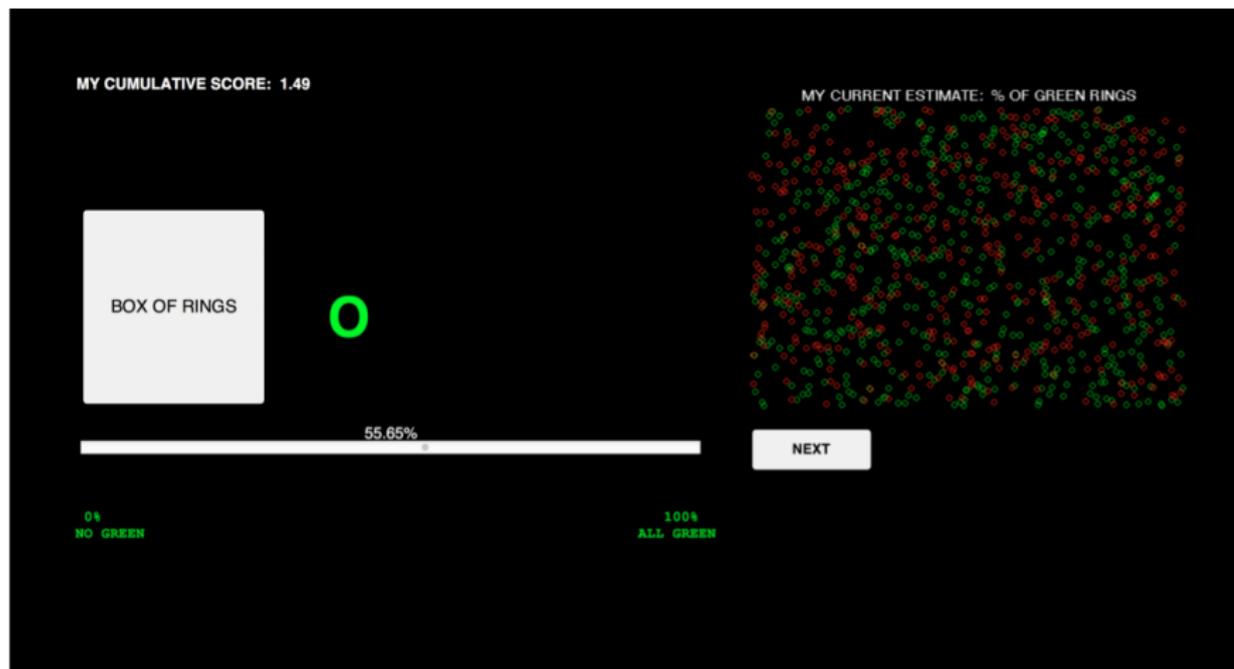
The Experiment

(Khaw, Stevens & Woodford, 2017)

- Variant of an experiment designed by Gallistel et al. (2014)
- Participants estimate the probability of drawing a green ring out of a box with green and red rings
- The Bernoulli parameter governing this probability changes occasionally in an unsignaled manner
- “Tracking” problem in which time-varying true state is continuously distributed
- Participants are told the data generating process of the rings and are rewarded based on their forecasting performance

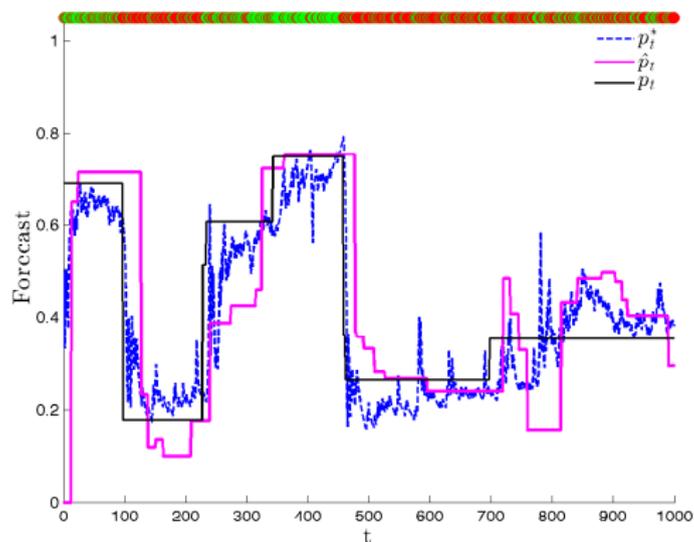
The Experiment

(Khaw, Stevens & Woodford, 2017)



Example of the Data

(Khaw, Stevens & Woodford, 2017)



optimal forecast: Subject 11, session 10
such discrete adjustment is pervasive across subjects

Summary: Application to Discrete Adjustment

(Khaw, Stevens & Woodford, 2017)

- Discreteness of adjustment seems to reflect cognitive constraints rather than any external costs of adjustment
- Evidence points to a model of rationally inattentive adjustment in which subjects choose attention to the task optimally
- Estimated proba of adjustment is nonzero everywhere, monotonic in the gap between value of current position and expected value obtained from adjustment
- Position choice is also inattentive (unlike in standard models of discrete adjustment)

EXPERIMENT I
APPLICATION TO PROBABILITY DISTORTIONS
Khaw, Stevens & Woodford (2020)

Experimental Evidence on Probability Distortions

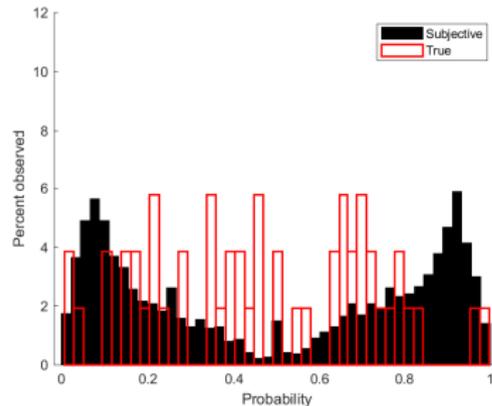
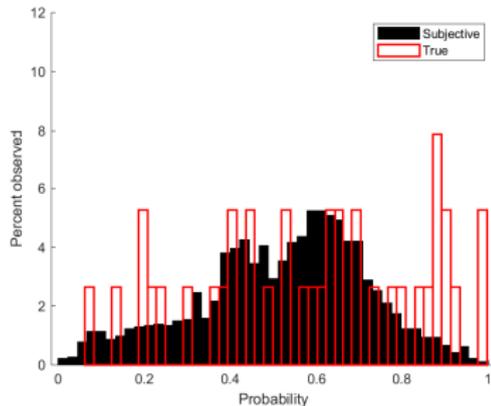
- Many probability estimation tasks generate **conservative estimates** that are biased toward the mean
 - overweight low proba events and underweight high proba events (Tversky & Kahneman, 1992; Camerer & Ho, 1994)
 - also common in perceptual judgements (sound, length, speed)
 - especially if there is **low perceptible difference** between stimuli (Wei & Stocker, 2017)
 - or when the sensory evidence is associated with **high uncertainty**, relative to a priori knowledge (Petzschner, Glasauer, Stephan, 2015)

Experimental Evidence on Probability Distortions

- Other tasks generate **polarized estimates** where extreme values are disproportionately reported
 - 'hot-hand' effect causing over-reaction (Offerman & Sonnemans, 2004)
 - over-weighting recent observations (Murdock Jr, 1962)
 - similar to sequence effects found in perceptual judgments (Cross, 1973)

Biased Forecast Distributions

(Khaw, Stevens & Woodford, 2020)



both conservative and polarized distributions for the same task

Estimating Probability Distortions

We characterize the data using the **linear log odds representation** of probabilities (Zhang & Maloney, 2012):

$$\log \left(\frac{R_{it}}{1 - R_{it}} \right) = \alpha_i + \beta_i \log \left(\frac{B_t}{1 - B_t} \right) + \varepsilon_{it}$$

R_{it} is an individual's reported probability

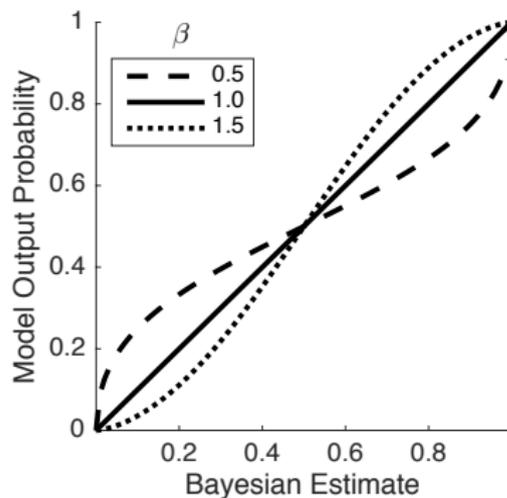
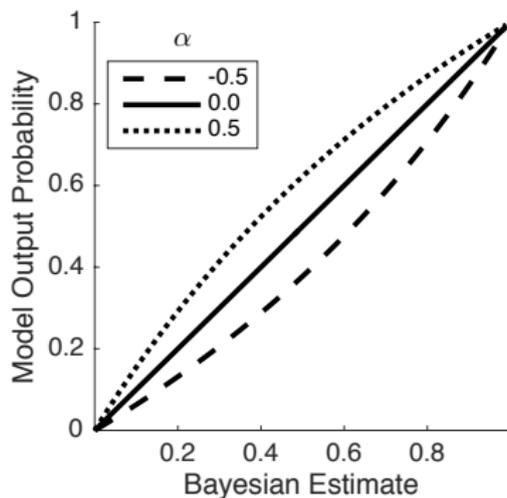
B_t is the optimal Bayesian forecast on trial t , given the DGP and the history of ring realizations

ε_{it} is random noise in the subjective estimates: $\varepsilon_{it} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma_i^2)$

α_i allows for systematic over/under-estimation

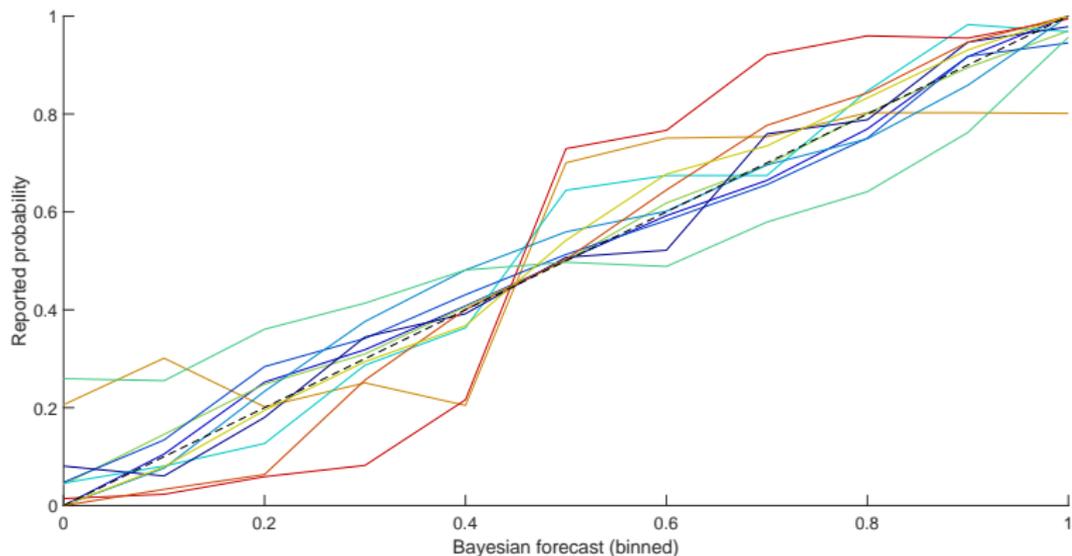
β_i allows for conservatism/repulsion with a flexible crossover (or indifference) point

Probability Distortions



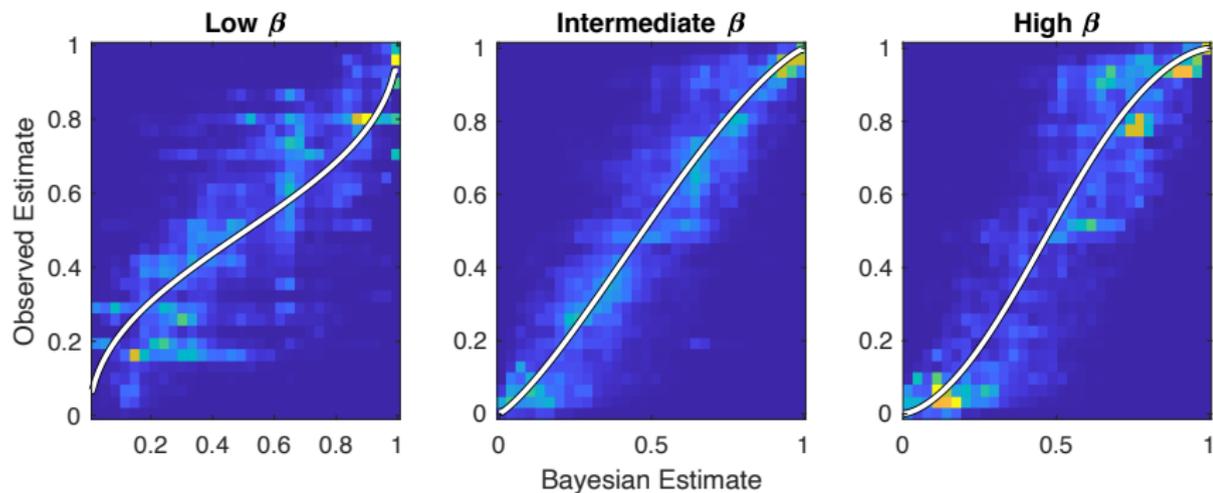
optimal forecast: $\alpha = 0$; $\beta = 1$

Probability Distortions in the Data



actual versus optimal Bayesian forecasts

Systematic Biases at the Individual Level

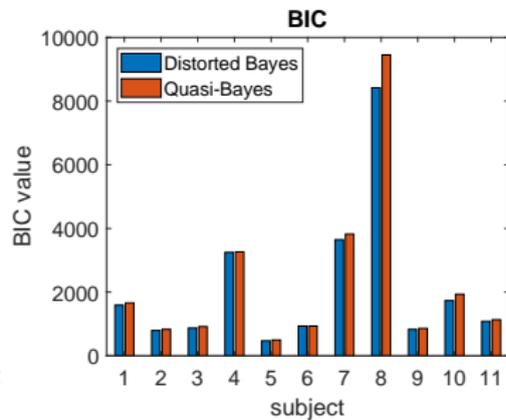
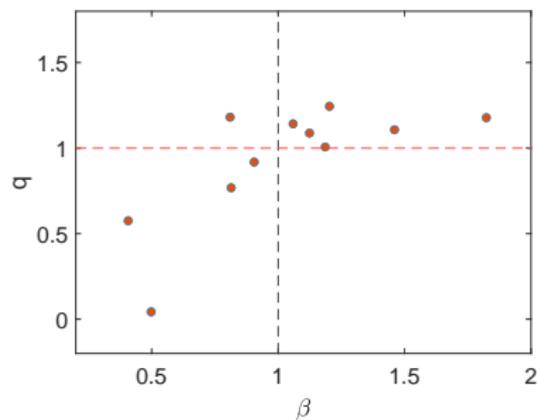


statistically significant differences in β

Quasi-Bayesian Forecasting

- The distorted Bayesian forecast can be given a deeper foundation, as approximating boundedly rational Bayesian probability estimation
- Consider one such quasi-Bayesian forecasting rule that puts an incorrect weight on the likelihood ratio in the application of Bayes rule
- Let q govern the weight put on the likelihood in the updating of the posterior distribution
 - $q = 1$ \Rightarrow Bayesian forecast
 - $q \in (0, 1)$ \Rightarrow under-weighting the likelihood
 - $q > 1$ \Rightarrow over-weighting likelihood relative to the prior
 - $q < 0$ \Rightarrow incorrect use of the data: posterior moves in the wrong direction given the data received

Quasi-Bayesian Forecasting



for each subject we estimate q_i and σ_i^2

Quasi-Bayesian Forecasting

- As predicted by the theory, the estimates of q are positively correlated with the non-linear bias β in the baseline model, potentially providing a potential explanation for the direction of the observed distortions
- But: the BIC values of the QB model are larger than those of the distorted Bayesian model for all participants [though modest for half of the participants]
- **Through its flexible inclusion of nonlinear biases, the distorted Bayesian benchmark appears to capture biases in probability forecasting beyond the incorrect weighting of information**

Experience-Based Alternatives

- A challenge of the Bayesian and quasi-Bayesian models of probability estimation is their significant computational complexity and biological implausibility
- One possibility is that DMs act *as if* they produced forecasts that are functions of the Bayesian forecast
- If so, even though models based on the Bayesian forecast may not be descriptively accurate, they can still be used to understand and predict subjective forecasts
- Alternatively, DMs may be using very different forecasting algorithms
- These may generate what *look* like biases in Bayesian forecasting but may in fact reflect completely different mechanisms

Constant Gain Learning

- One popular alternative: error-based updating of the forecast:

$$D_t = D_{t-1} + \delta (s_t - D_{t-1})$$

D_t is the model's forecast after the ring realization of trial t

s_t is that ring realization

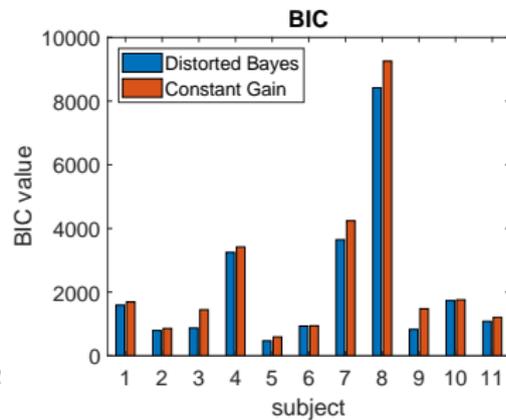
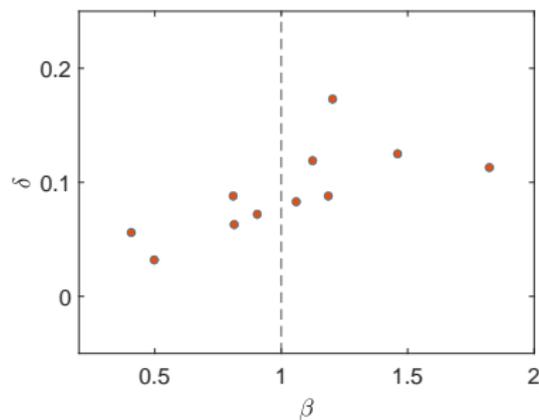
δ is the learning parameter

Constant Gain Learning

- Can generate faster or slower adaptation to change depending on the learning rate (trading off accuracy vs. speed of adaptation)
- Very popular as a descriptive model of what the brain is actually doing
- Advantage of low computational complexity
- Forsgren et al. (2020) show it outperforms Gallistel et al.'s (2014) model of step-wise updating !

“Our results suggest that trial-by-trial updating plays a prominent role in the cognitive processes underlying learning of non-stationary distributions”

Constant Gain Learning



for each subject we estimate δ_i and σ_i^2

Summary: Probability Distortions

(Khaw, Stevens & Woodford, 2020)

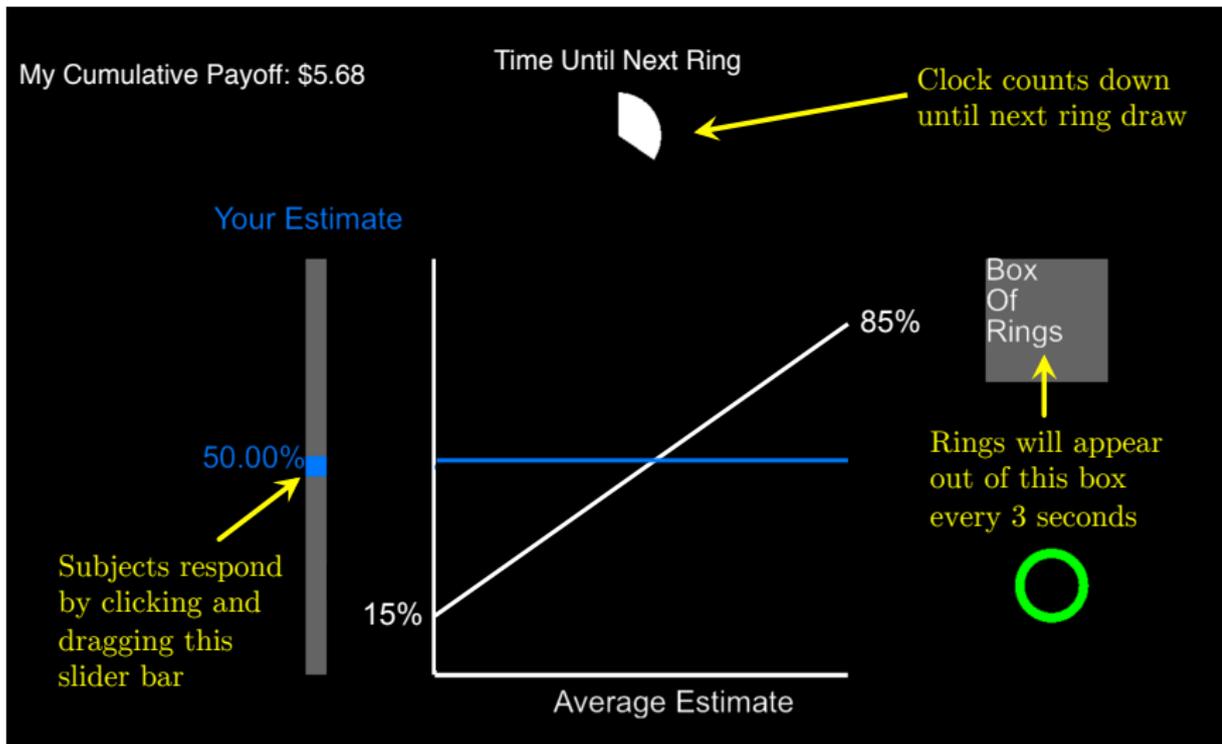
- Strong statistical evidence in support of biases at the individual level, given very large number of trials per subject
- It is possible that people have evolved to explore different forecasting strategies in a way that makes the group collectively approximately Bayesian, even though individuals deviate substantially from it (Wojtowicz, 2020)
- Additional studies needed to confirm prevalence of types and if subjects cluster at extremes of conservatism- polarization spectrum, or are aligned on a continuum
- Intriguing that biased Bayesian forecast remains most compelling way to characterize the forecasts in our data
- Future work should also aim to causally identify the factors that promote subject-level heterogeneity in distorted estimates, for instance by manipulating details of the experiment

EXPERIMENT 2
APPLICATION TO STRATEGIC SETTING
Khaw, Stevens & Woodford, 2021

The Experiment

- Subjects predict the percentage of green rings in virtual box with green + red rings
- Percentage depends in part on the **group forecast**:
 $p_t = z_t + \alpha \hat{P}_t$
- Exogenous z_1 drawn from uniform on
 $Z = \{0.05, 0.15, 0.25\}$
- After each ring draw, 0.05 probability of an intercept shift
 - if no shift, $z_{t+1} = z_t$
 - if shift, z_{t+1} is another independent draw from Z
- REE forecast: $p_t^{RE} = z_t / (1 - \alpha) \in \{0.17, 0.50, 0.83\}$

The Experiment

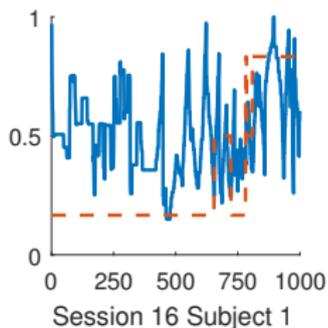
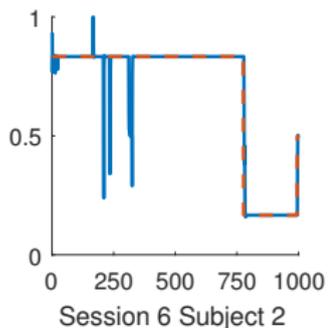
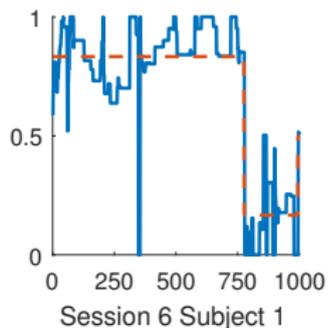
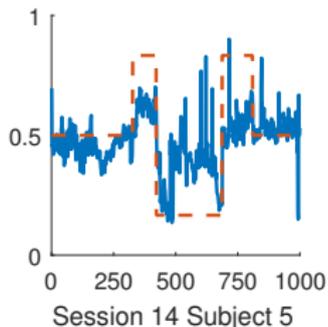
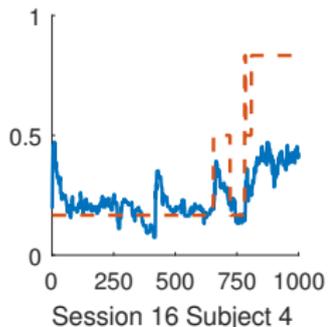
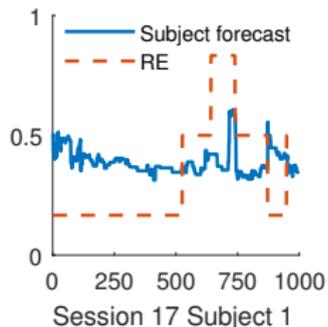


what the subjects see

The Experimental Results

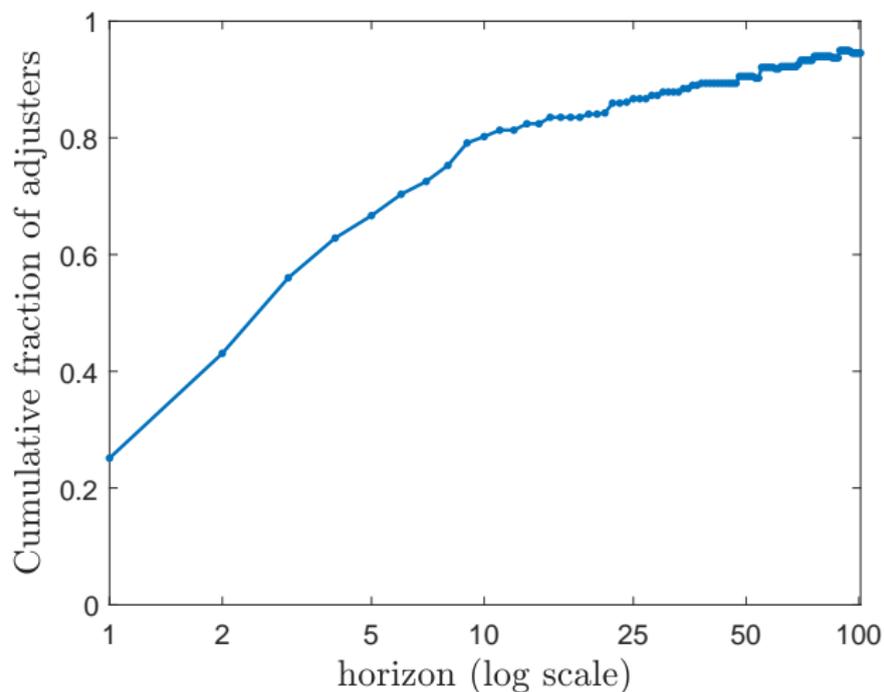
- Simple probability estimation task in which individual payoff depends on both exogenous term and the group's forecast
- Forecasts are **noisy** and **biased**
- Responses **converge slowly** to theoretical equilibria
- Behavior is too **heterogeneous** to be described by single rule
- **Low strategic sophistication** compared to games literature
- Prevalence of **experience-based** forecasting, rather than model-based method

Examples of the Data: Individual Subjects



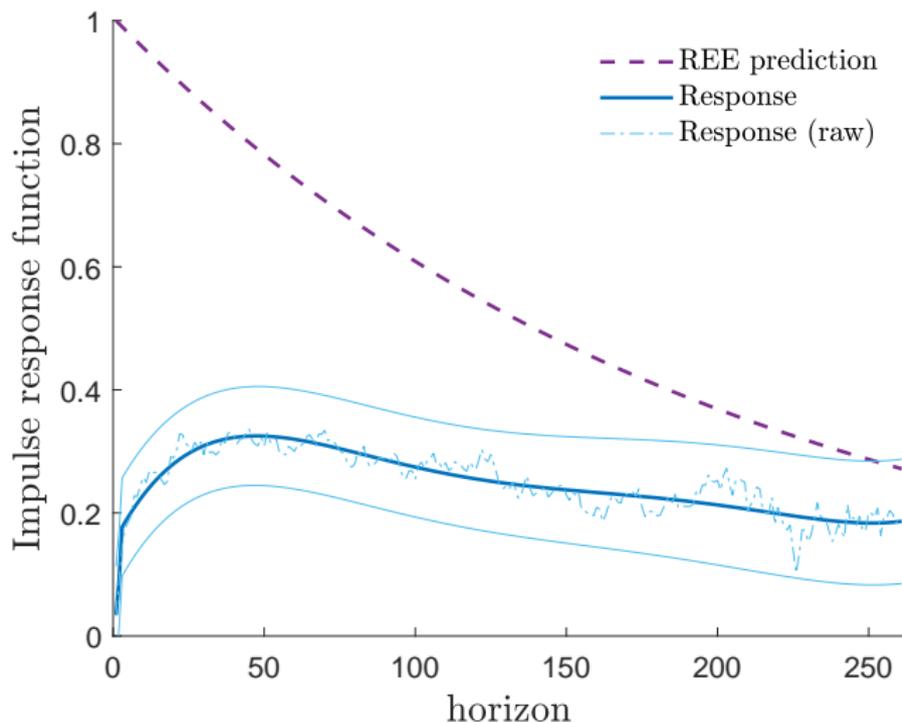
mix of strategies

Cumulative Fraction of Adjusters



RE model predicts 100% on impact
(unconditional frequency of adjustment = 28%)

Dynamics of Adjustment to Intercept Shift



very incomplete, sluggish adjustment
more so than in non-strategic settings

Modeling Subjects' Forecasts

- Log-odds model of the subjective probabilities, with noise

$$\log \left(\frac{R_{it}}{1 - R_{it}} \right) = \alpha_i + \log \left(\frac{M_t}{1 - M_t} \right) + \varepsilon_{it},$$

where R_{it} is individual's reported forecast on trial t , M_t is the model forecast, and $\varepsilon \sim \mathcal{N}(0, \sigma_i^2)$ is idiosyncratic noise

- **Level-k**: natural starting point for the deterministic model forecast

Stochastic Adaptive Level k

- Consider **modified Level-k**, in which Level 0 is not purely random
- Level 0 subjects **watch the rings** and form their best guess using **adaptive learning** algorithm with a constant gain coefficient γ_i
- Higher level subjects base their expectations of others on their own ring-watching

⇒ For each subject, we estimate the best-fitting set $(k_i, \sigma_i, \gamma_i)$

Strategies for Responding to the Fundamental

- For 44% of subjects the **deductive level-0 rule** is the best fitting rule
 - linear fn. of the history of rings, estimated at the subject level
 - very small average gain γ_i of 0.03
 - average additive bias α_i of 0.2
 - no strategic considerations
 - no discrete response to intercept shift
- The level-0 **“ring watchers”** generate gradual adjustment
- They exhibit wide **dispersion in the sensitivity** to the ring history and a wide range of noise levels

Strategies for Responding to the Fundamental

- For 42% of our subjects the best fitting rule is **level 1**
 - respond to intercept (with noise) but assume others set forecast based on history of rings
 - forecast of others is based on own subject-level estimate
 - larger gain on average, at γ_i of 0.05
 - lowest noise on average, seem to be best estimated
 - most simplistic form of strategic thinking
- For 12% of our subjects the best fitting is **level 2**
 - gain average of γ_i of 0.06
 - high noise on average, inattentive response to intercept shift
- 1 subject is essentially **RE** with noise

Summary: Strategic Considerations

- Adjustment is noisy and **very sluggish**
 - more so than in experiments with non-strategic contexts
- Strategic sophistication is dispersed and **limited**
 - more so than in experiments from strategic games literature
- Sophistication and accuracy depend a lot on **strategic context**
 - as in Fehr and Tyran (2008), etc.
- **Experience-based forecasting** seems much more empirically-relevant than model-based forecasting
 - has **implications for policy changes** s.a. raising inflation target
 - crucial to allow for both in the experimental design; otherwise may under-estimate how well people could do on a task (our subjects had the model representation but many chose not to use it)

Conclusions

- We've come a long way from FIRE
- More controlled lab experiments can provide useful diagnostics for selecting between alternative models

THANK YOU