

CEBI WORKING PAPER SERIES

Working Paper 20/20

DESIGNING INFORMATION PROVISION
EXPERIMENTS

Ingar Haaland

Christopher Roth

Johannes Wohlfart

ISSN 2596-44TX

CEBI

Department of Economics
University of Copenhagen
www.cebi.ku.dk

Designing Information Provision Experiments*

Ingar Haaland Christopher Roth Johannes Wohlfart

June 29, 2020

Abstract

We review methodological questions relevant for the design of information provision experiments. We first provide a literature review of major areas in which information provision experiments are applied. We then outline key measurement challenges and design recommendations that may be of help for practitioners planning to conduct an information experiment. We discuss the measurement of subjective beliefs, including the role of incentives and ways to reduce measurement error. We also discuss the design of the information intervention, as well as the measurement of belief updating. Moreover, we describe ways to mitigate potential experimenter demand effects and numerical anchoring arising from the information treatment. Finally, we discuss typical effect sizes in information experiments.

Keywords: Experimental Design, Beliefs, Information, Obfuscation

JEL Classification: C90, D83, D91, L82

*Ingar Haaland, University of Bergen and CESifo, email: Ingar.Haaland@uib.no; Christopher Roth, University of Warwick, briq, CESifo, CEPR, CAGE Warwick, email: Christopher.Roth@warwick.ac.uk; Johannes Wohlfart, Department of Economics and CEBI, University of Copenhagen, CESifo, Danish Finance Institute, e-mail: johannes.wohlfart@econ.ku.dk; The activities of the Center for Economic Behavior and Inequality (CEBI) are funded by the Danish National Research Foundation. We thank Philipp Doerrenberg, Lukas Hensel, Philipp Lergetporer, Michel Marechal, Ricardo Perez-Truglia, and Diego Ubfal for very useful suggestions. Hrishikesh Iyengar and Florian Schoner provided excellent research assistance.

1 Motivation

Standard economic theories usually understand choices as a combination of four factors: preferences, constraints, information, and beliefs. The goal of economic experiments is typically to change some features of the choice environment to study how choices are made. Information experiments achieve this by varying the information set available to economic agents. By only providing information to a random subset of the population of interest, information experiments have become a popular method to study how economic agents form beliefs and make choices. For instance, information experiments have been extensively used to study policy questions (Alesina et al., 2018c; Armantier et al., 2016; Coibion et al., 2020a, 2019b; Hjort et al., 2019) and test economic theories (Bursztyn et al., 2012, 2019a). They can be conducted in controlled lab settings, online, or in more natural settings where the population of interest is typically not aware of its participation in an experiment.

Information experiments enable clean identification of the questions of interest by only varying one feature of the information set. One very powerful application of information experiments is to generate exogenous variation in perceptions of real world environments, which themselves cannot be directly changed. For instance, it is impossible to change the characteristics of immigrants, but researchers can vary perceptions of the immigrant population by correcting people's misperceptions (Grigorieff et al., 2020). Similarly, researchers cannot manipulate intergenerational mobility or influence the state of the macroeconomy, but it is possible to change perceptions of intergenerational mobility (Alesina et al., 2018c) or the perceived likelihood of a recession (Roth and Wohlfart, 2019). Finally, researchers

cannot manipulate social norms, but information provision experiments can be used to study the causal effect of perceived social norms on behavior (Bursztyn et al., 2018).

In this article, we review the growing literature on information experiments in economics. In Section 2, we summarize areas in which information experiments have been widely applied. In Section 3, we outline best-practice recommendations for the measurement of beliefs. In Section 4, we discuss the design of the information intervention. In Section 5, we outline important aspects of the measurement of belief updating. In Section 6, we discuss best practice recommendations for mitigating concerns about experimenter demand effects. In Section 7, we discuss online samples that are commonly used for information provision experiments. In Section 8, we discuss typical effect sizes and recommendations for sample sizes in information provision experiments. Finally, we offer concluding remarks in Section 9.

2 Major applications

In this section we provide an overview of areas in economics in which information provision experiments have been widely applied. This review is necessarily incomplete, and focuses on applications in public economics, political economy, macroeconomics and household finance, and labor economics.¹

¹Our review does not include information provision experiments operating in a laboratory setting in which respondents receive information about features of the laboratory environment or the behavior of other participants in the lab. The review also does not feature work studying the role of the media in shaping beliefs and behavior.

Public Economics Information provision experiments are used in many areas of public economics. Chetty and Saez (2013) conduct an experiment with 43,000 EITC recipients in which a random subset received personalized information about the EITC schedule. Doerrenberg and Peichl (2018) examine how social norms affect tax morale, and Blesse et al. (2019) study how beliefs shape preferences for tax simplification. Bérgholo et al. (2017) and Doerrenberg and Schmitz (2015) examine how firms respond to information about audit probabilities, and Bott et al. (2019) study whether people's tendency to evade taxes responds to information about detection probability and moral appeals. Similarly, Perez-Truglia and Troiano (2018) examine how information on financial penalties, shaming penalties, and peer comparisons shape tax delinquents' future repayment rates. De Neve et al. (2019) study the impact of deterrence, tax morale, and simplifying information on tax compliance. Finally, a literature in behavioral public economics has studied how misperceptions about the fuel economy affect consumers' purchasing decisions (Allcott, 2013; Allcott and Knittel, 2019; Allcott and Taubinsky, 2015).

Political Economy Information experiments are also commonly used to study how beliefs affect policy attitudes, such as people's demand for redistribution (Alesina et al., 2018c; Chen et al., 2016; Cruces et al., 2013; Fehr et al., 2019a,b; Gärtner et al., 2019; Karadja et al., 2017; Kuziemko et al., 2015), their support for government spending (Lergetporer et al., 2018a; Roth et al., 2019), their views on educational inequality (Lergetporer et al., 2018c) and tuition fees (Lergetporer et al., 2016), their support for immigration (Alesina et al., 2018a; Bansak et al., 2016; Barrera et al., 2020; Facchini et al., 2016; Grigorieff et al., 2020; Haaland and Roth, 2019b; Hopkins et al., 2019; Lergetporer et al., 2017), their tendency

to discriminate against immigrants (Alesina et al., 2018b), or their support for affirmative action (Haaland and Roth, 2019a; Settele, 2019). In the context of the coronavirus pandemic, Settele and Shupe (2020) study the role of beliefs for supporting lockdown measures, and Rafkin et al. (2020) study determinants of inference from official government projections.

Information experiments are also conducted to better understand the demand for news and the implications of media on behavior. Chopra et al. (2019) study how perceived informativeness affects people's demand for economic and political news. Bursztyn et al. (2020b) study how the common knowledge of rationales (which are usually supplied by the media) affects the public expression of xenophobia through the lens of a signaling model.

In the context of natural field experiments, researchers have used information treatments to study voting behavior (Cruz et al., 2017, 2018; Kendall et al., 2015; Orkin, 2019) or to study strategic behavior of political activists (Hager et al., 2019a,b) and protesters (Cantoni et al., 2019; Hager et al., 2019c). Finally, researchers have studied whether information about improved public services can help build trust in state institutions and move people away from non-state actors in Pakistan (Acemoglu et al., 2019), and showed that, in the context of social service program delivery, mailing cards with program information to targeted beneficiaries increases the subsidy they receive from a subsidized rice program (Banerjee et al., 2018).

Macroeconomics Information provision experiments are widely used in macroeconomics to study expectation formation of households and firms in the context of beliefs about inflation (Armantier et al., 2016; Binder and Rodrigue, 2018; Cavallo et al., 2016a; Coibion et

al., 2019a, 2020a, 2019b,c, 2018a), house prices (Armona et al., 2019; Fuster et al., 2018; Qian, 2019), recessions (Roth and Wohlfart, 2019), and exposure to macroeconomic risk (Roth et al., 2020). In the context of the Covid-19 pandemic, Coibion et al. (2020b) study how provision of information about policy responses shapes households' macroeconomic expectations and spending plans. Binder (2020) provides evidence on the effect of providing Fed communication about its coronavirus response on household expectations.

Household Finance Research in household finance has studied the effects of information provision on retirement savings (Beshears et al., 2015; Dolls et al., 2018). Moreover, Bursztyn et al. (2019b) examine how moral appeals affect debt repayment. Bursztyn et al. (2012) study the mechanisms underlying peer effects in financial decisions. Bottan and Perez-Truglia (2020) study the causal effect of home price expectations on the timing of home sales using a large-scale field experiment, featuring administrative data. Laudenbach et al. (2020) use an information experiment to study the causal effect of subjective beliefs about the stock market and stock returns on investment choices. In the context of the coronavirus pandemic, Hanspal et al. (2020) provide experimental evidence that beliefs about the duration of the stock market recovery shape households' expectations about their own wealth and their planned investment decisions and labor market activity.

Labor, education and health economics Information provision experiments have been applied to study job search (Abebe et al., 2020; Altmann et al., 2018; Belot et al., 2018a,b; Carranza et al., 2019; Franklin, 2017), social norms (Bursztyn et al., 2018), educational aspirations (Lergetporer et al., 2018b), schooling decisions (Jensen, 2010), major choice

(Bleemer and Zafar, 2018; Conlon, 2019; Wiswall and Zafar, 2014) as well as school choice (Andrabi et al., 2017). Researchers have shown that information about school quality affects parental investment decisions (Greaves et al., 2019) and that parents' beliefs about children's ability affect their educational investments (Dizon-Ross, 2019). Coffman et al. (2017) highlight that information about peers' choices can affect job choice. Researchers in behavioral labor economics have also studied how information provision about peers affects people's work morale and labor market behavior (Card et al., 2012; Cullen and Perez-Truglia, 2018). In agricultural economics, information provision experiments are also widely applied, for example, Hanna et al. (2014) study the effects of information on farmers' behavior.

Finally, information provision experiments have been used to study information relevant for health behaviors. For example, Nyhan and Reifler (2015) and Nyhan et al. (2014) study the effects of information about vaccines. Fitzsimons et al. (2016) find that information provision to mothers in Malawi increases children's food consumption. Carneiro et al. (2019) study an intervention targeting early life nutrition, which also provides nutritional information. Dupas (2011) studies the effect of HIV information on sexual behavior. Kerwin (2018) examines how information about HIV risks affects sexual behaviors, Barari et al. (2020) study public health messaging and social distancing in the context of the coronavirus pandemic, while Fetzer et al. (2020) study perceptions of the pandemic risk factors.

3 Measuring Beliefs

Information provision experiments aim to study the effect of information on people's beliefs or to generate exogenous variation in beliefs to study the effect of beliefs on other outcomes. This section discusses whether one should measure prior beliefs before the information and posterior beliefs after the information, as well as issues related to the measurement of beliefs, including advantages and disadvantages of measuring qualitative or quantitative point beliefs versus probabilistic beliefs, the use of external benchmarks for the elicitation of beliefs, the framing of belief elicitation and techniques on how to deal with measurement error. Finally, we review the measurement of beliefs using hypothetical vignettes.

3.1 Eliciting prior and posterior beliefs?

There are several advantages to eliciting prior beliefs in information provision experiments. First, in designs with a pure control group (that is, a control group that does not receive any information), the expected directional response to the treatment depends on people's prior beliefs about the information. Second, to interpret effect sizes and examine mechanisms, it is important to measure the extent to which people update from the information.

Eliciting posterior beliefs is important in settings where there is a direct interest in studying the effect of information on these beliefs. Moreover, measuring posterior beliefs is necessary to learn about the size of the first stage in settings where information provision experiments are used to study the causal effect of beliefs on other outcomes. In settings where respondents are provided with information about facts (e.g., Roth et al., 2019),

eliciting posteriors primarily serves to measure attention to the information, and is less strictly needed than in designs where respondents receive a potentially noisy signal about a variable (e.g., Roth and Wohlfart, 2019), where posteriors are used to assess how informative respondents find the provided signal.

A potential downside of designs measuring both priors and posteriors is that such within-designs potentially lead to stronger experimenter demand effects (see Section 6). Alternatively, respondents may be subject to consistency bias in their survey responses (Falk and Zimmermann, 2012), leading to a muted effect of information in within-designs. However, Roth and Wohlfart (2019) do not find any significant effect of eliciting priors on the estimated average learning rate in the context of information about macroeconomic risk. Moreover, in designs with a pure control group, being asked the same question twice might confuse respondents in the control group. One remedy is to use a different elicitation mode for the posteriors compared to the priors (Coibion et al., 2019b).

3.2 Qualitative, quantitative or probabilistic beliefs?

How exactly should one measure beliefs? Should one measure beliefs using qualitative or quantitative survey questions? Should one measure point estimates on quantities or probabilistic beliefs in which people attach probabilities to different states of the world occurring?

Qualitative beliefs: verbal response scales One way to measure beliefs is to present respondents with verbal response scales, e.g. reaching from “very low” to “very high”, or from “strongly agree” to “strongly disagree”. Such belief measures have the simple

advantage that the response options are easy to understand for respondents, but the clear disadvantage that they are not easily interpersonally comparable, which can result in severe identification challenges (Bond and Lang, 2019). For instance, in the context of measuring beliefs about the size of the immigration population, people might hold systematically different views on whether a given fraction of immigrants in the population is “very low” or “very high.” That is, if respondents interpret the verbal answer options differently, making comparisons across individuals is challenging. Moreover, verbal response scales are relatively crude and therefore limit the extent of information that can be conveyed (Manski, 2018). Furthermore, with qualitative beliefs, it is often theoretically ambiguous in which direction people should update their beliefs in response to an information treatment. For instance, to manipulate perceptions about the size of the immigration population in the United States, one could inform treated respondents that 12 percent of the US population are immigrants (Grigorieff et al., 2020; Hopkins et al., 2019). Without a quantitative pre-treatment beliefs measure, it is not clear whether treated respondents should revise their beliefs about the size of the immigration population upwards or downwards in response to this information.

Qualitative beliefs: open-ended questions It is also possible to use open-ended questions to measure beliefs (Bursztyn et al., 2020b; Stantcheva, 2020). The key advantage of such open-ended questions is that respondents are not primed by the available answer categories. In other words, open-ended questions enable researchers to directly measure what “comes to mind”. For example, Stantcheva (2020) examines what considerations people have in mind when thinking about a given policy. Bursztyn et al. (2020b) use such

an open-ended elicitation to study inference about the motives for xenophobic expression. Using a pre-registered text analysis procedure and handcoding of the qualitative responses by research assistants, they use this data for studying inference. They validate their open-ended question with a structured belief measure and establish strong correlations. In the context of macroeconomic expectation formation, Leiser and Drori (2005) and D'Acunto et al. (2019) study people's associations with inflation using open-ended text questions.

Quantitative point beliefs Quantitative point beliefs have the advantage of interpersonal comparability, but they do not allow for individuals to express their uncertainty about outcomes. It is therefore good practice to add a second qualitative question on how sure or confident people were in their previous answer. A second disadvantage of point beliefs is that it is unclear which moment of their subjective belief distribution over potential future outcomes respondents report. While researchers often implicitly or explicitly interpret point beliefs as the mean over the respondent's subjective distribution, respondents may report other moments such as their median or mode belief.² Furthermore, if beliefs are elicited with monetary incentives, people might rationally submit their beliefs about the mode rather than their beliefs about the average. Lastly, people's point beliefs might be sensitive to question framing (Eriksson and Simpson, 2012).

Probabilistic beliefs In probabilistic belief elicitation, respondents state probabilities for the occurrence of different mutually exclusive events. Such probabilistic elicitations have

²For instance, De Bruin et al. (2011) show that survey respondents' point forecasts about future inflation or future wage growth are not consistently associated with means constructed from individual-level subjective probability distributions over future inflation or wage growth, but are often associated with the median or other moments of respondents' reported distribution.

the advantage that there is a well-defined absolute numerical scale that is interpersonally and intrapersonally comparable (Manski, 2018). Probabilistic elicitations were pioneered by Manski (2004) and have been widely and successfully applied in some areas in economics, such as labor economics, and the economics of education in particular (Attanasio et al., 2019; Boneva and Rauh, 2017; Boneva et al., 2019; Wiswall and Zafar, 2016, 2017). These measures allow researchers to directly compute a measure of uncertainty as well as the mode and the mean. More recently, the direct measurement of uncertainty has received additional attention in the context of abstract choice and updating tasks as well as survey expectations (Enke and Graeber, 2019). Enke and Graeber (2019) propose the measurement of cognitive uncertainty and show that when people are cognitively uncertain, they implicitly compress probabilities towards a cognitive default of 50:50 in binary state spaces.

One drawback of probabilistic scales is that a large fraction of the population has difficulties in understanding and interpreting probabilities (Tversky and Kahneman, 1974). A second drawback is that people's stated beliefs are typically influenced by how the outcomes are categorized (Benjamin et al., 2017). A third drawback is that probabilistic questions are more time-consuming and taxing for respondents, which makes the experiment longer and potentially induces higher attrition or a higher fraction of missing responses. Some survey providers might also object to the use of probabilistic questions as they might confuse respondents. Probabilistic elicitations are thus primarily recommended in settings where it is very important to precisely quantify people's uncertainty, especially when the population of interest is relatively numerate.

3.3 Benchmarks

One approach measures beliefs about objects of interest for which there is an objective external benchmark. For instance, in the context of income inequality, one can elicit beliefs about the income share going to the top 1 percent income earners rather than a general question about whether income inequality is “high” or “low.” Measuring subjective beliefs about quantities with well-defined benchmarks has several advantages. First, by eliciting beliefs about a well-defined benchmark the experimenter fixes beliefs about the environment and imposes additional structure on the responses. This in turn may lower heterogeneity in how the question is interpreted and thereby reduce measurement error and make responses across participants more comparable. Second, it allows to characterize the extent of biases in beliefs compared to the benchmark. Third, it enables one to incentivize the belief elicitation in a transparent way. Fourth, the availability of benchmarks allows for the provision of information treatments that are tightly connected with the belief elicitation. Recent applications of belief elicitation reliant on benchmarks are studies on social norm elicitation (Krupka and Weber, 2013), racial discrimination (Haaland and Roth, 2019a), intergenerational mobility (Alesina et al., 2018c), immigration (Alesina et al., 2018a; Grigorieff et al., 2020; Haaland and Roth, 2019b), or infectious disease spread (Fetzer et al., 2020).

3.4 Framing of belief elicitation

In settings in which respondents are relatively experienced they are capable of accurately assessing economic quantities. For example, respondents are relatively good at assessing

the price of gas (Ansolabehere et al., 2013). However, in settings in which respondents are relatively unfamiliar, there will be higher levels of measurement error especially when respondents are unsure about the response scale, for example in the context of unemployment estimates. However, careful framing of questions can reduce measurement error. For example, the provision of anchors which convey information about the response scale can reduce measurement error (Ansolabehere et al., 2013). For instance, Roth et al. (2019) measure beliefs about the debt-to-GDP ratio in the US using different historical or cross-country anchors, and show that the provision of an anchor reduces the dispersion of beliefs and rounding.

3.5 Multiple measurement

Many belief elicitation questions pose respondents difficult questions. The cognitive strain in turn may induce measurement error. How can researchers mitigate the extent of measurement error? Gillen et al. (2019) propose an IV approach, which leverages multiple measurements to deal with classical measurement error. We believe that this is particularly important in the context of (quantitative) belief measurement. When reducing classical measurement error is important, researchers ideally should measure their belief of interest using (i) a qualitative survey question, (ii) a quantitative point estimate, and (iii) a probabilistic question in which respondents attach probabilities to mutually exclusive states of the world. This multiple measurement in turn can be leveraged to employ the IV methods that help to deal with measurement error (Gillen et al., 2019). For instance, Giglio et al. (2020) apply such an IV approach in the context of survey expectations about stock

returns, using both point beliefs and subjective probability distributions. However, since multiple measurements might be cognitively taxing for respondents, their benefits must be weighed against the risk of increasing survey fatigue or higher attrition rates. Moreover, this approach cannot be used in the case of non-classical measurement error.

3.6 Incentives

Do prediction incentives lower measurement error in belief elicitation? There is little systematic evidence on the relevance of prediction incentives in the elicitation of beliefs. In the context of economic beliefs, incentives have been shown to reduce partisan bias in people's stated beliefs about economics and politics (Bullock et al., 2015; Prior et al., 2015). For example, the partisan gap in beliefs about the current unemployment rate shrinks when respondents receive prediction incentives. Relatedly, Settele (2019) shows that gender differences in reported beliefs about the gender wage gap shrink in the presence of incentives. In the context of macroeconomic forecasting, it has been shown that unincentivized survey reports strongly correlate with incentivized belief measures (Armantier et al., 2015), and that incentives do not have any statistically significant effects on reported beliefs (Roth and Wohlfart, 2019). Finally, Grewenig et al. (2020) provide mixed evidence on the relevance of incentives in shaping accuracy. Their evidence highlights that incentives have similar effects as a prompt to google the statistic of interest. This highlights the potential undesirable side-effects of incentives when the information of interest is publicly available. Taken together, incentives seem to have little effect on beliefs in non-political domains and in which the responses cannot be easily googled. Danz et al.

(2020) provide evidence that incentives can actually lower truth-telling in the context of abstract prediction tasks. This further underscores the potential negative side-effects of incentives.

3.7 Hypothetical vignettes

Another approach to measuring beliefs is to ask respondents to make predictions about an outcome under different hypothetical scenarios. The use of such hypothetical vignettes is an increasingly popular approach to measure beliefs in contexts that are difficult to study in a real-world setting, such as in the context of education and human capital (Attanasio et al., 2019; Boneva and Rauh, 2017, 2018; Delavande and Zafar, 2018; Kiessling, 2019; Wiswall and Zafar, 2017), preferences over wealth taxation (Fisman et al., 2017), and in the context of beliefs about the effects of macroeconomic shocks (Andre et al., 2019). Hypothetical vignettes, also referred to as conjoint experiments, are widely used to study preferences over different types of immigration (Bansak et al., 2016; Hainmueller and Hopkins, 2014; Hainmueller and Hiscox, 2010). Hainmueller et al. (2015) show that the responses in the vignettes are highly predictive of real world behaviors.

Hypothetical vignettes have the advantage of allowing the researcher more control over the context specified to respondents. Potential disadvantages of hypothetical vignettes include that the hypothetical nature may lower respondents' effort or induce experimenter demand effects. Finally, it may be cognitively challenging for respondents to think in hypotheticals, which could in turn increase measurement error and reduce external validity.

4 Designing the information intervention

In this section, we discuss issues related to the design of the information intervention. First, we discuss different types of information that have been provided in prior work. Second, we discuss which sources of information are commonly used. Third, we review issues related to the presentation of the information. Fourth, we discuss ways through which researchers can more credibly identify the effects of information rather than the effects of priming individuals on an issue. Finally, we review commonly used methods that employ probabilistic information treatments.

4.1 Types of information

Quantitative information Many survey experiments provide respondents with quantitative information, such as statistics based on official census data (Bottan and Perez-Truglia, 2017; Grigorieff et al., 2020; Kuziemko et al., 2015; Roth et al., 2020) or forecasts about the future of the economy (Armantier et al., 2016; Coibion et al., 2019b; Roth and Wohlfart, 2019). While quantitative information may be hard to understand for a large fraction of the population, it often facilitates the interpretation of experimental findings in the context of a theoretical framework. Moreover, together with elicited priors and posteriors numerical information allows for the calculation of learning rates (see section 8). Many times researchers provide statistical information about the behavior of others (Allcott, 2011; Coffman et al., 2017; Duflo and Saez, 2003). A commonly used strategy provides a random subset of respondents with information about others' effort choices (Cantoni et al., 2019; Hager et al., 2019a,b) or others' beliefs, preferences and actions (Bursztyn et al., 2018, 2019a;

Coibion et al., 2018a).

Anecdotal evidence, stories, and narratives Another highly relevant and important, but different type of information relies on qualitative anecdotes, stories or narratives.³ This information is not based on statistics, but instead provides qualitative information which closely resembles case studies. Experiments systematically studying the role of stories, anecdotal evidence and narratives are still very scarce, and we believe a fruitful area for future research. Anecdotal information can also be communicated via pictures and videos, which may be more effective in conveying information. A literature in development economics has studied how inspirational videos change people's beliefs and economic behavior (Bernard et al., 2014; Riley, 2017).⁴

4.2 Sources of information

In this section we discuss a series of possible sources for information that prior research has used to exogenously vary respondents' beliefs and expectations. Researchers commonly provide respondents with official government statistics (for instance, about the unemployment rate of immigrants (Grigorieff et al., 2020)), research evidence (for instance, about the labor market effects of immigration (Haaland and Roth, 2019b), racial discrimination (Haaland and Roth, 2019a), or intergenerational mobility (Alesina et al., 2018c)). In the context of forward-looking expectations, one method to exogenously vary expectations is the provision of expert forecasts. In the context of macroeconomic forecasts, Roth and

³Bénabou et al. (2018) study the role of narratives from a theoretical perspective.

⁴This is also related to a literature studying how the media affects people's beliefs and their behavior (Banerjee et al., 2019; Bursztyn et al., 2020a, 2017; DellaVigna and Kaplan, 2007; La Ferrara et al., 2012; Martinez-Bravo and Stegmann, 2017; Yanagizawa-Drott, 2014).

Wohlfart (2019) provide respondents with different forecasts about the likelihood of a recession and Hager et al. (2019c) provide different expert forecasts about the anticipated turnout to different protests. In experiments which aim to change perceptions of social norms, researchers provide respondents with information about the views of respondents as measured in other surveys (Bursztyn et al., 2018). Moreover, researchers have also explored the effects of randomly providing news articles or statements from policymakers on people's beliefs and expectations (Coibion et al., 2019b). In general, it is important to consider how credible respondents find the source of information.

It is possible that recipients of information think that the information source is biased, for example, Republicans thinking that the BLS statistics were biased during the Obama era. In that case people will update but taking the perceived bias into account (Cavallo et al., 2016b). For an application of this idea in the context of inflation expectations in Argentina, see Cavallo et al. (2016b). As another example, Jacobsen (2019) provides evidence on how different sources differentially affect belief updating and policy views. It is good practice to include direct questions on how credible and accurate people found the provided information at the end of the survey.

4.3 Presentation of the information

How should researchers present the information in order to maximize the effectiveness of the information intervention? To minimize concerns about demand effects, the treatment should ideally be short and neutrally framed. At the same time, to generate a successful first stage on beliefs, it is important to present the information in a way that maximizes

understanding among respondents. One way to increase the understanding of the treatment message is to supplement the text with a graphical illustration of the information. In designs in which researchers elicit prior beliefs, an intuitive way of presenting the information graphically contrasts prior beliefs with the value from the information treatment (e.g. see Roth and Wohlfart (2019)).

4.4 Priming versus information

One key challenge in information experiments is to disentangle the effects of priming from genuine belief updating.⁵ Common methods to mitigate concerns about priming include (i) eliciting prior beliefs of respondents in both the treatment and control group, (ii) separate the information provision from the main outcomes with follow-up studies, and (iii) to include an active control group (that is, the control group also receives (differential) information). The first approach guarantees that both respondents in the treatment and the control group are primed on the issue of interest. The second approach ensures that any short-lived priming effects are no longer relevant when the main outcomes are elicited. The third approach ensures that respondents across all conditions receive information on the issue of interest, but the information differs in terms of its content. In the following, we discuss the use of active control groups in more depth.

Active versus passive control Many information provision experiments measure prior beliefs on an issue and then provide the treatment group with information on that issue, while a pure control group receives no information at all. An alternative design would

⁵For an excellent review on priming in economics, see Cohn and Maréchal (2016).

measure prior beliefs and then provide the treatment and control group with different information (this approach of using an active control group was pioneered by Bottan and Perez-Truglia (2017); for other recent examples of papers implementing active control groups in information provision experiments, see Hager et al. (2019c); Roth and Wohlfart (2019); Roth et al. (2020); Settele (2019)).

Providing the control group with information has several advantages for studying the causal effect of expectations on behavior. In a design with a pure control group the variation hinges on prior beliefs. The identification mostly comes from individuals with larger misperceptions *ex ante*. An active control group design generates variation in the relevant belief also among individuals with more accurate priors and therefore identifies average causal effects of beliefs on outcomes for a broader population. Moreover, receiving an information treatment may have side effects, such as uncertainty reduction, attention, and emotional responses (especially in designs where respondents have been corrected). Such side effects should arguably be constant across groups that receive different pieces of information. Finally, since prior beliefs may be measured with error and correlated with both observables and unobservables, causal identification and the interpretation of heterogeneous treatment effects are more difficult in designs with a pure control group.

There are also some advantages to having a pure control group. First, having a pure control group makes it easier to interpret correlations between the pre-treatment beliefs and the outcome of interest as beliefs among control group respondents are not affected by the treatment. Second, sometimes the policy relevant question of interest is concerned with the effect of providing a particular piece of information compared to not providing this information. See a discussion of these issues in Roth and Wohlfart (2019) in the context

of experiments on macroeconomic expectations or in Hager et al. (2019c) in the context of strategic interactions in political movements. Furthermore, sometimes it is not possible to have an active control group without deceiving respondents, in which case it is better to have a pure control group or employ a probabilistic design as discussed below.

4.5 Probabilistic information treatments

Researchers have started to use probabilistic information treatments to compare belief updating to Bayesian benchmarks (Eil and Rao, 2011; Mobius et al., 2015; Thaler, 2019; Zimmermann, 2019). In probabilistic information treatments, respondents are told that with a probability p they will learn the truth about a fact, and with $(1 - p)$ they will learn the opposite of the truth. Employing probabilistic information treatments provides researchers with fully exogenous variation in beliefs in settings where only one piece of truthful information about a benchmark is available and otherwise one would have to revert to a design with one treatment group and a control group. It also provides researchers with a Bayesian benchmark for the belief updating. However, it introduces motivated beliefs into the updating process, which could in turn lower the effectiveness of the information treatment (Eil and Rao, 2011; Mobius et al., 2015; Thaler, 2019). Probabilistic information treatments are usually applied to study motivated reasoning in belief updating, rather than studying causal effects of information and beliefs on behaviors. A downside of probabilistic information treatments is that they are more artificial and less natural for respondents.

5 Measuring belief updating

In order to understand the mechanisms through which an information treatment operates it is essential to measure a rich set of beliefs which capture the theoretical mechanisms that may be at play. We first discuss how to circumvent issues related to numerical anchoring. Second, we argue that measurement of beliefs about the provided information should be more commonly used to better understand and interpret the effects of information.

Numerical anchoring An additional methodological concern for quantitative outcome measures elicited after the information provision, such as posterior beliefs about the statistic, is unconscious numerical anchoring. There are several best practices for alleviating concerns about numerical anchoring. First, one can provide irrelevant numerical anchors and test their effects on the posterior belief of interest in order to gauge the importance of such anchoring (Cavallo et al., 2016a; Coibion et al., 2018b; Roth and Wohlfart, 2019). Second, one should measure at least some quantitative beliefs on a scale that differs from the scale on which the information is communicated. Third, one should also employ qualitative measures of beliefs, which are naturally immune to numerical anchoring.

Follow-up surveys Follow-up surveys, conducted a few weeks after the initial information intervention, are an important tool used to mitigate concerns about numerical anchoring, which is a short-lived phenomenon.

Follow-up surveys also alleviate concerns about consistency bias in survey responses (Falk and Zimmermann, 2012). Follow-up surveys to study whether information provision has persistent effects on beliefs, preferences and behaviors are increasingly common and

were pioneered by Kuziemko et al. (2015), Cavallo et al. (2016a) and Coppock (2016) in the context of survey experiments. Usually follow-ups in the context of information experiments take place one to eight weeks after the initial information provision. An exception are Fehr et al. (2019b) whose follow-up takes place one year after the initial information provision.

Measuring beliefs about the information Finally, in order to obtain a better understanding of the effects of the information treatment, we think that researchers should measure trust in and other beliefs about the provided information. For example, Haaland and Roth (2019b) elicit a rich set of beliefs about the research evidence provided to respondents. Naturally, such explicit questions may induce significant experimenter demand effects. One way to mitigate concerns about such experimenter demand effects is to elicit incentivized measures of willingness to pay for the information of interest (Fehr et al., 2019b; Haaland and Roth, 2019a; Hjort et al., 2019).

Cross-learning Another recurring issue in information provision experiments is cross-learning. Specifically, respondents may not only update beliefs about the object of interest, but at the same time change their beliefs about other outcomes. For instance, Coibion et al. (2019a) find that provision of information about inflation not only changes respondents' inflation expectations but also their beliefs about GDP growth. On the one hand, such cross-learning can be seen as a natural by-product of experimental changes in beliefs, as changes in beliefs due to natural variation are similarly often correlated across variables. On the other hand, cross-learning can complicate the interpretation of IV estimates exploiting

randomized information provision, as such estimates are often compared to theoretical benchmarks which do not account for cross-learning. One way to over-come the issue of cross-learning is to hold fixed beliefs about other variables by providing the same information about other variables to respondents in both control and treatment groups. However, simultaneous provision of several pieces of information will arguably reduce attention to the main piece of information and lead to a weaker first stage. In any case, researchers should include measures for beliefs about other variables which could be shifted by the treatment in their survey in order to be able to detect cross-learning and to gauge its extent and implications.

6 Dealing with experimenter demand effects

One concern with information provision experiments are demand effects (de Quidt et al., 2018; Mummolo and Peterson, 2019; Zizzo, 2010). While recent empirical evidence suggests a limited quantitative importance of experimenter demand effects in the context of online surveys (de Quidt et al., 2018; Mummolo and Peterson, 2019), it is still possible that in some contexts treatment effects are confounded by experimenter demand effects as people in the different treatment arms may make differential inference about the experimenter's expectations. In this section we outline best-practice recommendations to mitigate concerns about experimenter demand effects.

Obfuscated follow-ups Haaland and Roth (2019a,b) propose the use of obfuscated follow-ups to mitigate concerns about experimenter demand effects. Obfuscated follow-up

surveys are follow-up studies with the same respondents as in the initial experiment, which are presented as an independent study to participants. Since no treatment is administered in the follow-up study, differential experimenter demand between the treatment and control group is unlikely to be a concern unless respondents nonetheless realize that the follow-up is connected to the main study. Haaland and Roth (2019a,b) take several steps to hide the connection between their main study and their obfuscated follow-up study. First, they collaborated with a market research company where respondents regularly receive invitations to participate in surveys. The marketing company sent generic invitations that only reveal information about pay and expected completion time. Second, they employed two different consent forms for the two surveys. Third, to give the impression that the follow-up is an independent study, they first ask respondents a series of questions about their demographics. Fourth, to further obfuscate the purpose of the follow-up, they pose questions about unrelated issues before asking any of the actual questions of interest. Following the approach proposed by Haaland and Roth (2019a,b), Settele (2019) uses an obfuscated follow-up survey in the context of attitudes towards affirmative action.

Anonymity Anonymity has been argued to be a powerful tool against experimenter demand effects in experimental research (Hoffman et al., 1994). In the context of policy preference experiments, researchers have recently relied on the use of anonymous online petitions in order to mitigate concerns about experimenter demand effects (Grigorieff et al., 2020). A commonly used additional tool are “list methods” which aim to veil the answers of individual respondents and are increasingly applied throughout the social sciences (Bursztyn et al., 2018; Chen and Yang, 2019; Coffman et al., 2016; Lерgetporer et al., 2017).

Incentivized outcomes Over the last few years researchers have started using incentivized outcomes in the context of survey experiments. A commonly used approach is to elicit incentivized donations to political organizations which capture specific policy preferences (Bursztyn et al., 2019a; Grigorieff et al., 2020; Roth et al., 2019). Presumably, demand effects should be lower in tasks in which real money is at stake.

Field outcomes A small number of studies manage to link information provision with natural outcomes from the field, such as policy choices of politicians (Hjort et al., 2019), campaign donations (Perez-Truglia and Cruces, 2016), voting behavior (Cruz et al., 2018; Kendall et al., 2015) and canvassing activity using an online application (Hager et al., 2019a,b), home sales (Bottan and Perez-Truglia, 2020) or stock trading choices of retail investors (Laudenbach et al., 2020). The key advantage of these studies is that they provide unobtrusive behavioral outcome data from a natural setting. Experimenter demand effects are of no concern in many of these natural settings as respondents are often not aware of the fact that they are part of an experiment. In general, given that decisions in the field involve much higher stakes than survey responses, it is unlikely that changes in these outcomes reflect demand effects.

Neutral framing How should researchers frame the information treatments? One way to minimize the relevance of experimenter demand effects is to adopt a neutral framing of the experimental instructions. The neutral framing of instructions usually makes the purpose of the experiment less transparent and draws less attention of respondents to the expectations and wishes of the experimenter.

Obfuscated information treatments One way to mitigate experimenter demand effects is to obfuscate the information treatments. In other words, researchers can try to obfuscate the purpose of the study by providing respondents with additional pieces of information, which are irrelevant, or by giving respondents tasks that give the impression that the purpose of the study is completely unrelated to the actual goal. One possibility is to give people an unrelated reason for why they receive the information of interest. For instance, researchers could tell respondents that they need to proofread or summarize pieces of information. For an example in the context of immigration attitudes, see Facchini et al. (2016). Furthermore, in experiments in which the researcher elicits incentivized prior beliefs, the purpose of the information treatment may be naturally concealed by framing the information treatment as feedback on whether the respondent's answer qualified for an extra payment.

Demand treatments de Quidt et al. (2018) propose the use of demand treatments in order to measure the sensitivity of behavior and self-reports with respect to explicit signals about the experimenter's expectations. For example, they tell respondents that they "expect that participants who are shown these instructions" will act in a particular way. The idea behind their approach is that one can use explicit signals of experimenters' wishes in order to bound the natural action. Roth and Wohlfart (2019) and Mummolo and Peterson (2019) apply demand treatments in the context of survey experiments on macroeconomic expectations and in political science, respectively, and confirm the finding that responsiveness to demand effects is quite moderate.

Measuring beliefs about the study purpose Many research studies in economics and psychology measure beliefs about the study purpose. Demand effects are less likely a concern in an experiment or survey if participants cannot identify the intent of the study (Allcott and Taubinsky, 2015). Allcott and Taubinsky (2015) measure perceptions of study intent, and show that there is strong dispersion in perceived intent within treatment groups, suggesting that it is unclear in which way demand effects might affect behavior.

Heterogeneity by self-monitoring scale Allcott and Taubinsky (2015) argue that if demand effects are driving behavior in experiments, then they should be more pronounced for respondents who are more able to detect the intent of the study and are more willing to change their choices given the experimenter's intent. Allcott and Taubinsky (2015) employ the self-monitoring scale by Snyder (1974), and find no evidence that self-monitoring ability moderates the treatment effect.

Summary Overall, demand effects have been shown to be of limited quantitative importance in online experiments (de Quidt et al., 2018). However, the importance of demand effects will vary a lot across settings. We believe that they can be a concern particularly in sensitive domains and are probably less important in less charged domains such as macroeconomic expectation formation. It is best practice to include some of the above outlined checks, especially in sensitive domains.

7 Samples

We provide an overview of commonly used samples with a particular focus on the US.

7.1 Online panels

We now discuss the advantages and disadvantages of three different types of online samples that are commonly used for conducting information provision experiments: (i) probability-based samples, (ii) online panels representative in terms of observables, and (iii) online labor markets, such as Amazon Mechanical Turk.

Probability-based samples The most representative samples are probability-based panels. The idea behind probability-based samples is that respondents should have a known, non-zero probability of being recruited to the panel. Probability-based samples have the clear advantage that they come with sampling weights, which allows researchers to make more externally valid inferences about the whole population with a known sampling error. The disadvantages of probability-based samples are that they are typically much costlier than convenience samples and that they typically offer the least degree of flexibility in survey design and implementation.

In the United States, a widely used probability-based panel is AmeriSpeak by NORC at The University of Chicago. The panel uses NORC's National Frame, which is designed to provide at least 97 percent sample coverage of the US population. The NORC National Frame is used for several landmark studies in the US, including the General Social Survey, which is one of the most frequently analyzed data sets in the social sciences. Other probability-based samples of the US population open to academic researchers include The RAND American Life Panel, the Understanding America Study at the University of Southern California, and the Ipsos KnowledgePanel (formerly administered by GfK).

Representative online panels The second type of available online sample provide samples that are representative in terms of observables. These survey providers rely on convenience samples where participants typically sign up to join the panel in exchange for monetary rewards. The main advantages of these panels is that the samples can be made representative in terms of some important observable characteristics, such as age, income, race, and gender, and are much more affordable than probability-based panels. Furthermore, they allow for the use of obfuscated follow-up studies. The main disadvantage of these panels is that inferences may be less externally valid and there is a concern that respondents who self-select into online panels are very different from the broader population. However, using German data Grewenig et al. (2018) show that the online and the offline population hardly differ in terms of survey responses in the context of political views and opinions, once the survey method and observable respondent characteristics are controlled for. Two large providers that are widely used in the social sciences are Dynata (formerly Research Now and Survey Sampling International) and Lucid (Wood and Porter, 2019).

Coppock and McClellan (2019) find that samples from Lucid score similarly to the American National Election Study's (ANES) on the Big-5 personality inventory, show similar levels of political knowledge, and recover framing effects similar to the ones observed in a probability-based sample (the General Social Survey). Haaland and Roth (2019a) find similar experimental results using a sample from a representative online panel provider and a probability-based sample. Other comparable providers are YouGov, Respondi, and the Qualtrics panel.

Amazon Mechanical Turk The third type of available online sample are online labor markets, such as Amazon Mechanical Turk, which are widely used in the social sciences and economics (Kuziemko et al., 2015). Coppock (2018) conducts 15 replication experiments and finds a very high degree of replicability of survey experiments in the field of political science with MTurk as compared to nationally representative samples. Horton et al. (2011) replicates several well-known lab experiments using MTurk, concluding that online experiments on MTurk are just as valid as traditional physical lab experiments. However, recent studies suggest that data quality on MTurk has been declining over time, partly through the proliferation of bots (automated computer programs) and non-serious respondents, which threatens the data quality on the platform if sensible screening procedures are not implemented (Ahler et al., 2019; Chmielewski and Kucker, 2019). To maximize data quality on MTurk, one should only allow workers that have completed a large number of previous tasks with a high completion rate. Furthermore, in the actual survey, one should include fraud detection tools to rule out bots, such as a CAPTCHA at the beginning of the survey. While MTurk is less representative than most other survey platforms, the platform has some important advantages. First, data collection speed is typically very fast and it offers researchers maximum flexibility in terms of research design. Second, since users sign up for MTurk with their own credit card, it is also possible to incentivize respondents with real money (respondents from more representative panel platforms are typically paid in panel currencies that can be converted into gift vouchers). Third, it is possible to conduct follow-up studies with low attrition rates (Grigorieff et al., 2020).

7.2 Measuring attention in online surveys

Screeners One concern in online surveys is that respondents are inattentive and speed through the surveys (Krosnick, 1991). We recommend using multiple attention checks in online surveys. Recent research suggests that the inclusion of attention checks does not affect estimated treatment effects, but it allows researchers to study how measured attention affects behavior (Berinsky et al., 2014; Kane and Barabas, 2019). One example of an attention screener is the following:

The next question is about the following problem. In questionnaires like ours, sometimes there are participants who do not carefully read the questions and just quickly click through the survey. This means that there are a lot of random answers which compromise the results of research studies. To show that you read our questions carefully, please enter turquoise as your answer to the next question. What is your favorite color?

There are at least two features of attention checks that we consider important: first, it is important for attention checks to explain to participants why researchers use these attention checks. This explanation can mitigate concerns about negative emotional reactions to the use of attention checks on the part of participants. Second, we think that attention checks should be simple to understand and should not be too cognitively demanding. Therefore, having an unambiguous and easy-to-understand question which is not too cognitively challenging is important. For an excellent review on attention checks, see Berinsky et al. (2014).

Open-ended questions Bots have been identified as a threat to online surveys. On top of standard bot protections, such as CAPTCHAS, we recommend using at least two open-ended questions in the survey, e.g. to inquire about feedback about the survey or to ask about the study purpose. These open-ended questions are a useful tool to assess data quality and to identify bots that may provide identical (and/or non-sensical) responses to different open-ended questions.

8 Typical effect sizes and recommended sample sizes

In this section, we briefly discuss typical effect sizes from information provision experiments.

Learning rates Typical information experiments usually measure belief updating using either qualitative or quantitative questions. In the context of quantitative beliefs, papers usually calculate learning rates. To calculate such learning rates, we require both prior and posterior beliefs in order to quantify updating. Moreover, typically we observe both a treatment group which receives information and a control group, which does not receive any information. To quantify the extent to which the respondents update their beliefs towards the signal they receive during the information treatment one can estimate the following specification:

$$\text{Updating}_i = \beta_0 + \beta_1 \text{Treatment}_i \text{Perc.-gap}_i + \beta_2 \text{Treatment}_i + \beta_3 \text{Perc.-gap}_i + \varepsilon_i$$

where Updating_i is defined as the difference between the respondent's posterior and

prior about the quantity of interest. The perception gap, Perc.-gap_i , is the difference between the true signal and the respondent's prior belief about the signal. The key coefficient of interest, β_1 , captures the extent of belief updating toward the provided signal among respondents in the treatment group, on top of any updating that also happens for respondents in the control group. β_2 captures the average treatment effect on respondents' beliefs to the extent it does not depend on individual priors. β_3 measures changes in beliefs, which depend on the perception gap in the control group.

To give a sense of effect sizes for learning rates, we discuss the estimated learning rates of a few selected papers. Armantier et al. (2016) find a learning rate of 0.393 for 1-year inflation forecasts in response to a professional forecast. Armona et al. (2019) estimate an instantaneous learning rate of 0.18 for house price growth in response to information about past house price growth. In a 2 month follow-up, they estimate a learning rate of 0.13, indicating a high degree of persistence. Roth and Wohlfart (2019) estimate a learning rate of 0.318 for recession expectation in response to a professional forecast. In a two week follow-up they document a learning rate of 0.129, indicating a modest degree of persistence. Taken together, these papers document that people learn from the information provided, and that effects become weaker over time.

Effect sizes on beliefs versus preferences Effect sizes on self-reported attitudes and behavioral measures are typically much smaller in magnitude than effect sizes on belief updating in response to information treatments. For instance, Alesina et al. (2018c) employ an information treatment to generate exogenous variation in perceptions of social mobility. While perceptions about the probability of remaining in the bottom quintile of the income

distribution increase by 9.7 percentage points—thus making treated respondents substantially more pessimistic about the social mobility process—the authors find essentially no average impact on policy preferences. Similarly, an experiment by Kuziemko et al. (2015) provides respondents with accurate information about the income distribution. They find a large effect on beliefs about income inequality: treated respondents are 12 percentage points more likely to believe that income inequality has increased. By contrast, policy preferences are largely unaffected by the treatment. Haaland and Roth (2019b) report results from an experiment where effect sizes on beliefs and preferences are quite similar in magnitude. Specifically, they provide respondents with research evidence showing no adverse labor market impacts of low-skilled immigration. Treated respondents become 17.1 percent of a standard deviation more optimistic about the labor market impacts of low-skilled immigrants and 14.1 percent of a standard deviation more in favor of low-skilled immigration.⁶

Behavioral elasticities One way to illustrate effect sizes is to express findings in terms of behavioral elasticities. In this approach researchers instrument the endogenous belief of interest with the information treatment. For example, Cullen and Perez-Truglia (2018) find that increasing the perceived manager salary by 10% would increase the number of hours worked by 1.5%. Bottan and Perez-Truglia (2020) find that a 1 percentage point increase in home price expectations reduces the probability of selling within 6 months by 2.45 percentage points. Roth and Wohlfart (2019) find that a 10 percentage point increase in

⁶It is worth noting that Nyhan and Reifler (2010) find that corrections of beliefs frequently fail to reduce misperceptions among the targeted ideological group. Indeed, Nyhan and Reifler (2010) document several instances of a “backfire effect” in which corrections actually increase misperceptions among the group in question.

the perceived likelihood of a recession leads to a decrease in planned consumption growth by 13 percent of a standard deviation. The key advantage of this approach is that it makes it easier to compare results across settings. The key disadvantage is that the exclusion restriction needed for such an approach may not hold as the information provided may change several beliefs simultaneously.

Sample sizes In light of the rather small or moderate effect sizes on preference measures typically observed in information provision experiments, we recommend employing relatively large samples. Given that typical effects are usually around 15 percent of a standard deviation, and usually lower in the subsequent follow-up surveys, we recommend employing a sample size of at least 700 respondents per treatment arm of interest. Furthermore, since many information experiments yield small or modest effects, it is important to have relatively large samples in order to identify a precise null finding. Naturally, the required sample size will vary greatly across different contexts and needs to be tailored accordingly.

9 Concluding remarks

Our review provides an overview of methods used to study the causal effect of information on beliefs, behaviors and preferences. Our review outlined key measurement challenges and issues surrounding the measurement and the experimental manipulation of beliefs. The key focus of the review centers on methods to deal with (i) the design of information treatments, and (ii) undesirable side effects arising from information treatments, such as numerical anchoring and experimenter demand effects.

Some of the key open questions in this literature surround the exact mechanisms through which information affects beliefs and behaviors. For example, the role of attention and memory in information provision experiments is not well-understood.⁷ New methods which shed more detailed light on the role of attention will thus likely be at the center of future research in this area.

References

Abebe, Girum Tefera, Stefano Caria, Marcel Fafchamps, Paolo Falco, Simon Franklin, and Simon Quinn, “Anonymity or Distance? Job Search and Labour Market Exclusion in a Growing African City,” *Review of Economic Studies*, 2020.

Acemoglu, Daron, Ali Cheema, Asim I Khwaja, and James A Robinson, “Trust in State and Non-State Actors: Evidence from Dispute Resolution in Pakistan,” *Journal of Political Economy*, 2019.

Ahler, Douglas J, Carolyn E Roush, and Gaurav Sood, “The Micro-Task Market for Lemons: Data Quality on Amazon’s Mechanical Turk,” *Working Paper*), 2019.

Alesina, Alberto, Armando Miano, and Stefanie Stantcheva, “Immigration and redistribution,” Technical Report, National Bureau of Economic Research 2018.

—, **Michela Carlana, Eliana La Ferrara, and Paolo Pinotti**, “Revealing Stereotypes: Evidence from immigrants in schools,” Technical Report, National Bureau of Economic Research 2018.

—, **Stefanie Stantcheva, and Edoardo Teso**, “Intergenerational Mobility and Preferences for Redistribution,” *American Economic Review*, 2018, 108 (2), 521–554.

Allcott, Hunt, “Social norms and energy conservation,” *Journal of Public Economics*, 2011, 95 (9-10), 1082–1095.

⁷For recent formal models of attention, see work by Bordalo et al. (2016, 2017); Gennaioli and Shleifer (2010).

- , “The welfare effects of misperceived product costs: Data and calibrations from the automobile market,” *American Economic Journal: Economic Policy*, 2013, 5 (3), 30–66.
 - **and Christopher Knittel**, “Are consumers poorly informed about fuel economy? Evidence from two experiments,” *American Economic Journal: Economic Policy*, 2019, 11 (1), 1–37.
 - **and Dmitry Taubinsky**, “Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market,” *American Economic Review*, 2015, 105 (8), 2501–38.
- Altmann, Steffen, Armin Falk, Simon Jäger, and Florian Zimmermann**, “Learning about job search: A field experiment with job seekers in Germany,” *Journal of Public Economics*, 2018, 164, 33–49.
- Andrabi, Tahir, Jishnu Das, and Asim Ijaz Khwaja**, “Report cards: The impact of providing school and child test scores on educational markets,” *American Economic Review*, 2017.
- Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart**, “Subjective Models of the Macroeconomy: Evidence from Experts and a Representative Sample,” *Available at SSRN 3355356*, 2019.
- Ansolabehere, Stephen, Marc Meredith, and Erik Snowberg**, “Asking about numbers: Why and how,” *Political Analysis*, 2013, 21 (1), 48–69.
- Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar**, “The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment,” *Review of Economics and Statistics*, 2016, 98 (3), 503–523.
- , **Wändi Bruine de Bruin, Giorgio Topa, Wilbert Klaauw, and Basit Zafar**, “Inflation Expectations and Behavior: Do Survey Respondents Act on their Beliefs?,” *International Economic Review*, 2015, 56 (2), 505–536.
- Armona, Luis C, Andreas Fuster, and Basit Zafar**, “Home Price Expectations and Behavior: Evidence from a Randomized Information Experiment,” *Review of Economic Studies*, 2019.

Attanasio, Orazio, Teodora Boneva, and Christopher Rauh, “Parental Beliefs about Returns to Different Types of Investments in School Children,” *Working Paper*, 2019.

Banerjee, Abhijit, Eliana La Ferrara, and Victor H Orozco-Olvera, “The entertaining way to behavioral change: Fighting HIV with MTV,” 2019.

—, **Rema Hanna, Jordan Kyle, Benjamin A Olken, and Sudarno Sumarto**, “Tangible information and citizen empowerment: Identification cards and food subsidy programs in Indonesia,” *Journal of Political Economy*, 2018, 126 (2), 451–491.

Bansak, Kirk, Jens Hainmueller, and Dominik Hangartner, “How economic, humanitarian, and religious concerns shape European attitudes toward asylum seekers,” *Science*, 2016, 354 (6309), 217–222.

Barari, Soubhik, Stefano Caria, Antonio Davola, Paolo Falco, Thiemo Fetzer, Stefano Fiorin, Lukas Hensel, Andriy Ivchenko, Jon Jachimowicz, Gary King et al., “Evaluating COVID-19 public health messaging in Italy: Self-reported compliance and growing mental health concerns,” *medRxiv*, 2020.

Barrera, Oscar, Sergei Guriev, Emeric Henry, and Ekaterina Zhuravskaya, “Facts, alternative facts, and fact checking in times of post-truth politics,” *Journal of Public Economics*, 2020, 182, 104123.

Belot, Michele, Philipp Kircher, and Paul Muller, “How wage announcements affect job search—a field experiment,” 2018.

—, —, and —, “Providing advice to jobseekers at low cost: An experimental study on online advice,” *Review of Economic Studies*, 2018, 86 (4), 1411–1447.

Bénabou, Roland, Armin Falk, and Jean Tirole, “Narratives, imperatives, and moral reasoning,” Technical Report, National Bureau of Economic Research 2018.

Benjamin, Daniel J, Don A Moore, and Matthew Rabin, “Biased Beliefs About Random Samples: Evidence from Two Integrated Experiments,” Working Paper 23927, National Bureau of Economic Research October 2017.

- Bérgolo, Marcelo L, Rodrigo Ceni, Guillermo Cruces, Matias Giacobasso, and Ricardo Perez-Truglia**, “Tax audits as scarecrows: Evidence from a large-scale field experiment,” Technical Report, National Bureau of Economic Research 2017.
- Berinsky, Adam J, Michele F Margolis, and Michael W Sances**, “Separating the Shirkers from the Workers? Making Sure Respondents Pay Attention on Self-Administered Surveys,” *American Journal of Political Science*, 2014, 58 (3), 739–753.
- Bernard, Tanguy, Stefan Dercon, Kate Orkin, Alemayehu Taffesse et al.**, “The future in mind: Aspirations and forward-looking behaviour in rural Ethiopia,” *Working Paper*, 2014.
- Beshears, John, James J Choi, David Laibson, Brigitte C Madrian, and Katherine L Milkman**, “The effect of providing peer information on retirement savings decisions,” *Journal of Finance*, 2015, 70 (3), 1161–1201.
- Binder, Carola**, “Coronavirus Fears and Macroeconomic Expectations,” *Review of Economics and Statistics*, 2020, pp. 1–27.
- **and Alex Rodrigue**, “Household Informedness and Long-Run Inflation Expectations: Experimental Evidence,” *Southern Economic Journal*, 2018, 85 (2), 580–598.
- Bleemer, Zachary and Basit Zafar**, “Intended college attendance: Evidence from an experiment on college returns and costs,” *Journal of Public Economics*, 2018, 157, 184–211.
- Blesse, Sebastian, Florian Buhlmann, and Philipp Doerrenberg**, “Do people really want a simple tax system? Evidence on preferences towards income tax simplification,” *Evidence on Preferences Towards Income Tax Simplification*, 2019, pp. 19–058.
- Bond, Timothy N and Kevin Lang**, “The Sad Truth About Happiness Scales,” *Journal of Political Economy*, 2019, 127 (4), 1629–1640.
- Boneva, Teodora and Christopher Rauh**, “Socio-economic gaps in university enrollment: The role of perceived pecuniary and non-pecuniary returns,” 2017.

- **and** –, “Parental Beliefs about Returns to Educational Investments—The Later the Better?,” *Journal of the European Economic Association*, 2018, 16 (6), 1669–1711.
- , **Marta Golin, Christopher Rauh et al.**, “Can Perceived Returns Explain Enrollment Gaps in Postgraduate Education?,” Technical Report 2019.
- Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer**, “Stereotypes,” *Quarterly Journal Economics*, 2016, 131 (4), 1753–1794.
- , **Nicola Gennaioli, and Andrei Shleifer**, “Memory, attention, and choice,” Technical Report, National Bureau of Economic Research 2017.
- Bott, Kristina M, Alexander W Cappelen, Erik Ø Sørensen, and Bertil Tungodden**, “You’ve got mail: A randomized field experiment on tax evasion,” *Management Science*, 2019.
- Bottan, Nicolas L and Ricardo Perez-Truglia**, “Choosing Your Pond: Revealed-Preference Estimates of Relative Income Concerns,” *Available at SSRN 2944427*, 2017.
- **and** –, “Betting on the House: Subjective Expectations and Market Choices,” *Available at SSRN*, 2020.
- Bruin, Wändi Bruine De, Charles F Manski, Giorgio Topa, and Wilbert Van Der Klaauw**, “Measuring Consumer Uncertainty about Future Inflation,” *Journal of Applied Econometrics*, 2011, 26 (3), 454–478.
- Bullock, John G, Alan S Gerber, Seth J Hill, and Gregory A Huber**, “Partisan bias in factual beliefs about politics,” *Quarterly Journal of Political Science*, 2015, 10 (4).
- Bursztyn, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott**, “Misinformation during a pandemic,” *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2020, (2020-44).
- , **Alessandra L González, and David Yanagizawa-Drott**, “Misperceived social norms: Female labor force participation in Saudi Arabia,” Technical Report, National Bureau of Economic Research 2018.

- , **Daive Cantoni, Patricia Funk, and Noam Yuchtman**, “Polls, the press, and political participation: the effects of anticipated election closeness on voter turnout,” Technical Report, National Bureau of Economic Research 2017.
 - , **Florian Ederer, Bruno Ferman, and Noam Yuchtman**, “Understanding Peer Effects in Financial Decisions: Evidence from a Field Experiment,” *Econometrica*, 2012.
 - , **Georgy Egorov, and Stefano Fiorin**, “From Extreme to Mainstream: How Social Norms Unravel,” Working Paper 23415, National Bureau of Economic Research May 2019.
 - , **Ingar K Haaland, Aakaash Rao, and Christopher P Roth**, “I Have Nothing Against Them, But...,” Technical Report, National Bureau of Economic Research 2020.
 - , **Stefano Fiorin, Daniel Gottlieb, and Martin Kanz**, “Moral incentives in credit card debt repayment: Evidence from a field experiment,” *Journal of Political Economy*, 2019, 127 (4), 000–000.
- Cantoni, Davide, David Y Yang, Noam Yuchtman, and Y Jane Zhang**, “Protests as strategic games: experimental evidence from Hong Kong’s antiauthoritarian movement,” *Quarterly Journal of Economics*, 2019, 134 (2), 1021–1077.
- Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez**, “Inequality at work: The effect of peer salaries on job satisfaction,” *American Economic Review*, 2012, 102 (6), 2981–3003.
- Carneiro, Pedro, I Rasul, G Mason, L Kraftman, and M Scott**, “The Impacts of a Multi-faceted Pre-natal Intervention on Human Capital Accumulation in Early Life,” *Working Paper*, 2019.
- Carranza, Eliana, Robert Garlick, Kate Orkin, and Neil Rankin**, “Job search and hiring with two-sided limited information about workseekers’ skills,” Technical Report, Duke university working paper 2019.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia**, “Inflation Expectations, Learning and Supermarket Prices: Evidence from Field Experiments,” *American Economic Journal: Macroeconomics*, 2016, 9 (3), 1–35.

– , – , and – , “Learning from potentially-biased statistics: Household inflation perceptions and expectations in Argentina,” *Brookings Papers on Economic Activity*, 2016, pp. 59–108.

Chen, Yuyu and David Y. Yang, “The Impact of Media Censorship: 1984 or Brave New World?,” *American Economic Review*, 2019, 109 (6), 2294–2332.

– , **Hui Wang, and David Yang**, “Salience of History and the Preference for Redistribution,” *Available at SSRN 2717651*, 2016.

Chetty, Raj and Emmanuel Saez, “Teaching the tax code: Earnings responses to an experiment with EITC recipients,” *American Economic Journal: Applied Economics*, 2013, 5 (1), 1–31.

Chmielewski, Michael and Sarah C Kucker, “An MTurk crisis? Shifts in data quality and the impact on study results,” *Social Psychological and Personality Science*, 2019, p. 1948550619875149.

Chopra, Felix, Ingar Haaland, and Christopher Roth, “Do People Value More Informative News?,” Working Paper 2019.

Coffman, Katherine B, Lucas C Coffman, and Keith M Marzilli Ericson, “The Size of the LGBT Population and the Magnitude of Anti-Gay Sentiment Are Substantially Underestimated,” *Management Science*, 2016, 63 (10), 3168–3186.

Coffman, Lucas C, Clayton R Featherstone, and Judd B Kessler, “Can social information affect what job you choose and keep?,” *American Economic Journal: Applied Economics*, 2017, 9 (1), 96–117.

Cohn, Alain and Michel André Maréchal, “Priming in economics,” *Current Opinion in Psychology*, 2016, 12, 17–21.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, and Maarten van Rooij, “How Does Consumption Respond to News about Inflation? Field Evidence from a Randomized Control Trial,” *National Bureau of Economic Research Working Paper*, 2019.

- , – , – , and **Michael Weber**, “Forward guidance and household expectations,” Technical Report, National Bureau of Economic Research 2020.
- , **Yuriy Gorodnichenko**, and **Michael Weber**, “Monetary policy communications and their effects on household inflation expectations,” Technical Report, National Bureau of Economic Research 2019.
- , – , and – , “Does Policy Communication During COVID work?,” *Chicago Booth Research Paper*, 2020, (20-15).
- , – , and **Tiziano Ropele**, “Inflation expectations and firm decisions: New causal evidence,” *Quarterly Journal of Economics*, 2019.
- , – , **Saten Kumar**, and **Jane Ryngaert**, “Do You Know That I Know That You Know...? Higher-Order Beliefs in Survey Data,” Technical Report, National Bureau of Economic Research 2018.
- , – , – , and **Mathieu Pedemonte**, “Inflation Expectations As a Policy Tool?,” *Working Paper*, 2018.
- Conlon, John J**, “Major Malfunction: A Field Experiment Correcting Undergraduates’ Beliefs about Salaries,” *Journal of Human Resources*, 2019, pp. 0317–8599R2.
- Coppock, Alexander**, “The Persistence of Survey Experimental Treatment Effects,” *Working Paper*, 2016.
- , “Generalizing from Survey Experiments Conducted on Mechanical Turk: A Replication Approach,” *Political Science Research and Methods*, 2018, pp. 1–16.
- and **Oliver A McClellan**, “Validating the demographic, political, psychological, and experimental results obtained from a new source of online survey respondents,” *Research & Politics*, 2019, 6 (1), 1–14.
- Cruces, Guillermo, Ricardo Perez-Truglia, and Martin Tetaz**, “Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment,” *Journal of Public Economics*, 2013, 98, 100–112.

- Cruz, Cesi, Philip Keefer, and Julien Labonne**, “Buying Informed Voters: New Effects of Information on Voters and Candidates,” Technical Report, Working paper 2017.
- , – , – , and **Francesco Trebbi**, “Making policies matter: Voter responses to campaign promises,” Technical Report, National Bureau of Economic Research 2018.
- Cullen, Zoë and Ricardo Perez-Truglia**, “How much does your boss make? the effects of salary comparisons,” Technical Report, National Bureau of Economic Research 2018.
- Danz, David, Alistair J Wilson, and Lise Vesterlund**, “Belief Elicitation: Limiting Truth Telling with Information or Incentives,” 2020.
- de Quidt, Jonathan, Johannes Haushofer, and Christopher Roth**, “Measuring and Bounding Experimenters Demand,” *American Economic Review*, 2018, 108 (11), 3266–3302.
- Delavande, Adeline and Basit Zafar**, “University Choice: The Role of Expected earnings, Non-pecuniary Outcomes and Financial Constraints,” *Journal of Political Economy*, 2018.
- DellaVigna, Stefano and Ethan Kaplan**, “The Fox News effect: Media bias and voting,” *Quarterly Journal of Economics*, 2007, 122 (3), 1187–1234.
- Dizon-Ross, Rebecca**, “Parents’ beliefs about their children’s academic ability: Implications for educational investments,” *American Economic Review*, 2019, 109 (8), 2728–65.
- Doerrenberg, Philipp and Andreas Peichl**, “Tax morale and the role of social norms and reciprocity. evidence from a randomized survey experiment,” 2018.
- and **Jan Schmitz**, “Tax compliance and information provision—A field experiment with small firms,” *ZEW-Centre for European Economic Research Discussion Paper*, 2015, (15-028).
- Dolls, Mathias, Philipp Doerrenberg, Andreas Peichl, and Holger Stichnoth**, “Do retirement savings increase in response to information about retirement and expected pensions?,” *Journal of Public Economics*, 2018, 158, 168–179.
- Duflo, Esther and Emmanuel Saez**, “The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment,” *Quarterly Journal of Economics*, 2003, 118 (3), 815–842.

- Dupas, Pascaline**, “Do teenagers respond to HIV risk information? Evidence from a field experiment in Kenya,” *American Economic Journal: Applied Economics*, 2011, 3 (1), 1–34.
- D’Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber**, “IQ, expectations, and choice,” Technical Report, National Bureau of Economic Research 2019.
- Eil, David and Justin M Rao**, “The good news-bad news effect: asymmetric processing of objective information about yourself,” *American Economic Journal: Microeconomics*, 2011, 3 (2), 114–138.
- Enke, Benjamin and Thomas Graeber**, “Cognitive Uncertainty,” Technical Report, National Bureau of Economic Research 2019.
- Eriksson, Kimmo and Brent Simpson**, “What do Americans know about inequality? It depends on how you ask them,” *Judgment and Decision Making*, 2012, 7 (6), 741–745.
- Facchini, Giovanni, Yotam Margalit, and Hiroyuki Nakata**, “Countering Public Opposition to Immigration: The Impact of Information Campaigns,” 2016.
- Falk, Armin and Florian Zimmermann**, “A Taste for Consistency and Survey Response Behavior,” *CESifo Economic Studies*, 2012, 59 (1), 181–193.
- Fehr, Dietmar, Daniel Muller, Marcel Preuss et al.**, “Social Mobility Perceptions and Inequality Acceptance,” Technical Report, mimeo 2019.
- , **Johanna Mollerstrom, and Ricardo Perez-Truglia**, “Your Place in the World: The Demand for National and Global Redistribution,” Technical Report, National Bureau of Economic Research 2019.
- Ferrara, Eliana La, Alberto Chong, and Suzanne Duryea**, “Soap operas and fertility: Evidence from Brazil,” *American Economic Journal: Applied Economics*, 2012, 4 (4), 1–31.
- Fetzer, Thiemo, Lukas Hensel, Johannes Hermle, and Christopher Roth**, “Coronavirus perceptions and economic anxiety,” *Review of Economics and Statistics*, forthcoming, 2020.

- Fisman, Raymond, Keith Gladstone, Ilyana Kuziemko, and Suresh Naidu**, “Do Americans want to tax capital? Evidence from online surveys,” Technical Report, National Bureau of Economic Research 2017.
- Fitzsimons, Emla, Bansi Malde, Alice Mesnard, and Marcos Vera-Hernandez**, “Nutrition, information and household behavior: Experimental evidence from Malawi,” *Journal of Development Economics*, 2016, 122, 113–126.
- Franklin, Simon**, “Location, search costs and youth unemployment: experimental evidence from transport subsidies,” *The Economic Journal*, 2017, 128 (614), 2353–2379.
- Fuster, Andreas, Wiederholt Mirko Perez-Truglia Ricardo, and Basit Zafar**, “Expectations with Endogenous Information Acquisition: An Experimental Investigation,” Technical Report, National Bureau of Economic Research 2018.
- Gärtner, Manja, Johanna Mollerstrom, and David Seim**, “Income Mobility, Luck/Effort Beliefs, and the Demand for Redistribution: Perceptions and Reality,” 2019.
- Gennaioli, Nicola and Andrei Shleifer**, “What comes to mind,” *Quarterly Journal of Economics*, 2010, 125 (4), 1399–1433.
- Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus**, “Five Facts about Beliefs and Portfolios,” *Working Paper*, 2020.
- Gillen, Ben, Erik Snowberg, and Leeat Yariv**, “Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study,” *Journal of Political Economy*, 2019, 127 (4), 1826–1863.
- Greaves, Ellen, Iftikhar Hussain, Birgitta Rabe, Imran Rasul et al.**, “Parental Responses to Information About School Quality: Evidence from Linked Survey and Administrative Data,” Technical Report, Institute for Social and Economic Research 2019.
- Grewenig, Elisabeth, Philipp Lerner, Katharina Werner, and Ludger Woessmann**, “Incentives, search engines, and the elicitation of subjective beliefs: evidence from representative online survey experiments,” *Journal of Econometrics*, 2020.

– , – , **Lisa Simon, Katharina Werner, and Ludger Woessmann**, “Can Online Surveys Represent the Entire Population?,” 2018.

Grigorieff, Alexis, Christopher Roth, and Diego Ubfal, “Does Information Change Attitudes Toward Immigrants?,” *Demography*, 2020, pp. 1–27.

Haaland, Ingar and Christopher Roth, “Beliefs About Racial Discrimination and Support for Pro-Black Policies,” Working Paper 7828, CESifo 2019.

– **and** – , “Labor Market Concerns and Support for Immigration,” Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3017420 2019.

Hager, Anselm, Lukas Hensel, Johannes Hermle, and Christopher Roth, “Does Party Competition Affect Political Activism?,” *Working paper*, 2019.

– , – , – , **and** – , “Political activists as free-riders: Evidence from a natural field experiment,” *IZA working paper*, 2019.

– , – , – , **and** – , “Strategic Interdependence in Political Movements and Countermovements,” 2019.

Hainmueller, Jens and Daniel J Hopkins, “The Hidden American Immigration Consensus: A Conjoint Analysis of Attitudes Toward Immigrants,” *American Journal of Political Science*, 2014, 59 (3), 529–548.

– **and Michael J Hiscox**, “Attitudes Toward Highly Skilled and Low-skilled Immigration: Evidence from a Survey Experiment,” *American Political Science Review*, 2010, 104 (01), 61–84.

– , **Dominik Hangartner, and Teppei Yamamoto**, “Validating vignette and conjoint survey experiments against real-world behavior,” *Proceedings of the National Academy of Sciences*, 2015, 112 (8), 2395–2400.

Hanna, Rema, Sendhil Mullainathan, and Joshua Schwartzstein, “Learning through noticing: theory and experimental evidence in farming,” *Quarterly Journal of Economics*, 2014.

- Hanspal, Tobin, Annika Weber, and Johannes Wohlfart**, “Exposure to the COVID-19 Stock Market Crash and its Effect on Household Expectations,” 2020.
- Hjort, Jonas, Diana Moreira, Gautam Rao, and Juan Francisco Santini**, “How Research Affects Policy: Experimental Evidence from 2,150 Brazilian Municipalities,” Technical Report, National Bureau of Economic Research 2019.
- Hoffman, Elizabeth, Kevin McCabe, Keith Shachat, and Vernon Smith**, “Preferences, property rights, and anonymity in bargaining games,” *Games and Economic behavior*, 1994, 7 (3), 346–380.
- Hopkins, Daniel J, John Sides, and Jack Citrin**, “The Muted Consequences of Correct Information About Immigration,” *Journal of Politics*, 2019, 81 (1), 315–320.
- Horton, John J, David G Rand, and Richard J Zeckhauser**, “The Online Laboratory: Conducting Experiments in a Real Labor Market,” *Experimental Economics*, 2011, 14 (3), 399–425.
- Jacobsen, Grant D**, “How do different sources of policy analysis affect policy preferences? Experimental evidence from the United States,” *Policy Sciences*, 2019, 52 (3), 315–342.
- Jensen, Robert**, “The (perceived) returns to education and the demand for schooling,” *Quarterly Journal of Economics*, 2010, 125 (2), 515–548.
- Kane, John V and Jason Barabas**, “No Harm in Checking: Using Factual Manipulation Checks to Assess Attentiveness in Experiments,” *American Journal of Political Science*, 2019, 63 (1), 234–249.
- Karadja, Mounir, Johanna Mollerstrom, and David Seim**, “Richer (and Holier) Than Thou? The Effect of Relative Income Improvements on Demand for Redistribution,” *Review of Economics and Statistics*, 2017, 99 (2), 201–212.
- Kendall, Chad, Tommaso Nannicini, and Francesco Trebbi**, “How do voters respond to information? Evidence from a randomized campaign,” 2015, 105 (1), 322–53.

- Kerwin, Jason**, “Scared straight or scared to death? The effect of risk beliefs on risky behaviors,” *The Effect of Risk Beliefs on Risky Behaviors* (February 9, 2018), 2018.
- Kiessling, Lukas**, “Understanding Parental Decision-making: Beliefs about Returns to Parenting Styles and Neighborhoods,” *Working Paper*, 2019.
- Krosnick, Jon A**, “Response strategies for coping with the cognitive demands of attitude measures in surveys,” *Applied cognitive psychology*, 1991, 5 (3), 213–236.
- Krupka, Erin L and Roberto A Weber**, “Identifying Social Norms Using Coordination Games: Why Does Dictator Game Sharing Vary?,” *Journal of the European Economic Association*, 2013, 11 (3), 495–524.
- Kuziemko, Ilyana, Michael I Norton, Emmanuel Saez, and Stefanie Stantcheva**, “How Elastic are Preferences for Redistribution? Evidence from Randomized Survey Experiments,” *American Economic Review*, 2015, 105 (4), 1478–1508.
- Laudenbach, Christine, Annika Weber, and Johannes Wohlfart**, “Beliefs About the Stock Market and Investment Choices: Evidence from a Field Experiment,” *Working Paper*, 2020.
- Leiser, David and Shelly Drori**, “Naïve understanding of inflation,” *Journal of Socio-Economics*, 2005, 34 (2), 179–198.
- Lergetporer, Philipp, Guido Schwerdt, Katharina Werner, Martin R West, and Ludger Woessmann**, “How information affects support for education spending: Evidence from survey experiments in Germany and the United States,” *Journal of Public Economics*, 2018, 167, 138–157.
- , **Katharina Werner, and Ludger Woessmann**, “Does ignorance of economic returns and costs explain the educational aspiration gap? Evidence from representative survey experiments,” 2018.
- , – , **and –** , “Educational inequality and public policy preferences: Evidence from representative survey experiments,” 2018.

- , **Ludger Woessmann et al.**, “The political economy of university tuition fees: Information provision and income contingency in representative survey experiments,” in “Annual Conference 2016 (Augsburg): Demographic Change” number 145703 Verein für Socialpolitik/German Economic Association 2016.
- , **Marc Piopiunik, Lisa Simon et al.**, “Do Natives’ Beliefs About Refugees’ Education Level Affect Attitudes Toward Refugees? Evidence from Randomized Survey Experiments,” Technical Report, CESifo Group Munich 2017.
- Manski, Charles F**, “Measuring Expectations,” *Econometrica*, 2004, 72 (5), 1329–1376.
- , “Survey measurement of probabilistic macroeconomic expectations: progress and promise,” *NBER Macroeconomics Annual*, 2018, 32 (1), 411–471.
- Martinez-Bravo, Monica and Andreas Stegmann**, “In Vaccines We Trust? The Effects of Anti-Vaccine Propaganda on Immunisation: Evidence from Pakistan,” *CEMFI*, 2017.
- Mobius, Markus M, Muriel Niederle, Paul Niehaus, and Tanya S Rosenblat**, “Managing self-confidence: Theory and experimental evidence,” Technical Report, National Bureau of Economic Research 2015.
- Mummolo, Jonathan and Erik Peterson**, “Demand Effects in Survey Experiments: An Empirical Assessment,” *American Political Science Review*, 2019, 113 (2), 517–529.
- Neve, Jan-Emmanuel De, Clement Imbert, Johannes Spinnewijn, Teodora Tsankova, and Maarten Luts**, “How to Improve Tax Compliance? Evidence from Population-wide Experiments in Belgium,” *Working Paper*, 2019.
- Nyhan, Brendan and Jason Reifler**, “When Corrections Fail: The Persistence of Political Misperceptions,” *Political Behavior*, 2010, 42, 303–300.
- **and** – , “Does correcting myths about the flu vaccine work? An experimental evaluation of the effects of corrective information,” *Vaccine*, 2015, 33 (3), 459–464.
- , – , **Sean Richey, and Gary L Freed**, “Effective messages in vaccine promotion: a randomized trial,” *Pediatrics*, 2014, 133 (4), e835–e842.

- Orkin, Kate**, “Everybody Loves a Winner: A Field Experiment Providing Information on Polls in South Africa,” *Working Paper*, 2019.
- Perez-Truglia, Ricardo and Guillermo Cruces**, “Partisan interactions: Evidence from a field experiment in the united states,” *Journal of Political Economy* (forthcoming), 2016.
- **and Ugo Troiano**, “Shaming tax delinquents,” *Journal of Public Economics*, 2018, 167, 120–137.
- Prior, Markus, Gaurav Sood, Kabir Khanna et al.**, “You cannot be serious: The impact of accuracy incentives on partisan bias in reports of economic perceptions,” *Quarterly Journal of Political Science*, 2015, 10 (4), 489–518.
- Qian, Wei**, “House Price Expectations and Consumption - A Survey-based Experiment,” 2019.
- Rafkin, Charlie, Advik Shreekumar, and Pierre-Luc Vautrey**, “Public Inference from Crisis Policy: Evidence from the COVID-19 Pandemic,” *Available at SSRN 3613446*, 2020.
- Riley, Emma**, “Role models in movies: The impact of Queen of Katwe on students’ educational attainment,” 2017.
- Roth, Christopher and Johannes Wohlfart**, “How Do Expectations About the Aggregate Economy Affect Personal Expectations and Behavior? Experimental Evidence,” *Review of Economics and Statistics*, 2019.
- , **Sonja Settele, and Johannes Wohlfart**, “Beliefs about Public Debt and the Demand for Government Spending,” available at SSRN: <https://ssrn.com/abstract=2809655> 2019.
- , – , **and –** , “Risk Exposure and Attention to the Macroeconomy,” *Working Paper*, 2020.
- Settele, Sonja**, “How Do Beliefs About the Gender Wage Gap Affect the Demand for Public Policy?,” *Available at SSRN 3382325*, 2019.

- **and Cortnie Shupe**, “Lives or Livelihoods? Perceived Tradeoffs and Public Demand for Non-Pharmaceutical Interventions,” *Perceived Tradeoffs and Public Demand for Non-Pharmaceutical Interventions* (April 30, 2020), 2020.
- Snyder, Mark**, “Self-monitoring of expressive behavior,” *Journal of personality and social psychology*, 1974, 30 (4), 526.
- Stantcheva, Stefanie**, “Understanding Economic Policies: What do People Know and How Can they Learn?,” Technical Report, Harvard University Working Paper 2020.
- Thaler, Michael**, “The “Fake News” Effect: An Experiment on Motivated Reasoning and Trust in News,” Technical Report, Working Paper 2019.
- Tversky, Amos and Daniel Kahneman**, “Judgment Under Uncertainty: Heuristics and Biases,” *Science*, 1974, 185 (4157), 1124–1131.
- Wiswall, Matthew and Basit Zafar**, “Determinants of college major choice: Identification using an information experiment,” *Review of Economic Studies*, 2014, 82 (2), 791–824.
- **and** – , “Human capital investments and expectations about career and family,” Technical Report, National Bureau of Economic Research 2016.
- **and** – , “Preference for the Workplace, Investment in Human Capital, and Gender,” *Quarterly Journal of Economics*, 2017, 133 (1), 457–507.
- Wood, Thomas and Ethan Porter**, “The elusive backfire effect: Mass attitudes’ steadfast factual adherence,” *Political Behavior*, 2019, 41 (1), 135–163.
- Yanagizawa-Drott, David**, “Propaganda and conflict: Evidence from the Rwandan genocide,” *Quarterly Journal of Economics*, 2014, 129 (4), 1947–1994.
- Zimmermann, Florian**, “The Dynamics of Motivated Beliefs,” *American Economic Review*, 2019.
- Zizzo, Daniel John**, “Experimenter Demand Effects in Economic Experiments,” *Experimental Economics*, 2010, 13 (1), 75–98.