CONTENTS

Illustrations vii

Acknowledgments xi

Introduction 1

PART I MEASURING EXPECTATIONS THROUGH SURVEYS 11

1 How Should We Measure Economic Expectations? 13

2 What Makes a Good Survey of Expectations? 29

3 What Have We Learned from Existing Surveys of Expectations? 52

4 Mapping Survey Expectations into Theory 70

PART II CREATING EXOGENOUS VARIATION IN EXPECTATIONS 89

5 Designing Information Treatments 91

6 Quantifying and Understanding the Effects of Information Treatments on Expectations 103

7 What Have Information Treatments Taught Us about Policy Communication? 118

8 Mapping Information Treatments into Theories of Learning 133

PART III CHARACTERIZING THE EFFECTS OF SUBJECTIVE EXPECTATIONS ON ECONOMIC DECISIONS 153

9 Measuring Decisions and Outcomes 155

10 Quantifying the Effects of Beliefs on Decisions 167

11 Alternative and Complementary Approaches to RCTs 187

12 Mapping Expectational Pass-throughs to Theory and Applications 200

PART IV APPLICATIONS 219

13 How Does Economic Uncertainty Affect the Decisions of Households and Firms? 221

14 How Do Expectations Affect Individual Portfolio Decisions? 237

15 How Do Beliefs about Peers Shape Decisions? 256

16 Expected Home Prices and Household Decisions 275

Conclusion: Where Do We Go from Here? 292

References 299

Index 317

Introduction

ON THE MORNING of December 21, 2015, we received an email from an editor of *The American Economic Review*, one of the top academic journals in the profession, giving us the opportunity to resubmit a revised version of a paper we had previously submitted to them, but only if we could implement the changes he and the anonymous reviewers requested. We were both thrilled and dismayed. Thrilled because publishing in top journals is what we academics live for and the opportunities to do so are few and far between. Dismayed because we had absolutely no idea how we could possibly do what they were asking us to do.

What was the problem? We had written, along with a wonderful new coauthor in New Zealand, Saten Kumar, a paper that described some properties of firms' inflation expectations using new surveys that we had run together in New Zealand. We were excited about the paper because there was essentially no evidence prior to this about how firms formed expectations about future inflation. And the results that came out were rather striking and unexpected (at least to us): the firms in the survey were remarkably uninformed about inflation and their beliefs about future inflation seemed out of line with what standard macroeconomic models would have predicted. The editor told us in his decision letter that while the results of the survey were intriguing, he didn't think inflation expectations were all that important by themselves. He asked us to provide ironclad evidence that the expectations actually mattered for firms' decisions. This seemed like a "mission impossible" request because we would have to tackle one of the major challenges that has long limited progress in macroeconomics.

When we write down models to describe how the economy functions, we start by characterizing the decisions of individual actors in the economy, such as how households choose how much to consume or save, how firms choose

how much to invest or what prices to set, and so on. In most cases, those decisions are inherently forward-looking in nature: how much we choose to spend today, for example, depends on what we think our job prospects are and what we think will happen to prices in the future, amongst other expectations.

The resulting importance of what people believe about the future for macroeconomic outcomes has long been emphasized in the field. Keynes (1933), for example, described how "animal spirits," or investors' expectations about future returns on capital, could drive investment and economic activity. Economists in the 1970s speculated about how so-called "adaptive expectations" could generate persistent increases in inflation: if firms see that aggregate prices have gone up recently and infer that they will likely keep rising, they will choose to raise their own prices to keep them in line with their competitors', thus fueling more inflation. Wage-price spirals are another mechanism through which overall expectations, and inflation expectations in particular, can drive macroeconomic outcomes. If workers or labor unions expect that prices will rise more rapidly in the future, then they will try to bargain for larger wage increases, which will increase firms' costs and lead them to raise their prices, which would lead workers to want even larger wage increases, and so on.

But these expectations have historically gone unmeasured. The Bureau of Economic Analysis, for example, which provides exceptionally detailed measures of Gross Domestic Production, Consumption and Investment for the U.S. economy, does not publish a measure for investors' animal spirits or workers' inflation expectations. So one approach macroeconomists have long had to rely on is to simply *assume* an expectations formation process to plug into their models. In the 1960s and 1970s, this was typically backward looking: our belief for next year's inflation rate was assumed to be equal to this year's inflation rate. But if people are entirely backward-looking, policymakers will be able to systematically fool them, leading to outcomes that seemed implausible and counterfactual. In the 1980s and 1990s, macroeconomists instead started assuming that people formed beliefs about the future in a way that was consistent with the world in which they lived: the "rational expectations" revolution. But since the world in which these economic actors lived depended in part on how they formed their beliefs about the future, there was an inherent circularity that, when applied to formal models of the economy, sometimes implied multiple equilibria or indeterminacy. And even though macroeconomists could not seem to agree on how to describe the world in which we lived, the agents in our models were assumed to be much smarter and know precisely how that world functioned. Some macroeconomists thought such

extreme rationality and informedness were not any more realistic than the assumption that people were completely backward-looking. But in the absence of evidence on how people actually formed beliefs about the future, macroeconomists were left with arguing about what assumptions to use with little reason to choose one over the other.

A natural solution to this would be to start measuring people's beliefs about the future to figure out exactly how they are formed. And this is exactly what some economists (including us) were actively doing in the 2000s and 2010s. This line of work used historical surveys of expectations and created new and larger surveys to help us better understand the nature of the expectations formation process. As we discuss in Chapters 3 and 4, this literature repeatedly and systematically rejected the benchmark assumption of full-information rational expectations (FIRE) that had become the dominant way of modeling expectations in macroeconomics. In doing so, it also provided support for models that allowed for a slow and/or imperfect adjustment of beliefs to changing economic conditions.

But surveys have many practical issues associated with them, making the measurement of beliefs a challenge. And when the results of the surveys turned out to look very different from what our standard models with FIRE were predicting, it wasn't obvious whether the issue was with the surveys or with the theoretical models. This was essentially the challenge posed to us by the editor: if you want me to believe that firms form expectations as your survey suggests, you need to convince me through direct empirical evidence that those survey expectations you are measuring are actually the relevant ones that affect the economic decisions of these firms.[1]

This required us to somehow address a fundamental identification issue to assign causality: how do beliefs affect decisions? Simply measuring expectations well and looking at how they correlate with firms' decisions is not enough. For example, suppose we observe a positive correlation between firms' inflation expectations and their actual price changes. This could be because if firms expect higher inflation, they anticipate that their prices will

1. To be fair, the editor gave us another option: develop a model of expectations that could explain all of the patterns that we observed and could then be used to consider the implications for policy. We didn't know how to deliver this option, and indeed it subsequently took years for smarter and younger economists (many of them our PhD students!) to write many papers that explained different parts of the results. See for example Afrouzi (2024), Yang (2022) and Kamdar and Ray (2024).

depreciate relative to their competitors' over time, and therefore they might choose to set higher prices in the first place. In this case, causality runs from expectations to prices. But causality could easily run in the opposite direction as well. Suppose a labor market shortage in Alabama caused wages there to rise; then firms in Alabama would tend to raise their prices more than in other states. If those firms in Alabama believe that the wage increase is occurring nationwide, then they would expect all other firms in the economy to progressively raise their prices as well. This would deliver a positive correlation between the prices and inflation expectations of firms across the country, but the expectations would not be *causing* the price changes at all. The editor was asking us to rule out the latter and prove that the former was actually at work, a tall order.

Assessing causality in economics is not new, of course. Applied microeconomists have been at the forefront of doing this using micro-level data for a long time. Joshua Angrist, David Card and Guido Imbens received a Nobel Prize in Economics in 2021 for devising different strategies for doing exactly that, by exploiting natural experiments, quasi-experimental variation, discontinuities, and other approaches. But in most settings, there is no such identifiable variation, or one has to be much more clever than we are to find it. We certainly did not see an easy way to apply these strategies to our surveys of firms in New Zealand. Our thought was then, can we *create* the exogenous variation in the data ourselves that would provide the identification necessary to address concerns about causality? In that respect, there is a gold standard out there that is well known to scientists: randomized control trials (RCTs).

In a typical RCT, participants are randomly assigned to one of two groups: a control group and a treatment group. The treatment group receives the "medicine" that is being tested, which can be a drug, a type of therapy, etc. The control group either does not receive any treatment or, at most, is given a placebo. Because of the randomization procedure used to allocate people to each group, the two groups will be essentially identical on average in terms of characteristics of individuals, but the treatment they receive will differ, which means any difference in outcomes across the two groups can be attributed exclusively to the treatment. In that sense, causality is clear and unambiguous; no fancy statistics are needed.

Applying this type of approach in macroeconomic contexts is not obvious. In a Bloomberg interview in October 2024, Donald Trump described running the Federal Reserve as "the greatest job in government. You show up at the office once a month and you say, 'Let's see, flip a coin,' and everybody treats

you like you're a god." If his description of the Federal Reserve decision-making process as that of flipping a coin were correct, this would be an example of a repeated RCT applied to a macroeconomic context. Unfortunately for monetary economists interested in studying the effects of monetary policy on the economy, this description is far from correct, and most of the interest rate decisions by the Federal Reserve can instead be accounted for as an endogenous response to economic conditions (Romer and Romer 1989), and not as the result of a coin flip. The coin flip approach to monetary policy would obviously negatively affect millions of people, as would many applications of RCTs to macroeconomic environments. Only in development economics have RCTs become quite common, thanks in large part to the work of Esther Duflo and Abhijit Banerjee, for which they were awarded the Nobel Prize in Economics in 2019. RCTs were feasible in development and health economics because one could assign "treatments" (e.g., handing out malaria nets) to randomly selected villages and not others, allowing one to assess how these treatments affected outcomes across villages over time.

So how could we apply something similar in the case of firms' inflation expectations? Conceptually, the ideal experiment would be one in which, for each firm in our sample, we could flip a coin and if heads came up, we would raise their inflation expectations, but if tails came up, we would leave them alone. This would create new variation in inflation expectations that would be systematically unrelated to firms' characteristics. Our first draft of the paper included a simple experiment like this, in which some randomly selected firms had been provided with information about inflation or the central bank's inflation target. We used it to study how firms' expectations responded to the new information. Other papers had applied similar strategies to households (Armantier et al. 2016, Cavallo et al. 2017), but this by itself did not establish that those expectations were important for firms' actions in any way.

To address the editor's question, we had to take it one step further and determine whether changes in expectations actually had an effect on firms' actions. This would require an additional ingredient, namely, tracking the decisions of firms over time. If we could show that those firms whose inflation expectations were randomly changed subsequently set their prices or employment in a manner that differed from firms whose expectations were left untouched, this would provide clear causal evidence that those expectations mattered for their decisions. So we designed a two-stage experiment in which we would first run an RCT to change the inflation expectations of a randomly selected subset of firms. In the second stage, we would track their decisions

over time. Through this combination, we would be able to estimate the degree to which changes in inflation expectations causally passed through into firms' decisions.

Excited at having this plan of attack, we went to our coauthor in New Zealand and proposed that we run two new survey waves: one to implement the RCT in which we would try to change some firms' inflation expectations and a subsequent one to enable us to assess whether their decisions had been altered relative to other firms whose expectations we did not affect. He agreed that this strategy might work but raised what turned out to be the first of many issues we've since learned can make this approach more challenging than how it first sounded: the cost! His funding had run out and implementing two more waves (and these needed to be large surveys to have enough power to detect effects on beliefs and decisions) would cost around $100,000, an amount that seemed insurmountably high. But after some time and the amazing support of our home universities (Go Bears! and Hook 'em Horns!) and the National Science Foundation, we were able to move forward and test out our new empirical strategy.

To be honest, we were not 100% confident that it would work. There were so many ways things could go wrong. The provision of information could fail to change the inflation expectations of firms in the treated group. The sample sizes could be too small to have a detectable effect on decisions even if there was one. Self-reported outcomes could be too imprecise. And, in our heart of hearts as trained macroeconomists, we shared some of the skepticism of our editor: if these expectations were important enough to actually matter for firms' decisions, wouldn't they have chosen to be well informed about inflation in the first place? Didn't the fact that they let themselves be uninformed about inflation imply that these beliefs did not matter? To cover our anxiety as we waited out the long months before the completion of the experiment, we would joke among ourselves about how this might turn out to be the most expensive failure to reject the null hypothesis of no effect ever, a little econ humor that provided us with little comfort.

Finally, the data came in and, to our great relief, the experiment had worked! We found that the information treatment was successful in changing the inflation expectations of firms that were initially uninformed, and that the change in their inflation expectations led them to alter their decisions over subsequent months. Firms that raised their inflation expectations due to the treatment (and relative to the control group) ended up hiring more workers and doing more investment than they had initially planned. And the implied

pass-through of expectations into decisions was economically large along some margins: every 1 percentage point increase in inflation expectations raised employment growth by 2 percentage points and investment growth by 3 percentage points, for example. In contrast, the pass-through into prices was economically small and we could not reject the possibility that it was zero.

With this new causal evidence in hand showing the relevance of firms' inflation expectations for at least some of their decisions, we resubmitted the paper and it was ultimately accepted. Thanks to the skepticism of this editor, we ran our first RCT linking expectations and decisions, the first of what would turn out to be many. Since then, we have had the incredible good fortune to work with different research teams around the world who had access to surveys in which RCTs could be implemented. In addition to our initial surveys of firms in New Zealand with Saten Kumar, we have been able to work with surveys of firms in the U.S. with Brent Meyer, in Uruguay with Serafin Frache and Rodrigo Lluberas, in Italy with Tiziano Ropele and in France with Erwan Gautier and Frederique Savignac. We also began running similar experiments on surveys of households, initially in the U.S. with Michael Weber but then also in the Netherlands with Maarten van Rooij, and in the Euro area with Dimitris Georgarakos and Geoff Kenny. Through these research partnerships as well as through the work of our current and former graduate students and many other research teams working on similar topics and methods in parallel, we have learned much about the challenges and opportunities that come with using RCTs to study how beliefs affect economic decisions.

This book is the fruit of this research and is meant to provide guidance to current and future economists who may be interested in what this approach can bring to the table. Some of this guidance is simply a list of things not to do, lessons we have learned the hard way over time. Some of it is practical steps to raise the probability of success. Throughout, our goal is to help make this approach a more standard part of macroeconomists' toolkit, though the method is much broader and not restricted to macroeconomic questions. Since expectations affect decisions for households, firms, traders, and many other actors/participants in the economy, possible (and existing) applications include industrial organization, labor economics, public finance, and political economy among other fields.

We've divided the book into four parts. The first three focus on methodology, while the fourth describes different applications in the literature. The first part of the book discusses surveys of expectations, issues related to the design of questions to measure expectations, and some of the other challenges that

arise in running surveys of expectations. The literature on expectations surveys has exploded in recent years and we provide an overview of some of the main empirical results to have come out of this literature, although due to space limitations we cannot do justice here to the breadth of this literature. Instead, we focus primarily on inflation expectations, both because of their particular relevance to macroeconomic policymaking and because they have been at the forefront of macroeconomic surveys and research. Finally, we discuss how these measures of expectations can be used to guide and discipline theoretical models of expectations formation.

The second part of the book turns to the RCT component of the strategy. This is the stage at which information is provided to randomly selected subsets of participants to assess whether and how it affects their economic expectations. We focus on how to implement these treatments and under what conditions we should expect them to be successful. Information treatments are not guaranteed to work, and indeed there are many cases where we should *not* expect them to work. Understanding ahead of time which treatments will likely work (and when) has been an important lesson from the very rapidly growing literature on information treatments in macroeconomics. Since the information treatments generate the exogenous variation in expectations that form the basis for identification, this is an essential component of the empirical strategy. We describe some of the main lessons from this line of research, for both policymaking and communication, as well as how results from information treatments can speak to theoretical models of learning and expectations formation.

The third part of the book describes how to estimate the passthrough of expectations into decisions. Ultimately, this is the question that we are interested in answering with our empirical strategy. In addition to the variation in beliefs created via the RCT, we need information on the decisions of economic agents to be able to relate changes in beliefs to decisions. We describe the different ways that one can try to measure outcomes, from the ideal scenario of matching the survey to external administrative datasets to the more common case where outcomes are measured using follow-up survey waves or even planned decisions from the same wave as that of the RCT. One chapter discusses different econometric approaches that can be used to estimate passthrough coefficients, as well as the many issues that can arise in both the implementation and the interpretation of estimated pass-through coefficients. We also discuss in a separate chapter how scenario-based questions (hypotheticals) can be used to identify the same theoretical object as the full-blown RCT

strategy, but more quickly and at significantly lower cost. Finally, another chapter delves into how these estimates can be mapped into theoretical models, and the potential value of doing so.

Through all three of these initial parts of the book, we focus primarily on applications and examples involving inflation expectations to illustrate each of the components within a common framework. This reflects the fact that much of the literature on surveys, information treatments and estimated pass-throughs has been focused on inflation expectations. But the role of beliefs in affecting economic decisions is much broader than this specific application.

The final part of the book therefore covers other topics to which this approach has been applied. One such application is how uncertainty affects decision-making, an important line of research following the seminal work of Bloom (2009). We discuss how one can use RCTs to create exogenous variation in the economic uncertainty perceived by firms or households to study the different channels through which uncertainty affects their behavior, an essential ingredient for understanding the aggregate effects of uncertainty. In another application we consider how peer effects matter for individual decisions. "Keeping up with the Joneses" forces have long been emphasized for spending, but relative income standing could matter for other consumer decisions as well. We review recent papers, especially those relying on this type of RCT approach, that assess how perceived relative income standings affect spending decisions, labor supply decisions, investment decisions, political preferences and even the happiness of individuals in a causal framework. The final two applications are on how individuals' beliefs about future financial asset returns and future housing prices affect their decisions, respectively. In each case, recent research has utilized RCTs to generate variation in expectations of the future asset price changes and then related these changing beliefs to households' decisions, not only for their investment choices but also along other margins such as spending. Across the board, expectations are found to matter for economic decisions, although the quantitative magnitudes of the estimated effects in some cases are large whereas in other cases they are quite small. In all cases, the estimated pass-throughs can be used to discipline economic models in which these expectations channels are at work, and may also relate directly to policy communication objectives and tools.

Throughout, our primary objective is to present the main building blocks of this empirical strategy and identify potential pitfalls. As a result, much of the book concentrates on empirical topics and issues. But because another objective is to encourage macroeconomists in particular to adopt the strategy

as part of their toolbox, we devote significant space to discussing how different steps of the approach can be mapped to theoretical models and the potential usefulness of doing so. There have been a few attempts to do so, and what we present should be thought of only as stepping stones to more systematically integrating the results from RCT studies into macroeconomic theory and modeling, an area we hope to see blossom in coming years. For non-macroeconomist readers, the theory chapters (Chapters 4, 8 and 12) may be of less relevance, and could be skipped without the reader losing understanding of the empirical methods.

Ultimately, this book is meant for students and future economists. Macroeconomics has changed dramatically in the twenty years since we were in graduate school ourselves. The field has become much more empirically driven than it has ever been and this trend is unlikely to reverse itself anytime soon. We view the methods in this book as an additional powerful set of tools that can be used to strengthen the empirical foundations of macroeconomics and guide the development of future models. We hope that students and practitioners of the field will embrace them and build upon them rapidly, hopefully making this book quickly outdated, as good science should do.

INDEX

Active control, 101–102, 168, 287
Adaptive expectations, 2, 13, 83
Administrative data, 42, 46, 101, 155–158, 161–163, 170, 256–257, 265, 292–293
Anchored expectations, 67–69, 198, 201
Anchoring effects, 95–96
Asset price expectations, 237–255
Average inflation targeting, 42, 119, 295

Bank run, 130–131, 250–254
Bayesian learning, 133–135, 267
Business cycles, 70, 74, 81, 100–101, 192–197, 221–222, 298
Business Inflation Expectations (BIE) Survey, 13, 15, 25, 46, 116

Calvo pricing, 83, 201–210
COVID-19 pandemic, 57, 62, 131–132, 136, 161, 190, 211
Conditional expectations, 20–23
Consumer Expectations Survey (CES), 41, 116, 159, 224, 228, 248
Cryptocurrency, 237–249

Decision plans, 8, 161–166, 171–173, 228, 232, 286
Demand effects, 95–96
Diagnostic expectations, 76–78
Disagreement in expectations, 31, 52–59, 68, 80, 105, 112, 146, 197, 206–217, 223, 239

Elicitation of subjective probability distributions, 17–23
External validity, 103, 114–117, 294
Extreme values, 113–114

Federal Reserve, 4–5, 94, 99, 104, 106, 119–121, 130–132, 142, 169, 250, 253, 295
Fixed cost of acquiring information, 72, 138–142, 145
Forecast horizon, 23–24, 48, 60, 68, 71, 74, 99, 202, 276, 279
Forward guidance, 58, 99, 122
Full-information rational expectations (FIRE), 2–3, 13, 48, 70–71, 83, 119, 297

General equilibrium, 191, 203, 226, 294–295
Great Depression, 56, 239
Great Recession, 84, 101

Happiness, 262–264, 272–273
Higher-order beliefs, 66–67, 146–150
Home price expectations, 275–291
Huber robust regression, 114, 127, 182
Hypothetical scenario questions, 160, 174–175, 177, 187–199, 235, 246, 254, 295–296

Identification, 3–4, 8, 34, 173–174, 177, 225–227, 243, 256, 264–265
Ifo Business Survey, 43
Indirect Consumer Inflation Expectations Survey (ICIE), 40

Inequality, 101, 268–269
Inflation targeting, 42, 119, 295–296
Information treatment, 91–102, 295
Instrumental variables, 170–171, 182, 214n3, 268

Kalman gain, 73, 77, 134, 149

Livingston Survey of Firms, 42–43

Manager, 22, 27–28, 48–50, 55, 62, 65–68, 122, 147–148, 157, 184, 222–224
Marginal cost, 46, 84, 139, 203–205, 208
Measurement error, 26, 32, 78, 109–110, 163–166, 181–182, 274
Menu cost of price adjustment, 205
Michigan Survey of Consumers (MSC), 14, 17, 24–26, 32–34, 35, 159, 223
Misallocation, 206–216
Multiple hypothesis testing, 184–185

Narratives, 31, 57–58, 124
Natural experiment, 91, 129, 223, 297
Nielsen Homescan Panel, 40, 50–53, 108–111, 115–116, 238–239, 247, 286–289
Noisy information, 72–79, 150

Online survey, 49–51, 96, 113, 181
Overreaction of expectations, 77–79

Panel conditioning, 33
Passive control group, 101–102
Peer effects, 256–274
Perceptions, 55, 159, 250, 258, 268, 273, 276
Phillips curve, 81–85, 200–206
Placebo, 4, 95–96, 101, 243
Policy communication, 102, 113, 118–132, 175, 216
Portfolio allocations, 50, 156–157, 160, 237–248
Predictability of forecast errors, 59–62, 73–75, 77–80
Preregistration, 184–185
Priming, 24–25, 289
Profit, 28, 157, 169–170, 185, 191–192, 226, 236

Rational inattention, 83, 137–146
Reduced-form regression, 167–170
Rounding in survey responses, 25–26, 163, 224

Self-reported data, 6, 28, 159–166, 181–184, 274, 292
Speeders, 50n2, 113
Sticky information, 71–76, 83
Survey of Consumer Expectations (SCE), 15, 18–19, 33, 35, 107–108, 159, 197, 224, 275–276
Survey of Firms' Inflation Expectations (SoFIE), 15, 43–44, 207
Survey of Inflation and Growth Expectations (SIGE), 47, 157, 160, 169, 184, 211–216
Survey on the Access to Finance of Enterprises (SAFE), 157, 160–161
Surveys of Professional Forecasters, 14, 48–49, 65

Total effect, 175–177, 188–189
Total factor productivity (TFP), 210, 216
Triangular questions, 19, 224, 228

UK Decision Makers Panel (DMP), 43, 165
Uncertainty, 17–20, 26, 62–66, 68, 99, 102, 133–146, 176–179, 190–195, 221–236
Underreaction of expectations, 62, 73, 78–80, 87
Unemployment, 22–23, 30, 52–53, 55, 58, 81–85, 92, 99–101, 122, 125, 176, 193–194, 197–198

Wage expectations, 46
Wage-price spiral, 2
Willingness to pay, 192–197
Worker, 2, 6, 81, 157, 201, 208, 269–270