Chapter 1

Why Study Price Stickiness? Why This Way?

Nothing astonishes men so much as common sense
and plain dealing.
—RALPH WALDO EMERSON

The Importance of Price Stickiness

In recent decades, macroeconomic theorists have devoted enormous amounts of time, thought, and energy to the search for better microtheoretic foundations for macroeconomic behavior. Nowhere has this search borne less fruit than in seeking answers to the following question: Why do nominal wages and prices react so slowly to business cycle developments? In short, why are wages and prices so "sticky"? The abject failure of the standard research methodology to make headway on this critical issue in the microfoundations of macroeconomics motivated the unorthodox approach of the present study.

No one should think the question unimportant. On the contrary, sticky prices are an essential element of Keynesian economics, which is sometimes called the economics of nominal rigidities. (It is called less polite things as well.) The sobriquet is an exaggeration, to be sure, but a forgivable one, for embedding the assumption of short-run price or wage rigidity into almost any macro model will make it produce characteristically Keynesian results—such as that an injection of money raises production. The simplest illustration is the quantity theory of money with fixed velocity:

$$MV = Py.$$

In the absence of nominal rigidities, real output, y, is essentially fixed on the supply side of the economy, so that changes in money must pass directly and proportionately into prices. But if P, the price level, is sticky in the short run, the very same equation implies that part of any change in M must first show up in y. Conversely, if a vertical aggregate supply curve (attributable, say, to instantaneous market clearing) is appended to an otherwise "Keynesian" IS-LM model, the real effects of fiscal and monetary policy disappear.

When a scientific discipline knows next to nothing about a question of paramount importance, it is in some trouble. How did macroeconomics get into such a predicament? Two obvious explanations can be dismissed immediately. First, it is not because macroeconomists have just discovered that wage-price stickiness is a central issue; we have known this since Keynes's *General Theory*, if not before. Second, the failure does not result from lack of effort. Scores, if not hundreds, of theorists have worked on this problem, producing many interesting theoretical explanations; and new ideas keep popping up all the time.[1] Progress has not been hampered by lack of imagination.

Nor, by the way, has it been hampered by lack of observation. Although much time and energy was wasted in the 1970s and 1980s arguing over whether or not the economy should be modeled as a giant auction hall with perfectly flexible prices, a small mountain of empirical evidence testifies to the fact that wages and prices adjust slowly to macroeconomic events. For example, the economist Robert Gordon (1990) summarizes the evidence that aggregate price indexes move sluggishly, while Yoram Weiss (1993) surveys some papers that provide similar evidence for individual prices. The tricky questions are two: How slow is slow? And what factors account for the sluggishness? This book is devoted to the second of these two questions. But a brief word on the first is in order, for it helps explain why conventional methods of economic inquiry—theory and econometrics—have yielded such meager results.

Why So Little Progress So Far?

When we say that wages or prices are "sticky," we generally mean that they move more slowly than would Walrasian market-clearing prices. Two curmudgeonly questions arise, each of which influenced the design of this study: Is the statement operational? And is it something we should care about? We take these up in reverse order.

Is Wage-Price Stickiness an Important Phenomenon?

The question here is basically whether wage-price stickiness has allocative significance. For example, if no one ever borrowed at the credit card interest rate, which remained around 19 percent for years, then the fact that this rate was extremely sticky would hardly have mattered.[2] Ever since the economist Robert Barro's (1977) ingenious paper, macroeconomists have worried that the sticky wages we see in many labor markets may in fact lack allocative significance. Specifically, firms may not equate the real wage to the marginal product of labor and workers may not equate the real wage to the the marginal utility of leisure. Why? Because, it is argued, employees and employers implicitly enter into long-term agreements to exchange labor for money. These contracts clear the labor market in a long-run sense, thereby tying long-run average labor supply closely to long-run average wage payments. But if workers dislike wage variability, the contract may pay steady wages month after month even if both the marginal product of labor and actual hours of work vary considerably over time.

This argument has persuaded many economists that it is hazardous to conclude from the observed stickiness of wages that there is pervasive disequilibrium in labor markets. It also calls into question the allocative significance of the limited variability of real wages over the business cycle. But is the argument empirically important? Do firms and workers actually enter into such agreements? No one really knows. When Alan Blinder and his student Don Choi (1990) asked a small sample of personnel managers what they thought of the implicit contract theory, the results were a bottle half full and half empty.[3] About half thought it "plausible or relevant," the rest did not.[4]

Irrespective of its empirical relevance, the theory has profoundly influenced the thinking of academic economists and has shifted the focus of research from wage stickiness to price stickiness. When Blinder initiated this research, he knew he would be preaching to the unconverted and was anxious to have at least some economists pay at least some attention to the unorthodox research findings. So he decided to study only price stickiness. Of course, the implicit contracts argument can be applied to prices, too. Sticky prices may just be installment payments on long-term agreements, rather than symptoms of non-clearing product markets. But everyone seems to agree that, however important or unimportant the implicit contract theory is in labor markets, it must surely be less important in product markets, where arms'-length, spot transactions are much more common.

Is Wage-Price Stickiness an Operational Concept?

The next question is whether and how we can breathe empirical life into the theoretical notion of price stickiness. In brief, how slow is slow? To state the issue perhaps a bit too boldly, a theory that predicts that prices adjust more slowly than market-clearing prices—*and nothing else*—is basically untestable, and therefore an empty theory. Why? Because economists have no agreed-upon metric to use in assessing the observed speed of adjustment of any particular price, let alone the aggregate price level. So if we find, say, that the price of candy bars changes every six months, how do we know whether this is slower or faster than the Walrasian norm?[5]

If a theory makes no prediction other than that prices move less rapidly than Walrasian prices, econometric testing is almost (but not quite) out of the question. To conduct a test, a complete model of supply, demand, and price adjustment must be specified, estimated, and used to derive a quantitative measure of the speed at which the market-clearing price moves. Then actual price movements can be compared to this norm. This research strategy is not a counsel of perfection; it can be implemented. Indeed, it is one of the ways that econometricians have demonstrated that prices and wages are sticky. But, of course, any such demonstration is conditional on the validity of the many maintained hypotheses used as the framework for estimation. So any such finding is open to dispute.

But the problem goes deeper than this. If we have a wide variety of models, each of which predicts that prices are sticky and nothing else, conventional econometrics will have a hard (if not impossible) time distinguishing among them, for there is no way to test one theory against another. This problem, we think, is the main reason why formal econometrics has made so little progress in weeding out invalid theories of wage-price stickiness. It is one of the two problems that originally drove Blinder—in desperation!—to the interview method.

The other problem is that several of the most prominent theories of price stickiness rely on variables that are either unobservable in principle or unobserved in practice.

One example is a theory that achieved wide popularity in the 1980s, which holds that firms hesitate to cut prices in slumps out of fear that customers will misinterpret any price cut as a reduction in quality—*when in fact there has been no such quality reduction*.[6] Notice that unobservability is crucial to the argument; if quality were easily observed, there would be no possibility of misinterpretation because everyone would recognize when quality had changed.

Another example is the "menu cost" theory, which says that firms change prices infrequently because they incur a fixed cost each time they do so.[7] In principle, such costs can be measured. In practice, however, we have few such measurements and are unlikely to get many.[8] It might seem that this particular theory does at least carry a clear collateral implication—namely, that fixed costs preclude small price changes, where the precise meaning of "small" is defined by the size of the menu costs. In principle, that is correct. But our inability to measure menu costs directly robs this implication of operational significance. Will the firm avoid 1 percent, 5 percent, or 10 percent price changes?

When theories rely on unobservables in essential ways, econometric testing is difficult, to say the least. Thus it is no accident that new theories of price stickiness have continued to proliferate faster than applied econometrics has been able to discard old ones. It was that unsatisfactory state of affairs that first set Blinder thinking about an alternative approach.

Time for a New Approach?

In pondering this dilemma, a curious "empirical regularity" emerged. Virtually every theory of price stickiness outlines a thought process that allegedly leads decisionmakers (generally modeled as profit maximizers) to conclude that it is against their best interest to change the price. But if people actually think the way one of these theories says, then they should be aware that they do—or so it seemed. Hence an idea: Why not ask them?

This naive idea must be approached with caution. If you confront a price-setter with an open-ended question like, "Why don't you cut your prices more (or more often) when sales sag?" you may get shrugs, blank stares, or incoherent answers. What you hear is unlikely to fit neatly into economists' theoretical boxes. But suppose you ask more pointed questions. Suppose you describe in plain English the chain of reasoning that, according to Theory X, goes through the minds of price-setters. If Theory X really describes their behavior, the decisionmakers ought to recognize and resonate to it. If they do not, then they are probably not behaving as the theory says. At least that was our methodological precept. If the true reasons for price stickiness are buried deep in the subconsciousnesses of decisionmakers, then interviews are unlikely to uncover them.

Here are two examples of the rather pointed way in which we posed questions about the theories.

One very old theory of why prices may be rigid over the business cycle, which enjoyed a strong revival in the 1980s, starts with the premise that profit-maximizing firms with market power set price (P) as a markup over marginal cost (MC), which markup depends on the elasticity of demand ($\epsilon > 1$).[9] Thus:

$$P = MC[\epsilon/(\epsilon - 1)],$$

where ϵ is defined to be a positive number. The theory then asserts that demand curves become less elastic as they shift in so that, even though MC falls as output contracts, the markup rises to compensate. The result may be approximate constancy of the profit-maximizing price over the business cycle.

In principle, this theory can be tested by conventional econometric means; all we need do is measure how the elasticity of demand varies over the business cycle in a variety of industries. In practice, however, any applied econometrician will recognize that as a tall order, unlikely to be filled with the nonexperimental data at our disposal. But now think about using the interview method as an alternative to time-series econometrics. If this theory is the real (or one real) reason for price rigidity, firms must both believe that their demand elasticities are procyclical and act on that belief. In that case, if you ask them about it—eschewing jargon like "elasticity," of course—they ought to recognize the idea and feel comfortable with it.

To test this theory, our questionnaire posed the following plain-English question:[10]

> B5(a). It has been suggested that, when business turns down, a company loses its least loyal customers first and retains its most loyal ones. Since the remaining customers are not very sensitive to price, reducing markups will not stimulate sales very much. Is this idea true in your company?

If the respondent answered yes, we then asked:

> B5. How important is it in explaining the speed of price adjustment in your company?

To preview some results that will be examined in more detail later, almost 60 percent of respondents accepted the premise that the elasticity of demand varies procyclically. But only about half of those (hence, about 30 percent of all firms) rated it a "moderately important" or "very important" source of price stickiness.

Our second example, the Okun (1981) "invisible handshake" theory, is an even more extreme example in that direct econometric testing seems out of the question even on conceptual grounds. The idea behind the theory is that firms have implicit understandings with their regular customers which proscribe price increases in tight markets, presumably in return for stable prices in weak markets. What observable variable can be used to measure the importance, or even the existence, of such implicit agreements? None, we fear—which may be why there has been so little econometric testing to date. But it seems to us that, if such tacit agreements exist, firms ought to know that they do—and should say so when asked.

Our questionnaire therefore "tested" this theory by posing the following question:

> B2(a). Another idea has been suggested for cases in which price increases are not prohibited by explicit contracts. The idea is that firms have implicit understandings with their customers—who expect the firms not to take advantage of the situation by raising prices when the market is tight. Is this idea true in your company?

About two-thirds answered yes. Furthermore, a large majority of those answering yes rated implicit contracts a "moderately important" or "very important" reason why prices adjust slowly in their companies. The theory evidently holds promise in certain sectors of the economy.

But Aren't Interviews Unreliable?

Economists are disposed to be skeptical that you can learn anything about economic behavior by asking people. Most believe you should not even try. Instead, you should observe what they *do* in markets (not what they say), model that behavior theoretically, and test the model econometrically.

The litany of objections to interviews is not without merit. Critics argue that responses may be terribly sensitive to the precise wording of the questions. We agree and hence devoted many hours to the form and structure of the questionnaire. While we do not pretend to have achieved perfection (and will, in fact, mention some problems in subsequent chapters), we invite skeptical readers to inspect the full questionnaire that was used in the field. It is included as appendix A of this book.[11]

Other critics will object that interviewees have no incentive to respond truthfully or thoughtfully, and so may refuse to cooperate or give misleading answers. Where the respondent has reason to conceal the truth or mislead the interviewer, this objection is, to our minds, a show stopper. In such cases, the interview method is simply not a promising mode of inquiry. Thus, for example, interviews may be a poor way to estimate the extent of tax evasion or the prevalence of collusion among businesses. But there are many interesting and important questions about which people have no particular reason to conceal the truth—unless they are pathological liars. In such cases, the interview method might help.

The thoughtfulness problem goes deeper. For example, people may not understand or be able to articulate their own motives or behavior very well. We all know the billiard-ball analogy: A good pool player makes excellent intuitive use of the laws of physics without understanding them intellectually. So if you ask expert players to explain how they shoot so well, they may not give you a coherent answer—and almost certainly will not give an answer that a physicist would mark correct. For this reason, we think, many economists are skeptical that you can learn anything by asking "economic players"—even good ones—about how they play the game.

In part, we agree. We do not, for example, think much is learned by asking corporate executives open-ended questions about the goals of their companies. They may just pick objectives that sound lofty or otherwise appealing. But more pointed questions, posed in plain English, can elicit more useful responses. For example, if you ask skilled billiards players whether they base their shots on the principle that the angle of incidence equals the angle of reflection, they will probably think you strange. But, if you take them to the table and ask about the angles they choose—using elementary physics to tailor your questions—they would probably respond in the affirmative. Closer to home, we should remember that most of our standard data come, in the first instance, from either face-to-face interviews or mailed questionnaires filled out by minor functionaries. How else do you think the Labor Department measures unemployment or the Commerce Department estimates Gross Domestic Product (GDP)?[12]

Thus, while many objections to the interview method have some validity, we should keep them in perspective. Economists above all should evaluate the usefulness of any suggested mode of inquiry—including interviews—by posing the classic question: Relative to what? The imperfect knowledge we pick up from questionnaires should not be compared to some epistemological ideal, but to the imperfect knowledge that nonexperimental scientists can deduce theoretically or glean from econometric investigations.

In doing so, it is important to remember that theory and econometrics also have their limitations, which are often inadequately appreciated. All too often theoretical deductions are untested and/or based on untested premises. Worse yet, either the conclusions or the assumptions may be *untestable*. Econometric evidence is often equivocal and/or subject to methodological dispute. Results may be fragile owing to small samples or multicollinearity. There may have been "regime changes" during the sample period. Appropriate instruments are scarce or nonexistent in time-series applications. And computers make data mining all too easy. Stacked up against competition of this caliber, the interview method may not look so bad after all—especially if viewed as a supplement to, rather than a replacement for, more conventional modes of economic inquiry.

In sum, we are more than willing to accept the methodological precept that economists should rely on the standard tools of inquiry—theory and econometrics—where they bear fruit. In cases where people's observed behavior conflicts with what they say they do, we, like most economists, would stick with the observations.

But we have just argued that conventional research techniques have made virtually no progress in explaining wage-price stickiness. Might we not, therefore, learn something by opening our eyes and ears and listening to the folks who populate the economies we study, the people who actually do the things we theorize about? Yes, it is true that physicists and chemists do not ask their subjects why they behave as they do. But, in our zeal to emulate the hard sciences, economists should not misinterpret that lesson. If molecules could talk, would chemists refuse to listen?

A Bird's Eye View of the Survey Design

Details of how the questionnaire was designed and tested, how the sample was selected, and how the interviews were conducted are provided in chapter 3. Here we offer the reader in a hurry a quick summary of the methodology.

Twelve theories of sticky prices—all but one of them culled from the theoretical literature on the microfoundations of macroeconomics—were selected for testing. The selection process was far from random and not entirely objective. First of all, we took it for granted that almost all firms in our economy (excluding farms) are price makers rather than price takers—an assumption amply justified by the survey responses. Second, the selection of theories for testing reflected Blinder's personal judgments about which of the many theories of price stickiness were most prominent in the academic literature, translatable into plain English, and sensible enough to be explained to business people with a straight face. These choices are justified, to some extent, in chapter 2.

Each theory so selected was turned into a question for part B of the questionnaire; two examples were offered above. Call these the *main questions*. (They are numbered B1 through B12 on the questionnaire.) Each main question was followed by up to ten additional questions tailored to the specifics of each theory. So part B of the questionnaire has twelve sections, one for each theory.[13] Part A requests a variety of factual information about the company, such as its size, how often it changes prices, to whom it sells, whether it has formal contracts, and so on. Even economists who are skeptical that we can learn anything of value by asking business executives their opinions on theories may find the answers to these factual questions enlightening. We certainly did.

Selecting the sample was a delicate task. Our philosophy was simple, though its execution was not. We took the goal to be to explain why the GDP deflator moves sluggishly, and hence sought to create a random sample of the GDP. Well, not quite. The sample was actually limited to the private, unregulated, nonfarm, for-profit GDP. The motives for making these exclusions—which together amount to about 29 percent of GDP—are perhaps obvious; if not, they are explained in chapter 3. That chapter also explains our sampling method in detail—namely, how a computerized "dartboard" was created, with each firm assigned a "slice" proportional to its value added.

Once this was done, a random sample of three hundred and thirty firms was drawn, interviews were requested with each, and two hundred agreed to participate.[14] The stunningly high response rate of 61 percent already suggests that any nonresponse bias is probably small, but some evidence in support of this point is offered in chapter 3. Interviews were conducted in person, generally in the respondent's office. A few were done by Blinder personally, but most were carried out by a team of thirteen Princeton graduate students specially selected and trained for the task. Interviews varied in length, depending mainly on how discursive the respondent was, but generally took forty-five to seventy minutes.

Since the probability that any firm would be selected into the sample was proportional to its value added, all averages and distributions reported in this book are unweighted. The weighting was embodied in the probability of being selected into the sample, so any further weighting of the responses would have amounted to double counting.[15]

Plan of the Book

The book is organized into eighteen chapters, but most of them are blissfully short. Chapter 2 reviews the theoretical, empirical, and survey antecedents to this research.

Chapter 3 offers details on the questionnaire, sample selection, and interview procedures, highlighting some problems that are not mentioned in the brief summary above. It also offers some evidence that there were no serious biases from nonresponse, only minor differences in coding across the fourteen interviewers, few significant effects of interview date (the data collection period included a recession), and almost no differences at all by geographical location.

Chapter 4 summarizes the findings from the "factual" part of the questionnaire (part A). We view these findings as answers to a set of fascinating and important questions that simply cannot be addressed with standard data sources. For example, we offer here what we believe to be the first estimates of the fraction of United States GDP sold under written contracts, the fraction sold to repeat customers, the average lag between a change in demand and the corresponding change in price, and so on. Many of the answers are surprising—both to us and to others.[16]

Chapter 5 provides an overview of the results for the twelve main theories. Which of the theories seem to have the most validity overall and within particular sectors? What do the survey results tell us about specific questions that cut across theories? For example, are prices more sticky downward than upward? Surprisingly, the answer appears to be no. Chapters 6 through 17, most of which are brief, delve into the details of each of the twelve theories. Each chapter explores a theory in detail, gives the survey results on its "popularity," explores which attributes of a firm help explain its affinity for the theory, and discusses the answers to follow-up questions designed to shed light on the validity, applicability, or other aspects of the theory. Because chapter 6 is the first of these chapters, it also includes some methodological discussion that is relevant for the subsequent chapters.

Finally, chapter 18 quickly sums up, assesses what we have and have not learned, points to some policy implications, and offers the inevitable suggestions for future research.

Chapter Summary

This entire book was motivated by a failure—the failure of conventional theoretical and empirical research tools to answer a question of overwhelming importance for macroeconomic theory and policy: Why are prices sticky? One important impediment to progress, we suggest, is that many of the competing theories are epistemologically empty, or nearly so. They may predict nothing other than that prices should move more slowly than some unknown Walrasian norm. Or they may be based on variables which are either unmeasurable in principle or unmeasured in practice.

Is there a way out of this methodological box? Some economists, perhaps most, would answer no. But, since each theory is based on a chain of reasoning that allegedly takes place inside a decision-maker's head, we suggest that an interview approach may hold promise. To be sure, there are many hazards in trying to learn about how economies work by asking participants. But skeptics who raise these (valid) objections to interviews often forget about the many hazards that plague conventional econometric and theoretical research. A pragmatic attitude, not methodological purity, is called for. And since the interview method is virtually uncharted territory, an initial exploration seemed worthwhile—and was made. This book is a report on that expedition.