

## American Economic Association

---

Why are Prices Sticky? Preliminary Results from an Interview Study

Author(s): Alan S. Blinder

Source: *The American Economic Review*, Vol. 81, No. 2, Papers and Proceedings of the Hundred and Third Annual Meeting of the American Economic Association (May, 1991), pp. 89-96

Published by: [American Economic Association](#)

Stable URL: <http://www.jstor.org/stable/2006832>

Accessed: 25/06/2014 10:32

---

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



American Economic Association is collaborating with JSTOR to digitize, preserve and extend access to *The American Economic Review*.

<http://www.jstor.org>

## Why are Prices Sticky? Preliminary Results from an Interview Study

By ALAN S. BLINDER \*

Those of us who teach Keynesian economics (an increasing number these days) attach great importance to the phenomenon of wage and/or price “stickiness.” It explains, for example, why recessions cure themselves only slowly and why changes in the money supply have real effects. In fact, Keynesian economics is sometimes characterized as the economics of nominal rigidities. Yet, more than a half-century after Keynes published *The General Theory*, the phenomenon itself remains poorly understood. Just why are wages or prices sticky?

It is not that economists have ignored these questions. One could literally fill many volumes with good empirical studies of wage and price stickiness, and many more with clever theories purporting to explain these phenomena. Yet, despite all this work, the range of admissible theories is wider than ever, and new theories continue to crop up faster than old ones are rejected. This lack of scientific progress makes one wonder

about the basic research strategy that economists have been pursuing. Is there a better way?

In Section I, I argue that traditional research strategies may be unpromising vehicles for learning about why prices are sticky; the time may have come to entertain new and unorthodox approaches. Section II defends the notation that interviews, in particular, can be a useful research tool in this area. Sections III and IV then describe the design of and summarize some preliminary findings from a large-scale interview study currently underway at Princeton. This paper is very much a progress report since the interviewing still has 9–12 months to go.

### I. Why Interviews?

The standard program of scientific research in economics is to (a) develop a theory of some phenomenon, (b) formulate it in econometric terms, and then (c) test it with actual data. The theory is then either rejected or provisionally accepted, that is, allowed to survive to the next test. Unfortunately, this program has been singularly unsuccessful in the area of wage-price stickiness. Most economists would, I think, agree that we know next to nothing about which of several dozen theories of wage-price stickiness are valid and which are not. We might have expected statistical tests to have weeded out the weaker theories by now, especially since many of them have been around a long time. But, in fact, the survivorship rate rivals that of congressional elections. Why?

Part of the reason, no doubt, is that economists are fonder of building theories than of testing them. But I think the main reason is that most of the theories are *empty*

\*Department of Economics, Princeton University, Princeton, NJ 08544. I owe many debts of gratitude: to the Russell Sage and Alfred P. Sloan Foundations for financial support; to Eric Wanner for persuading me to go ahead with the project; to Elie Canetti for extensive and excellent research assistance; to Phyllis Durepos for efficiently handling the considerable volume of office traffic; to my graduate student interviewers, Anthony Marcus, David Genesove, Katy Graddy, Dean Jolliffe, Harold Kim, John Leahy, Alec Levenson, John Penrod, Michael Quinn, Steve Schwartz, Tim Vogelsang, and David Zimmerman; to Bob Abelson, George Akerlof, Larry Ball, Olivier Blanchard, Elizabeth Bogan, Dennis Carlton, Russell Cooper, Bob Gordon, Bruce Greenwald, Glenn Hubbard, Tom Juster, Danny Kahneman, Anil Kayshap, Alan Krueger, David Romer, Julio Rotemberg, Bob Shiller, and Andy Weiss for helpful suggestions; and, especially, to the many business executives who took time out from their busy schedules to answer our questions.

in the following specific sense: Either they involve unobservable variables in an essential way, or they carry no real implications other than that prices are “sluggish” in some unmeasurable sense, or both. This makes econometric modeling a blunt, perhaps even useless, investigative tool. Let me illustrate.

To begin with, think about what we mean when we say that a theory predicts that prices are “sticky”: Often nothing more than that prices adjust less rapidly than Walrasian market-clearing prices. But since equilibrium price movements often go unmeasured, it is hard to know whether actual prices are moving faster or slower than this norm. More important, all the theories share *exactly the same* prediction: that prices are sticky in this sense. So how are we to discriminate among them? It seems difficult even to imagine what a decisive test would look like, much less to carry one out.

A natural idea is to use each theory to derive other, auxiliary predictions, and then test these. Unfortunately, often there are no such predictions—or at least none that can be checked against actual data. One reason is that many of the theories are based on variables that are *unobservable* either in principle or in practice.

As an example of the former, consider the theory that firms hesitate to cut prices because they fear that customers will interpret a price cut as signal that quality has been reduced—*when, in fact, there has been no quality reduction*. Clearly, the theory is predicated on the existence of *unobservable* quality differences for, if consumers could observe quality readily, price would not play a signalling role. As an example of the latter, consider the so-called menu cost theory. In principle, fixed costs of changing prices can be observed and measured. In practice, such costs take disparate forms in different firms, and we have no data on their magnitude. So the theory can be tested at best indirectly, at worst not at all.<sup>1</sup>

<sup>1</sup>For example, the menu cost theory predicts firms will never make “small” price changes, which appears to be a testable proposition. But does it rule out price increases of 1 percent or 10 percent? Without measuring the menu costs, we cannot answer this question.

One might argue that theories that make identical predictions are, in an operational sense, the same theory; so it does not matter which theory is correct. In fact, however, it does matter. It matters for the value of economics as a descriptive science; and it matters for the conduct of macroeconomic policy because not all sources of price rigidity open the door to welfare-improving policies.

If it really matters which theory is correct, but conventional modes of economic inquiry cannot adjudicate the dispute, economic science would appear to be in deep trouble. Fortunately, one other common characteristic of the theories suggests a way out: Virtually every theory of price rigidity *describes a chain of reasoning which allegedly leads the firm to conclude that a change in price is inadvisable*.

That is what gave me the idea for an interview study. If actual decision makers really think the way one of these theories says, they ought to know that they do. If you ask them open-ended questions like, “Why don’t you cut your prices when sales decline?,” you may get shrugs or incoherent answers. But, if you confront them with the chain of reasoning they actually follow, they ought to recognize and agree with it. Conversely, if they explicitly deny the relevance or validity of a particular argument, then it is probably not governing their behavior. At least that is my methodological precept.

## II. But Aren’t Interviews Unreliable?

Economists are skeptical that you can learn much by asking people. We are trained to study behavior by watching what people *do* (usually in markets), not by listening to what they *say*. For example, critics will point out that subjects of interviews have no incentive to respond truthfully or thoughtfully; so *homo economicus* might refuse to cooperate or even give misleading answers. If the respondent has reason to conceal the truth or mislead the interviewer, this objection is probably a show stopper. But, in the case of price stickiness, people have no particular reason to conceal the truth. As long as people are not pathological liars, interviews may elicit useful information.

The thoughtfulness problem goes deeper. We all know the billiard ball analogy: A good pool player makes excellent use of the laws of physics without understanding them, and certainly without being able to articulate them. For this reason, many economists doubt that much can be learned by asking "economic players" about how they play the game. But I believe that more pointed questions, posed in plain English, can elicit useful answers. For example, if you ask a skilled billiards player whether he bases his shots on the principle that the angle of incidence equals the angle of reflection, he will probably look at you quizzically. But, if you take him to the table and, pointing to the proper angles, ask: "Do you try to make this angle the same as that angle?," I imagine he would respond in the affirmative.

Legitimate questions can also be raised about the size and representativeness of interview samples. Detailed case studies of two or three companies can, at most, provide anecdotes, not useful statistical information. And samples that are unrepresentative of the underlying population give no basis for drawing inferences about population statistics. But these are familiar problems, well known to any user of data. They are not reasons to reject interview evidence out of hand.

We should remember that theory and econometrics have their limitations, too. Theoretical deductions are often untested or, worse yet, *untestable*. Econometric evidence is often equivocal and/or subject to methodological dispute. The imperfect knowledge we can pick up from interviews and questionnaires should therefore not be compared to some epistemological ideal, but to the imperfect knowledge that nonexperimental scientists can deduce theoretically or glean from econometric studies. By this more reasonable standard of evidence, data culled from interviews certainly look admissible.

### III. The Research Design

Along with a team of Princeton graduate students, I have now been in the business of interviewing executives about their pricing strategies since August 1988. Let me briefly

describe the research strategy and some of the tactics.

*Designing the Questionnaire:* The first step, obviously, was to decide which theories were to be tested and then to turn them into questions that ordinary people can understand. This selection process was not entirely objective; nor could it have been. In scouring the intellectual waterfront, I excluded theories that sounded silly or that actually *were* silly, or that seemed too complicated to be explained tersely in plain English. These omissions were not particularly costly; they did not disqualify any of the major competing theories.

Translating the equations and diagrams into English proved to be easy. My first stab at a questionnaire was tried out on many "guinea pigs" (including economists, business people, and a few scholars involved in survey research) and altered in literally hundreds of ways. But, in general, the current version of the questionnaire bears a striking resemblance to my original draft. Translating from technical journalese into English simply proved not to be very difficult. I took this to be a good omen.

However, one important caveat should be entered. Not infrequently, a question must be rephrased on the spot to make the respondent understand it. To do so successfully, an interviewer must be reasonably articulate, must be able to think on his or her feet, and, most importantly, must understand the economics well enough to paraphrase a question without changing its meaning. Thus professional pollsters will not do. Instead, my interviewers are carefully selected Princeton graduate students.

*Eliciting Cooperation from Businesses:* The next problem was getting our feet in the door and getting our questions answered by the right people. A low-response rate would obviously raise fears of selectivity bias.

A small-scale pilot study was used both to estimate the likely response rate in a large-scale study and to polish our questionnaire and interviewing techniques. We randomly selected 16 firms in the northeastern United States and wrote each an introductory letter requesting an interview. We then followed up with phone calls and/or further mail-

ings as necessary and, after considerable effort, successfully interviewed 8 companies. (These 8 companies are *not* included in the data reported below.) The estimated 50 percent response rate struck me as high enough to merit proceeding to the full study. So far, the response rate in the full study is running above 60 percent.

In addition, we had no trouble making connections with the person or persons in each company who could answer our questions. The number of "I don't know" responses has been small, and almost never do we hear, "You're asking the wrong person." Once in the door, we have found people more than willing to talk frankly. They find the questions interesting, mostly understandable, and not invasive of privacy.<sup>2</sup> Our experience, in fact, is that people are often eager to tell us things about which we would never have dared ask for fear that we would appear to be prying. We do, of course, promise confidentiality.

*Getting Economists to Pay Attention:* The hardest problem may yet lie ahead: Getting economists to pay attention to the results. Several concrete steps have been taken to increase the believability of the ultimate findings, whatever they may be. First, I wrote letters to the originators or major proponents of each theory. The letters explained the nature of the study, included a copy of the relevant portion of the questionnaire, asked for suggested improvements in the questions, and asked the "theorists" to suggest other testable implications of their theories. Most responded, sometimes offering useful ideas for modifying the questionnaire, which were adopted. Interestingly, however, not one person suggested a single further implication that could be tested in the questionnaire. This underscores the point I made earlier about the theories being empty.

Second, and much more important, we have taken great pains to ensure that the sample of firms is (a) large enough for seri-

ous statistical analysis, (b) randomly selected, and (c) representative of the private, for-profit GNP. To my knowledge, this is the first time anything like this has ever been attempted in an interview study. Since it is what makes the study unique, it is worth describing in some detail.

We purchased a tape listing firms with annual sales over \$10 million in the northeastern United States, an area which accounts for 40–45 percent of U.S. GNP. From this sampling universe, we eliminated government enterprises and nonprofits on the grounds that the theories we wanted to test were all about profit-seeking firms; this left almost 25,000 firms. We then assigned a sampling weight to each firm proportional to its value-added and drew a random sample of 400 firms, seeking to complete 200 interviews. (Remember, the estimated response rate was 50 percent.)

Notice that we excluded any company with annual sales under \$10 million, even though they employ nearly half the labor force. Why? Because there are so many of them, and the expense of reaching any sizable portion would be prohibitive. Clearly, an optimal experimental design would balance the value of the information obtained against the costs of obtaining it and would, therefore, undersample very small companies.<sup>3</sup> We approximated this crudely by not sampling any companies below the \$10 million threshold. Analogous cost considerations motivated the geographical truncation: costs depend on distance from Princeton, NJ; benefits do not.

These two exclusions, however, might compromise the representativeness of the sample. So we remedied them as follows. Suppose pricing behavior varies geographically only because industrial structure differs across states, not because, say, California firms are inherently different from New Jersey firms. Then we can (and did) create a *synthetic national sample* by reweighting each firm in the Northeast to reflect na-

<sup>2</sup>Considerable effort was expended to guarantee the latter. For example, I abjured any questions that might conjure up images of the U.S. Department of Justice.

<sup>3</sup>This presumes that the study aims to explain the behavior of the GNP deflator.

tional, rather than regional, shares in value-added. Firms in industries that are overrepresented in the Northeast, like banking, were appropriately undersampled; and firms in industries underrepresented in the region, like gas and oil extraction, were correspondingly oversampled. Thus, the industrial structure of our sampling frame matches that of the nation as a whole.

While we were adjusting sampling weights, it was a simple matter to reduce the potential bias from omitting small firms. For example, any sample that excludes firms with annual sales under \$10 million will underweight retailing. We simply oversampled the remaining retail firms enough to assign retailing its proper national weight—and similarly for every other two-digit industry. In this way, we eliminate from our sample the portion of any “large-firm bias” that stems from the different industrial structures of small versus large firms. However, to the extent that small firms really differ from large firms *in the same industry*, the bias remains.

#### IV. Preliminary Results

The interviewing is now in progress. Given the limited clerical and interviewer resources at my disposal, we are conducting them in eight waves of approximately 50 letters (25 interviews) each over a period of about 18 months. As of this writing, the first three waves are nearly complete and we have just begun the fourth.

The results reported here are based on tabulations of the 72 interviews completed as of mid-November 1990. Since this partial sample is too small to permit much useful disaggregation, I restrict my attention to national averages here. Once the full sample of 200 is available, however, there will be a great deal of information about, for example, how the validity of different theories varies by industry.

The questionnaire comes in two parts and usually takes about 45–60 minutes to administer. The first part gathers a variety of factual information about each firm, while the second inquires directly about the theories. The factual data will eventually be

TABLE 1—MEASURES OF PRICE STICKINESS<sup>a</sup>

A. Frequency of Price Change <sup>b</sup>		Percent of Companies
More than 12		10.1
4.01 to 12		4.3
2.01 to 4		10.1
1.01 to 2		20.3
1.0		37.7
Less than 1.0		17.4

  

B. <i>Question:</i> How much time normally elapses after a significant _____ before you raise (reduce) your prices?		
Event	Mean Lag <sup>c</sup>	Standard Deviation
Increase in demand	3.23	2.93
Decrease in demand	3.60	3.93
Increase in costs	3.17	2.90
Decrease in costs	3.97	4.47

<sup>a</sup>Based on 72 interviews completed as of mid-November 1990.

<sup>b</sup>Times per year.

<sup>c</sup>In months.

used for disaggregation, in cross tabulations, and as right-hand variables in regressions explaining the validity of the various theories. I make little use of these data in this paper, but one issue is important enough to deal with right now: Just how “sticky” are prices in the U.S. economy?

To find out, we ask each firm, “How often do the prices of your most important products change in a typical year?” The distribution is shown in panel A of Table 1. These results offer cheer for Keynesians who model prices as, say, rigid for one period or more. I like to joke that in macroeconomic theory we know that the length of “the period” is one quarter. According to these data, less than 15 percent of the GNP is repriced more frequently than quarterly, and fully 55 percent is repriced no more often than once per year. If we had to pick a single abstraction to represent the whole economy, annual price change (the median response) would appear to be the “right” model.

From the point of view of macroeconomic theory, frequency of price change may not be the right question to ask, for it depends as much on the frequency of shocks as on

firms' pricing strategies. We are more interested to know how long price adjustments lag behind shocks to demand and cost. Panel B summarizes the results from a series of questions of the form, "How much time normally elapses after a significant \_\_\_\_\_ before you raise [reduce] your prices?" We inquire about four events: an increase in demand, a decrease in demand, an increase in cost, and a decrease in cost. The table shows that, while the distributions are quite spread out, the mean lags cluster in the 3–4 month range. There is precious little evidence that prices increase faster than they decrease, and virtually none that firms respond to cost shocks more quickly than to demand shocks.

The bulk of the questionnaire inquires about the twelve theories themselves. We begin each section by succinctly summarizing one of the theories and then posing a question like, "How important is this in explaining the speed of price adjustment in your company?"<sup>4</sup> Respondents answer freely in their own words, but interviewers code the responses on the four-point scale indicated at the bottom of Table 2.

While the scale resembles the typical grading system used at universities, we should expect the "grades" to be more compressed. At the high end, no theory is ever going to score a perfect 4.0. A theory rated "very important" by half the firms and "of minor importance" by the other half would be a fine theory indeed. So I would interpret a score of 3.0 as excellent. At the low end, a score of 1.0 would mean that every firm in the sample dismissed the theory out of hand. This is not a grade of D, but a kind of super F. Indeed, any score below 1.5 is a remarkably poor performance.

If the respondent says a particular idea is "totally unimportant" to the speed of price adjustment, we ask no further questions about that theory. Otherwise, we ask a variety of follow-up questions (sometimes many,

<sup>4</sup>For many theories, there is a preliminary factual question. For example, before we ask whether judging quality by price deters prices increases, we first ask whether the firm thinks its customers actually do judge quality by price.

TABLE 2—SUMMARY EVALUATION OF THEORIES<sup>a</sup>

Rank	Theory	Mean Score <sup>b</sup>	Percentage "3" or Higher
1	Delivery lags/ service	2.86	65
2	Coordination failure	2.85	66
3	Cost-based pricing	2.72	57
4	Implicit contracts	2.52	56
5	Explicit nominal contracts	2.29	40
6	Costs of price adjustment	2.28	43
7	Procyclical elasticity	1.97	36
8	Pricing points	1.97	33
9	Inventories	1.72	28
10	Constant marginal cost	1.56	19
11	Hierarchies	1.54	18
12	Judging quality by price	1.45	12

<sup>a</sup>Based on 72 interviews completed as of mid-November 1990.

<sup>b</sup>Based on the following scale: 1 = totally unimportant; 2 = of minor importance; 3 = moderately important; 4 = very important; N = don't know or cannot answer.

sometimes few) geared specifically to each theory. In this interim report, I ignore follow-up questions and concentrate on the main questions about the perceived validity of each theory.

Table 2 summarizes the central results on the validity of the twelve theories, as perceived by actual decision makers, ranked in order of popularity. The second column names the theories briefly, and sometimes cryptically. (Explanations follow.) The third column records the mean score on the four-point scale; remember that we expect most theories to score in the 1.5–3.0 range. We can interpret a rating of 1 or 2 as meaning that the firm rejects the theory as an explanation of price rigidity, and a rating of 3 or 4 as meaning that the firm accepts it. So the fourth column gives the percentage of respondents that rate the theory as 3 or higher—an indicator of the fraction of the private, for-profit GNP to which this theory

applies. The two alternative ways of ranking the theories agree closely. Ranked from best to worst by mean score, the theories are:

1) *Delivery lags / service*:<sup>5</sup> The idea here is that price is but one of several elements that matter to buyers. Rather than cut (raise) prices when demand is low (high), firms might prefer to shorten (lengthen) delivery lags or provide more (less) auxiliary services. The mean score for this theory so far is a healthy 2.86. Seventy-six percent of the firms accept the premise, and 65 percent say it is an important factor in slowing down price adjustments.

2) *Coordination failure*: This very old idea has been revived and formalized in recent New Keynesian theorizing. The notion is that firms might like to raise or lower prices, but hesitate to do so unless and until other firms move first. Once other firms move, they follow quickly. The mean score for this theory of 2.85 and its "acceptance rate" of 66 percent are quite high indeed.

3) *Cost-based pricing* connotes another old Keynesian idea: that prices are based on costs and do not rise until costs rise. Well over half the firms in our sample so far rate this a moderately or very important factor in explaining the speed of price adjustment; the mean score is 2.72.

4) *Implicit contracts* connotes the "invisible handshake" theory that firms have implicit understandings with their customers which proscribe price increases when markets are tight. Though this theory obtains an average score of 2.52, it elicits strikingly bimodal responses. Sixty-one percent of the sample accepts the premise that implicit contracts exist and, within this subsample, respondents generally think such contracts are an important source of price stickiness: the mean response within this group is a stunning 3.42.

5) *Explicit nominal contracts* refers to the naive Keynesian idea that written contracts prohibit price adjustments while they remain in force. Most firms have such con-

tracts, for at least some of their products; but discounting appears to be common. Our rough estimate is that this is an important factor in price stickiness for only about 40 percent of the economy. The mean score (counting firms with no explicit contracts as 1's) is 2.29.

6) A number of theorists have suggested that firms incur special *costs of price adjustment* whenever they change prices. Both menu costs and convex costs are lumped together under this rubric, but follow-up questions distinguish between the two. About 70 percent of all firms report that they have such costs. But fewer than half think them an important factor in slowing down price responses. The average score is 2.28.

7) According to a very old idea revived in the mid-1980's, demand curves become less elastic when they shift in. Such *procyclical behavior of elasticity* would lead to countercyclical markups which would, in turn, rigidify prices. This is another case where the premise is often accepted (59 percent of the time), but the idea is thought to be a significant source of price rigidity in only 36 percent of the cases. The average score is just below 2, which connotes minor importance.

8) *Pricing points* refers to the idea that certain prices (such as \$19.95 for a shirt) are psychological barriers that firms are reluctant to cross. This also scores 1.97 on average. Again, a majority of the firms (53 percent) accepts the premise; but only a third think it explains much price stickiness.

9) *Inventories* refers to the theory that, when demand rises (falls), firms draw down (build up) their inventories rather than increase (decrease) prices. Note that this question is not asked of service companies, so the sample size so far is only 33. Of these, however, fewer than a third think inventories a significant factor in deterring price adjustments. The average score is a low 1.72.

10) *Constant marginal cost* is shorthand for the notion that prices are sticky because both marginal costs and markups are constant over the business cycle. Since its ratings are so poor (a mean score of only 1.56

<sup>5</sup>Space constraints prohibit me from giving scholarly references for each of the theories. I apologize to those slighted.

and an acceptance rate of just 19 percent), I should explain how we posed the question. First, we asked respondents how their “variable costs of producing additional units” (our plain-English translation of marginal cost) behave as output rises. Only 42 percent of the sample reported that their MC was constant. Among this minority, about 60 percent said that the constancy of MC was totally unimportant or of minor importance in explaining their speed of price adjustment.

11) *Hierarchies* is a code word for a “theory” that does not come from the academic literature. Rather, it was suggested to us by an executive of a large corporation. The simple idea is that price changes are slowed down by the difficulty of getting a large, hierarchical organization to take action. Its low average score of 1.54 may appear unsurprising, since it should apply only to giant companies. But, even among firms with annual sales over \$100 million, its mean score is only 1.66.

12) The worst theory, according to our practitioners, is *judging quality by price*, a theory I have mentioned before. It obtains a paltry mean score of 1.45 and an acceptance rate of just 12 percent. Less than a quarter of the firms in the sample so far think their customers would “interpret a price cut as a signal that the quality of the product has been reduced.” In our tabulations, we count the others as rating the theory “totally unimportant.” However, among the minority of firms whose customers *do* judge quality by price, roughly 45 percent say it is a moderately or very important factor in dis-

couraging price increases. Thus this theory may be of some importance within a narrow sector of the economy.

In summary, it seems to me that the theories divide themselves into three groups. The top four listed above distinguish themselves as especially promising and/or applicable to the U.S. economy. The leading theory (by a slim margin), delivery lags/service, has not received the attention it deserves. The other three look to me decidedly Keynesian. Cost-based pricing is, of course, old-fashioned Keynesian stuff. Co-ordination failure is a major strain of New Keynesian theorizing. And the “invisible handshake” is by now part of the Keynesian tradition.

At the bottom of the list, we find four theories that appear to be rejected by our respondents. Two of them have enjoyed great popularity in the 1980’s. If practitioners are to be believed, marginal costs are constant in only a minority of industries. And judging quality by price appears to be neither a common nor an important phenomenon.

It is true that these results are based on only 36 percent of the ultimate sample. But I will be surprised if the rankings change dramatically as more data come in; 72 observations is sufficient to answer most questions about national averages. The payoff to a large sample will come later, when I have enough data to answer disaggregated questions like, “What kinds of firms make ‘invisible handshakes’ with their customers?” I will report information of that sort at a later date.