11
Conclusion

What have we learned from this extensive and (to us) exhausting review of the minimum wage literature? It is only fair to the possibly equally exhausted reader that we briefly summarize our findings and provide some, but not too many, comments about the implications of these results. We include some comments about implications for the minimum wage as a tool of policy and about how economists go about studying controversial topics.

Evidence leads us to conclude that moderate increases in the minimum wage are a useful means of raising wages in the lower part of the wage distribution that has little or no effect on employment and hours. This is what one seeks in a policy tool, solid benefits with small costs. That said, current research does not speak to whether the same results would hold for large increases in the minimum wage. Our suspicion is that large increases could touch off the disemployment effects that are largely absent for moderate increases, but evidence for the United States is lacking because there have not been large increases in the last generation. Similarly, increases in the minimum wage are not the only policy needed to address issues of low income in the United States. As many others have argued, programs such as the EITC and Food Stamps play a critical role in placing a floor under incomes and consumption, and higher minimum wages are not a substitute for such programs. In other words, the minimum wage is a useful tool for policy and, as with most policy tools, must be used wisely and in coordination with other policies to achieve the desired end.

A SUMMARY OF OUR FINDINGS

Employment

Employment, long square one for disputes about the minimum wage and its effects, has been more intensively studied than any other variable

in the minimum wage, both in the NMWR and before. Unfortunately, much work exhibits one or the other of two statistical problems (Donald and Lang 2007; Moulton 1990) that make judgment of the reliability of the analyses impossible. Little of the best work since then has been able to detect a substantively significant response of employment, measured as the number of jobs, the number of people working, or the number of hours. Although this does not close the issue, the preponderance of the evidence currently leans that way. This is borne out by our own meta-analysis which, depending on the specific controls used, finds an overall elasticity in the range [−0.048, −0.018], closer to the right end of the range than the left.[1] When the correction for publication bias is not based on an assumption of symmetry, the range shifts toward zero, to [−0.027, −0.014], with only the estimate closest to the right end of the range statistically significant. Considering only teenagers, the range of elasticities over all estimates is [−0.020, −0.042], and none of the individual estimates (from the metaregressions) is statistically significant. Restricting the sample of estimates considered to those that likely do not suffer from either of the statistical problems alluded to above, use only data from the United States, and consider only employment and not hours, the interval is [−0.084, −0.035]: the left extreme of the range is neither statistically significant nor based on a robust specification of the metaregression. Excluding it gives the range [−0.065, −0.035] and a mean elasticity of −0.053. The corresponding range and mean for restaurants is [0.068, −0.012] and −0.033, respectively. While some of the estimates that define these ranges are statistically insignificant and others are statistically significant, none are large enough to be economically significant.

How long does it take for employment responses to increases in the minimum wage to play out? Baker et al. (1999) is one of the most ingenious and frequently cited analyses by someone other than a participant in the original conference. Providing an interpretation of earlier work, it suggested that the period in question is at least five to six years. Coming, however, before Bertrand, Duflo, and Mullainathan (2004) and the widespread recognition of the Moulton (1990) problem that followed, the robustness of their interpretation is not clear. Following the logic of the rational expectations hypothesis, Pinoli (2010) argues that much of the response to an anticipated increase occurs before the increase itself. From the little work that is relevant, this does not appear to be the case,

at least in the United States. Furthermore, analyses that directly examine the timing indicate that the employment response takes no more than three years to complete (Belman and Wolfson 2010). However, considerably more work, using data from a variety of situations and countries is necessary before drawing any conclusions with confidence.

Gross Flows

If employment is not affected, what is? For one, there is some evidence of declines in both accessions (hiring) and separations (quits and layoffs). This suggests that jobs do become harder to find following increases in the minimum wage, and that previously employed workers are sufficiently productive, at least afterward, that employers are not inclined to fire them. Explanations for this include a different allocation after the increase of the surplus that the employee creates in this job, a reorganization of the work process following the increase (known as the *shock effect*), and efficiency wage theory, where the higher minimum wage induces greater productive effort from the worker.

Unemployment

A few analyses examine the response of the unemployment rate to increases in the minimum wage. The only one that makes no recourse to economic models detects no effect on the unemployment rate of teenagers. Those that begin with well-articulated models report a moderate increase in the unemployment rate. This disagreement raises questions about the extent to which the results are baked in, that is, whether the theoretical models determine the result. One of those that indicates an increase in the unemployment rate attributes it to increases in the labor force participation rate—in other words, increases in the minimum wage induce those without jobs to begin looking for them.

Even if the unemployment rate does not move up and down with the minimum wage, the decline in hiring suggests that after increases it may be harder for those already unemployed to find work following increases. This would be reflected in the length of unemployment spells. The one study that examines this reports a decline in unemployment duration for male high school graduates and increases in duration for several groups: men without a high school diploma, women with a

high school diploma, low-skilled women, and women older than 24. The one decrease in unemployment duration is slightly more than a month, while the increases range between three and a half months and five months.

Wages and Benefits

What effect does the minimum wage have on wages and their distribution? It is evident that average wages rise along with increases in the minimum wage. Most studies that look specifically at workers who were previously earning less than the new higher minimum wage report higher wages after the fact. It is quite clear that wages of the lowest paid 10 percent of workers are higher following a minimum wage increase, and for women it appears that this is true for the lowest 30 percent. The extent of spillovers varies by country, research on spillovers in the UK suggest it reaches considerably higher up the wage distribution, and particularly the male wage distribution, than is the case for the United States. Current research also finds that while many incumbents quickly move out of minimum wage jobs after entering the labor market, a substantial fraction of U.S. workers spend much of their first decade at the minimum wage or at wage levels that are affected by the minimum wage. Finally, studies that look at the entire wage distribution report that the minimum wage does indeed raise wages at the bottom and reduce wage inequality.

It has long been suggested that employers may respond to minimum wage increases by reducing spending on training, fringe benefits, and working conditions valued by employees. Results for health insurance are mixed, as are those for training, though some of the evidence for a decline in training after increases in the minimum wage seems to be due to those in low-wage jobs receiving little training irrespective of minimum wage policy. One line of attack has been to study the response of quit rates to minimum wage increases; the underlying reasoning is that in accepting job offers after an increase, workers' expectations about fringe benefits and working conditions are based on correlations with the wage from before the increase, correlations that do not hold after the increase. They learn this only on the job, at which point, according to this argument, many quit. Here again, results are mixed, but overall it appears that quit rates decline following increases.

Enrollment

As with many types of policy, the minimum wage may have unintended consequences. One of the most serious that has received attention is that it may induce teenagers to leave school, interrupting or prematurely ending their formal education. An issue with this literature is how little discussion there is between those approaching the topic from economics and those coming from education policy. Another is statistical problems that likely overstate the precision of the estimates. These problems and the disagreement among results suggest that a definitive answer awaits further work, although a reasonable foundation exists on which to build it.

The Product Market

Broadly speaking, economists use two competing classes of models to understand the minimum wage and its effects in the labor market, competitive models and a variety of models referred to as monopsonistic models, the most important of which involves firms' searching for workers and workers' searching for jobs. Competitive and monopsonistic models have different implications for the response of both employment and product prices. Competitive models imply lower employment and higher output prices in response to a minimum wage increase. Monopsonistic models allow for the possibility of higher employment, which in turn implies lower prices in affected industries. It is quite clear that restaurant prices rise by a small amount following minimum wage increases. In industries that are both sensitive to the minimum wage and face foreign competition, the price response appears to be weaker. In the UK, there is some evidence that increases in the minimum wage reduce profits at affected firms, but it is difficult to detect this in exit rates, that is, in firms going out of business, suggesting that the response is small. Analyses of the value that financial markets in the United States and New Zealand place on firms that are affected by the minimum wage indicate no pass-through to profits. Perhaps the price response, while weak, protects profits enough to make it worthwhile to remain in business.

Who is Affected by the Minimum Wage?

Although the minimum wage has no negative net effect on employment and raises earnings among most groups, its effects are not identical across all groups. Evidence of a modest negative employment effect among teenagers exists, and the few studies that have focuses on lesser-educated women suggest an effect there as well. Strong evidence exists in support of a wage effect for low-wage workers and among women. There is also some evidence that the minimum wage leads to lower levels of health coverage among less-educated workers at small firms. While the difference by age, gender, and income suggests that the effects of the minimum wage are heterogeneous, research and policy analysis should not focus on the groups (e.g., teenagers) where researchers expect to find a result. The emphasis on teens in employment research has provided detailed if contradictory information into the effects on teens, but has not provided insight about other affected groups or the labor market as a whole. A balanced view requires systematically casting a broader net than has typified existing research.

WHAT WE KNOW WITH CONFIDENCE

What then can we be reasonably certain of with respect to the minimum wage? As the minimum wage increases considered in this research have been moderate, the conclusions that we draw are premised on moderate increases.

Under such conditions there is little evidence of negative labor market effects. Hours and employment do not seem to be meaningfully affected. Accessions and separations may slow after minimum wage increases. Decisive evidence that training or benefits responds at all to increases in the minimum wage does not exist.

The evidence on schooling is suggestive but not sufficient to draw conclusions for policy making. This reflects both limitations of the currently available studies, and the variety of outcomes across studies. If any conclusion can be drawn from extant studies, it is that any negative effects of the minimum wage on school enrollment is associated with allowing students to leave school in the first two years of high school.

Mandatory attendance laws that only allow students to leave in their junior or senior year of high school appear to eliminate any negative effect on school attendance.

There is strong evidence that the minimum wage boosts the earnings of the lowest wage workers, and it may boost the earnings of those earning moderately higher hourly wages. In almost every wage study, the effect is more marked for women, who are more likely than men to be in low-wage positions.

Considered together, increases in the minimum wage raise the hourly wage and earnings of workers in the lower part of the wage distribution and have very modest or no effects on employment, hours, and other labor market outcomes. The minimum wage can then, as originally intended, be used to modestly improve the conditions of those working in the least remunerative sectors of the labor market. While not a full solution to the issues of low-wage work, it is a useful instrument of policy that has low social costs and clear, if limited, benefits.

BIG IDEAS

Given the certain predictions of core economic theory, how is it possible that the minimum wage raises wages without the anticipated negative effect on employment and employment-related measures? We are reluctant to spend much time and ink on this as the ground has been covered repeatedly over the last two decades. Earlier we alluded to the two most widely discussed models of the low-wage labor market, the competitive model and monopsonistic/search models.

Another possibility, which Kaufman (2010) has ably synthesized, is that in the presence of transaction costs and uncertainty about the future, firms do not respond to wage changes mechanistically. In this view, an employment relationship is not the exchange transaction depicted in the core economic model, but rather a relationship that persists over time because of transaction costs in the labor market. Moreover, uncertainty about these costs means that firms do not face a precisely defined relationship between wage levels and employment.

In Kaufman's (2010) view, although the use of labor in the production process is universal and labor's place in production can be

organized many ways, the labor demand curve used in the core economic analysis presumes an employment relationship: "That is, firms are the 'employer' who go to the labour market and hire people as 'employees' to provide a certain amount of labour services and follow the directions of the employer in return for a certain amount of remuneration per time period" (p. 776).

Such a relationship cannot exist absent transaction costs because without such costs, labor markets become competitive markets for inputs from atomistic independent contractors. Rather than each supplier being an employee of a specific firm, each has a separate relationship with one or more firms to provide a product or service. There is no employment relationship and no labor market that differs in important ways from, say, the market for ketchup (see Summers [1995]).

Transaction costs arise from the combination of limited human rationality, imperfect information, and ambiguities in property rights. They require that employers and employees establish employment contracts that define the terms under which work will be performed, and so define labor supply curves.[2] Employment contracts are inherently incomplete; they cannot fully specify outcomes because changes in circumstances such as economic environment, technology, and consumer taste necessitate altering the terms of employment. For example, most employment contracts establish compensation but not hours or employment levels. Changes in the circumstances facing an employer can result in large fluctuations in employment and hours. Because of imperfect information and bounded rationality, marginal product schedules and the demand curves derived from those schedules are probabilistic. If future circumstance A occurs, then the demand curve is in position A, but if future circumstance B occurs, then the demand curve will be in the nearby but not identical position B. The firm faces a "set" of demand curves with different likelihoods attached to each curve. The firm does not know which curve will be realized tomorrow and may well be uncertain about which demand curve it is on today. This moves firms away from a deterministic relationship between wages and employment levels and provides some latitude for firms to set wages. Latitude in wage setting is increased because the costs of job search for individuals provide firms limited monopolistic power over employees. Because of the transactions costs, deviation from a market wage will result neither in instantaneous loss of labor if a firm pays a below market wage nor a

long line of individuals seeking employment if it offers an above market one.

Our own thoughts on why there is such a weak employment effect follow a different and possibly more practical path. Economic theory and models are developed to explore specific topics. This focus makes possible rigorous exploration and full development of the implications associated with an issue of interest. This approach does not replicate the situation of decision makers in the market. Decision makers are daily confronted not with a situation in which all is constant *except* the change in the minimum wage. Rather, they face a world in which little is constant from day to day, week to week, month to month, or year to year. Not only does the minimum wage change, so do prices of supplies, fuel, rental, and myriad other factors. Demand is constantly changing as economic conditions, changes in views and tastes, and chance influence consumers' choices.

In a situation where so much is in flux, the stylization of decision making used in economics does not reasonably approximate decision makers' actual situations. In determining how much to produce, employers cannot simply take the consumer demand curve as fixed and, having determined their price, know what quantity to produce. Instead, in the face of a shifting demand curve, one must determine both price and quantity and then accept either the excess that could not be sold or the lost profits due to less product than could have been sold at that price. Because the firm's labor demand curve is derived from its product demand curve, the human resources executive decision maker has no more certainty about the appropriate number of employees to hire than about the price and quantity needed to exactly satisfy demand.

In such a world, one dominated by change rather than comparative static exercises, economic actors are unlikely to make decisions on the knife edge depicted in economic diagrams. Rather, small changes in prices are unlikely to move the decision maker to action. Slight movements in rent, fuel prices, or wages are unlikely to cause decision makers to rethink their use of inputs. While a large price increase might have sufficient effect for one to reconsider how to use that and other inputs, small increases likely get lost in the change of day to day operations.

This view is consistent with our results. It does not require that demand curves neither exist nor slope downward; rather, it suggests that the downward slope is not a one-dimensional line but rather a line with

some width, implying that for any quantity of employment, the firm is willing to pay a wage within a defined range. Were the minimum wage to increase 50 percent, it would be beyond the range consistent with the current employment level, and firms would reduce their employment. However, when increases are moderate, are within the range consistent with current employment levels, decision makers are too engaged with the world to change their existing arrangements. It is also consistent with the finding that increases in the minimum wage reduce accessions. Both the formation of a firm and the decision to expand require positive action to bring new employees into a business. At such times, decision makers may well consider the cost of inputs, including labor inputs, and alternative arrangements.

If we supplement economic theory with this view of the situation facing economic decision makers, we are then likely to conclude that thresholds, which must be crossed before decision makers act, exist. These thresholds differ by market and individual and are unlikely to be stable over time. Without a doubt, the changes in the value of financial instruments that impel arbitragers and their computers to action are many times smaller than those required to attract the attention of a retailer or fast food franchiser to change their employment policies. The rise in the minimum wage needed to catch the attention of the fast food franchiser in Westchester County, New York, may be far larger than that for exurbia in Alabama. With such an understanding of the world, the lack of a relationship between moderate increases in the minimum wage and employment no longer stands in contradiction to core economic theory. Rather, it points to a research program to investigate the factors affecting thresholds of action in labor markets.

ISSUES OF INNOVATION AND CRAFTSMANSHIP

Although the purpose of this review is not to reflect on the work of the economics profession, our reading and rereading of hundreds of articles reveals a tension between innovation and craftsmanship. The drive toward innovation is strong in economics—witness the popularity of *Freakonomics* (Levitt and Dubner 2005)—and has been productive over the last 50 years with the regular rethinking and expansion

of economics thought. Innovators such as Gary Becker, Ronald Coase, William Baulmol, and George Akerlof have done much to advance our understanding of markets and of *homo economicus.*

Answering the important economic questions of the day requires more than novelty. While it can provide new understandings and approaches, unless founded on strong methodological approaches and placed in the context of prior work, it forgoes much of its opportunity to expand knowledge. Too often in our review we have been unable to reconcile results across journal articles because the authors have not systematically explored the sources of differences between what they present and prior work on the topic. Even when articles draw on the same data sources, differences in time period, technique, and measures preclude knowing the source of differences, sometimes dramatic, in the results. All too often, systematic investigation would have required no more than one table and a page or two of text. Furthermore, in too many cases, authors have failed to investigate important variations of their model to examine the robustness of their results. For example, although it is well established that the choice of control group can affect the estimates from difference-in-differences models, most authors choose to present estimates for one or possibly two control groups rather than for each of the obvious control groups.

Absent greater emphasis on craftsmanship, on the workmanlike investigation of an issue, economists limit their contributions to our understanding of a topic. Without knowing how differences in controls, data, time period, and method influence results, we are left with too many unreconciled findings. Craftsmanship is particularly important when topics are controversial, because there is a greater need to understand the sources of differences between studies and so limit the scope of passion.

How might economics place greater weight on craftsmanship? Partly by training graduate students to be more thoughtful about their research, but more so by reviewers and editors requiring authors to explicitly reconcile their work with prior work and to address reasonable variants on their models. This requires that editors and particularly reviewers be familiar with the topic under study. It also requires that they ask authors to compare and contrast their work with prior work explicitly, and to investigate differences. We believe that reconciliation would not require too much empirical effort or too many journal pages. It may

be challenging intellectually, as it is likely to bring to the fore issues of control groups, time periods, and measures.

Redressing the current imbalance between innovation and craftsmanship is then important to advancing our understanding of markets and the investigation of markets. Innovation provides the drive forward, while craftsmanship provides integrity. Both are required for the robust and credible investigation of markets.

Notes

1. Caveat lector: As we discuss in the section that describes the meta-analysis, it does not include all articles because too large a number do not report results in a way that makes them comparable to others.
2. Following Rosen (1974), the labor supply curve is defined not only by the wage, but all of the terms and conditions of employment available to employees. The presence of labor supply curves is then premised on there being an employment relationship that specifies those terms and conditions, and following back a step in Kaufman's logic, on transaction costs.