

THE IMPACT OF EMOTIONS ON WAGE SETTING AND UNEMPLOYMENT*

Lorenz Goette

University of Zurich and CEPR

David Huffman

IZA Bonn

Abstract

This chapter argues that the neglect of emotion in economic models explains their inability to predict important aggregate outcomes in the labor market. We focus on one example: firms frequently cut real wages, increasing nominal wages by less than the inflation rate, but they very seldom cut nominal wages. This pattern suggests that workers exhibit a special resistance to nominal wage cuts, which is hard to explain if they are purely rational as assumed in standard economic models.

We hypothesize that resistance to nominal wage cuts is best understood in terms of a model where salient features of a situation trigger emotional responses and sway judgment of the entire situation. Since a cut in the nominal wage leads to a very salient reduction in pay, we predict that the reaction of workers is dominated by emotions. On the other hand, an increase in the nominal wage may produce a more deliberative evaluation, because there is no immediately salient feature. The individual needs to compare the inflation rate to the wage change before it becomes clear whether the change increases or decreases utility, thus producing a more measured response. We present evidence from experiments showing that self-reported emotions do respond more strongly to nominal wage cuts, as opposed to cuts in the real wage achieved through increasing the nominal wage by less than the inflation rate.

Although emotions may benefit individual workers, by strengthening their bargaining position and preventing wage cuts, they may also lead to worse outcomes, in the form of higher unemployment.

* Goette gratefully acknowledges support from SNF and NCCR Affective Sciences.

I. Introduction

Traditionally, models of economic decision-making assume that individuals are rational and emotionless. This chapter argues that the neglect of emotion in economic models explains their inability to predict important aggregate outcomes in the labor market. We focus on one example, which is the fact that far fewer nominal wage cuts are observed in labor markets than predicted by traditional economic models. Firms frequently cut real wages, or purchasing power, of workers by increasing nominal wages by less than the inflation rate, but they seldom cut nominal wages. This pattern suggests that workers exhibit a special resistance to nominal wage cuts, which is hard to explain if they are purely rational.

Economists have tried to explain the infrequency of nominal wage cuts by maintaining the assumption that individuals are emotionless, but relaxing the assumption that they are cognitively sophisticated. In particular, if individuals suffer from a cognitive limitation known as "money" illusion, they fail to subtract the inflation rate from changes in the nominal wage, and thus interpret all nominal wage increases as leading to higher purchasing power. This mistake would make workers more resistant to a nominal wage cut in, say, a time of zero inflation, than to a nominal wage increase that occurs in a time of high inflation and implies exactly the same decrease in real wages. The problem with the money illusion explanation is that it requires that individuals are unable to perform subtraction, or that they underestimate the true inflation rate. Both of these requirements appear counterfactual. People are able to perform subtraction in many domains, and there is strong evidence that, on average, individuals actually have well-calibrated inflation expectations, even during times of very low inflation (see Mankiw, Reis, and Wolfers, 2003, for a review of the recent literature on inflation expectations).

We argue that strong resistance to nominal wage cuts is best understood in terms of a model where, consistent with evidence from psychology and neuroscience, salient features of a situation trigger emotional responses and sway judgment of the entire situation. Since a cut in the wage is a very salient feature, we argue that cutting the nominal wage leads to a reaction that is mainly dominated by emotions. On the other hand, we hypothesize that an increase in the nominal wage produces a more reflective evaluation, because there is no immediately salient feature: the individual needs to compare the inflation rate to the wage change before it becomes clear whether the change increases or decreases utility. Because the individual arrives at an understanding of the implications through a more reflective process, the ultimate emotional reaction is more tempered.

We test the predictions of this model by asking individuals how they would react to different hypothetical wage changes, in different economic environments. The subjects used in our experiment are individuals who are familiar with making these kinds of judgments. We find that cuts in the nominal wage trigger strong emotional reactions, above and beyond the emotional reaction caused by a real wage cut, achieved by increasing the nominal wage by less than the rate of inflation. On the other hand, varying the size of an increase in the nominal wage has no impact on emotional evaluations, holding the real wage constant. Thus, our evidence suggests that individuals are not "fooled" by nominal changes in general, as hypothesized in the money illusion explanation, but rather have an emotional reaction to the salient qualitative aspect of the situation. In summary, the results suggest that wage rigidity can be explained by a strong emotional reaction, triggered by the salient reduction in utility that occurs only in the case of a nominal wage cut.

We also review the existing empirical literature and find that it is consistent with the model outlined in this chapter. There is evidence that individuals strongly resent nominal wage cuts, and retaliate against the employer if the employer cuts their wage. It is also clear that employers shy away from using nominal wage cuts: looking at the distribution of wage changes in representative data sets, many workers receive wage freezes, but very few receive a cut in their nominal wage.

Taken together, the evidence discussed in this chapter strongly suggests that wages would be lower than they are now, in the absence of strong emotions triggered by wage cuts. Thus, it is tempting to conclude that emotions make workers better off, because they protect workers from nominal wage cuts. Such a conclusion would be premature, however. If nominal wages are kept high, this has consequences for the prices that firms charge for their products. It is straightforward to show in an economic model that the resulting price changes lead to lower employment. We review the evidence and conclude that downwardly rigid wages do indeed lead to higher unemployment, implying that the same emotions that prevent nominal wages from falling also lead, indirectly, to some other individuals losing their jobs. Thus, while it is clear that emotions play an important role in explaining stylized facts about the labor market, it is far from clear whether emotions improve outcomes in the labor market.

II. The Role of Emotions in Wage Setting

This section lays out a basic economic framework for analyzing employment relationships, and discusses evidence on the ways that emotions may influence wage formation.

A. The Basic Framework

A realistic model of the employment relationship incorporates three features. The first concerns the kinds of formal agreements that can be part of an employment contract. The other two capture important aspects of the preferences of employees.

(i) Contractual Incompleteness: Employment contracts are inherently *incomplete*. They do not specify all actions that an employee is required to take in all possible contingencies, in a way that is enforceable by a third party (Milgrom and Roberts, 1992). One cause of incompleteness is the cost of identifying and writing down all possible contingencies. Incompleteness can also arise even when it is clear to both the employer and employee what action is appropriate in a particular instance, because verification by an outside party is difficult or impossible. Together, these obstacles ensure that contracts are incomplete.

Contractual incompleteness confronts the employer with a motivational problem, which is how to motivate an employee to behave in the interests of the firm, when there are limited possibilities for enforcement. Several strategies have been proposed, for example paying a wage that makes the employee better off than her next best alternative, thereby creating an incentive to behave in the employer's interest in order to keep the job.¹

(ii) Reciprocity: paying high wages to motivate performance is particularly effective if employees have reciprocal preferences, i.e., if they respond to kind actions by the

¹ These mechanisms are known as efficiency wages. See Akerlof and Yellen (1986) for a survey. See also MacLeod and Malcolmson (1998) for a formal discussion of the role of bonus payments in this context.

employer (high wages) with kind actions (more effort than could be contractually enforced).

Akerlof and Yellen (1990) review the extensive literature in social psychology and incorporate key findings regarding reciprocity into a formal model of the labor market. It is notoriously difficult to verify the predictions of such a model using field data, however, as alternative explanations abound. A viable alternative is to test the behavioral predictions in laboratory experiments, which mimic a labor market with incomplete contracts. In a series of experiments, Fehr and colleagues show that reciprocity can indeed have an important impact on labor market outcomes (see Fehr, Kirchsteiger, and Riedl, 1993; Fehr, Gaechter, and Kirchsteiger, 1997. Fehr and Gaechter, 1997, survey the literature). In these experiments, firms consistently pay wages above the market clearing wage, and employees provide more effort than can be explained by selfish preferences alone.^{2,3}

(iii) Reference-Dependent Evaluations: The labor market experiments reviewed in Fehr and Gaechter (2002) show that reciprocity is a potentially important behavioral

² See also Brown, Fehr and Falk (2004), who show that labor market outcomes can change dramatically if even a small fraction of labor market participants have reciprocal preferences. Intuitively, this is because the presence of reciprocal types gives selfish agents an incentive to mimic reciprocal behavior, cooperating with the employer until they switch to a strategy of defection towards the end of their employment relationship.

³ See also Krueger and Mas (2004) for a study that uses field data. Their study shows that unkind actions of a firm can trigger strong behavioral responses. They look at data for a tire manufacturer, and document a surge in tire failures after the firm announced worse employment conditions for a part of the workforce. Given that technology and the quality of inputs appear to have remained unchanged, Krueger and Mas conclude that the tire failures were due to a retaliatory reduction in work effort.

force in the labor market. However, integral to the notion of fairness is a standard of comparison, or reference-point, which makes clear what types of actions are "kind" or fair.

Kahneman, Knetsch, and Thaler (1986) provided some of the first field evidence on the importance of reference points in fairness judgments. For example, they show that the fairness of a wage depends on how it compares to the wage paid to the employee in the past (see results for Questions 2 and 3 in their paper). They argue that transactions are coded as either gains or losses relative to the reference transaction, and that evaluations are characterized by loss aversion, i.e. a tendency for losses to be more painful than a gain of similar size is enjoyable (Tversky and Kahneman, 1991).

More recently, Shafir, Diamond, and Tversky (1997) show even more convincingly that the past wage is an important reference point in evaluating the current wage. Furthermore, both studies find that it is the past money wage, rather than the real wage, which appears to act as a reference point. For instance, Shafir, Diamond, and Tversky pose respondents with two scenarios: a small nominal wage cut when inflation is zero, and a small nominal wage increase when inflation is moderate. The parameters are chosen such that in both scenarios, the change in the purchasing power (the nominal wage change minus inflation) is the same. They find that individuals view the nominal wage cut as unfair, but evaluate the other scenario, with the same change in the real wage, much more favorably. Importantly, both studies appeal to money illusion in order to explain why nominal wage increases are coded as gains, even when they do not keep up with inflation. As alluded to in the introduction, economists have regarded this final argument with a fair amount of skepticism (see, e.g., McLaughlin, 1994). Use of the nominal wage as a reference point has the strong flavor of a very basic form of irrationality (and produces strong affective reactions in

economists), because it implies that individuals are incapable of calculating the real wage, i.e. subtracting the inflation rate from the nominal wage change. In support of this skepticism, introspection suggests that individuals are in fact able to perform additions and subtractions in many situations, and evidence from surveys shows that inflation expectations are actually quite accurate (Mankiw, Reis, and Wolfers, 2003). Thus, while reference dependence clearly plays an important role in determining reciprocal responses, it is desirable to provide an alternative explanation for why nominal wage increases do not provoke strong reactions, which does not involve an inability to perform addition or subtraction.

We propose to look at the issue from the perspective of dual-process models, in which the individual's decision-making depends on the interaction of two systems: a reflective system, and an impulsive, or affective system.⁴ This model maintains the assumption that the individual is cognitively sophisticated, e.g. the reflective system is assumed to be able to add and subtract. On the other hand, in certain circumstances the individual is assumed to experience emotions, or affect, which can override reflective decision-making. In line with evidence on the nature of affective evaluation, the affective system is assumed to respond more strongly to salient features of a situation. In the case of wage changes, we hypothesize that the nominal wage change is much more salient than the inflationary environment within which the wage change occurs. Thus, the affective system responds primarily to the nominal wage change. On the other hand, we propose that the reflective system primarily cares about the final purchasing power derived from a wage change, and we assume that the reflective system is able to calculate the real wage by subtracting the inflation rate.

⁴ This type distinction is common to many dual process models. See, for example, Strack and Deutsch (2004) and Satpute and Lieberman (2006).

We hypothesize that a nominal wage cut triggers an immediate and strong evaluation from the affective system, since the wage cut unambiguously reduces disposable income, without the need for further deliberation. When the wage change is positive, however, it requires some reflection to determine whether the wage change is an improvement or not: the individual must subtract the inflation rate from the nominal change, in order to determine whether the wage increase will make the individual better off or worse off than she was a year ago. Thus, we hypothesize that a nominal wage increase produces less of a response from the affective system, and instead activates the reflective system more strongly.⁵ Consistent with many examples from the literature on conflicts between affective and reflective decision-making, we assume that the affective system can strongly skew the response of the reflective system when it is highly aroused.⁶ On the other hand, when the affective system is not strongly aroused, behavior is assumed to correspond more closely to the goals of the reflective system. The resulting dual-process model can easily generate strong responses to a nominal wage cut. On the other hand, the same model can produce more measured responses to a nominal wage increase, where, clearly, the individual behaves as if she can add and subtract.

⁵ Hsee and Rottenstreich (2004) provide experimental evidence that is relevant for this point: emotionally-charged priming questions, e.g. questions about the emotions aroused by the word “baby,” lead subjects to evaluate a subsequent economic transaction in emotional terms (subjects’ willingness to pay for a bundle of music CDs was relatively insensitive to the amount of CDs in the bundle). A prime involving deliberative thinking, e.g. simple math problems, triggered evaluation of the transaction in deliberative terms (willingness to pay responded to quantity).

⁶ Loewenstein (1996) discusses various examples in which affective processes come to dominate deliberative decision-making, as arousal of affective processes increases.

B. Evidence on Emotions triggered by Wage Cuts

The Design. Undergraduates nearing graduation at the University of Zurich were recruited to participate in the study. Initially, subjects were asked to imagine a situation in which they had been working at a firm for about a year, during which time they had received no feedback regarding their performance. They were asked to imagine that it was now the end of the year, and that they had been called in for a meeting with their boss, to discuss their salary for the next year. The subjects were then presented with a first scenario, in which the boss proposed a specific wage change, and they were informed about the current inflation environment. They were asked to try their best to imagine how they would react to this wage change, and to indicate how well a list of emotions provided to them would characterize their reaction. They had to indicate their answers on a seven-point scale, with 1 indicating that the emotion in question did not describe their reaction at all, and 7 indicating that it described their reaction accurately. Following this rating, they were presented with two additional questions. In the first they were asked to imagine that, as in many workplaces, there were some tasks that needed to be done to ensure smooth operation of their firm, but that the one performing these tasks would go largely unrewarded for doing so. They were asked how likely they were to do these tasks following the meeting with their superior, again using a scale from 1 to 7. They were then asked how likely it was that they would look for a new job following the meeting, on a 1 to 7 scale. After this first scenario, subjects were posed with a second scenario in which the hypothetical wage increase was five percent higher, and whole procedure was repeated.

There were four different treatments in the first scenario. Table 1 provides an overview over the different conditions. Each subject was in either treatment A, B, C,

or D. If the subject was in treatment A in the first scenario, he or she was in A' in the second, and so on. Hence, all comparisons between A, B, C, D (and their primed counterparts) used between-subject variation. The others were identified by within-subject comparisons.⁷

Results: Figure 1 presents the emotional reactions to the different wage changes, plotting the mean score for each emotion. The top panel of Figure 1 shows scores for treatments A, B, C, and D. The bottom panel shows the scores for treatments A', B', C', and D'. Comparing treatments A and B, which involved a 2 percent decrease in the real wage, to C and D, where the real wage was cut by 4 percent, we see that a larger real wage cut leads to more anger and disappointment. However, comparing A to B, and C to D, we also see that the emotional reaction is consistently stronger in cases where there was a nominal wage cut, compared to cases where there was not, holding the change in the real wage constant. Linear regression results, presented in Table 2, tell a similar story. The dependent variable for each regression is the corresponding emotion score. As explanatory variables we include dummy variables for each of the possible real wage changes (from all scenarios), so that the effect of a nominal wage cut is identified by comparing A to B, and C to D. Because each subject participated in two scenarios, we adjusted the standard errors to correct for possible correlation of the error term across observations for the same subject. Despite this conservative calculation of standard errors, we find a highly significant effect of a nominal wage cut: anger and disappointment are clearly higher, and joy is clearly lower when a nominal wage cut is involved, holding the real wage change constant. Interestingly,

⁷ The most important results in this study are identified using between subject-variation. We are confident that our other findings, based on within-subject variation, would also hold up in a between subjects design as well.

the subjects also appeared more surprised by a real wage cut if it involved a nominal wage cut. Real wage changes clearly also matter for the evaluations, and the coefficients on the corresponding dummy variables indicate that these effects are quite large. For example, going from A to A', or B to B', we find large shifts in the evaluation of the wage scenario. Interestingly, when we compare A to C, and B to D, we find impacts that are much smaller, even though they entail the same difference in the real wage change. The difference between the two comparisons is that going from A to A' is comparing a real wage loss to a real wage gain, while going from A to C is comparing a large loss to a small loss (or a small gain in the case of comparing A' to C'). The difference in the results in these cases suggests the presence of strong loss aversion with respect to the real wage.⁸

The design of the study also allows us to examine whether there is something about a nominal frame *per se* that leads to a strong emotional reaction, or whether strong emotional reactions are confined to the case of nominal wage cuts. If a nominal frame affects evaluation in general, it should be true that the different nominal wage increases in the second set of scenarios lead to different emotional evaluations. However, inspection of Figure 1 already suggests that our results are not driven by an emotional reaction to nominal wage changes in general. For example, while we find a clear effect of a nominal wage cut on the joy ratings in C relative to D, we find no effect on joy ratings of a difference in the size of a nominal wage increase, as shown by a comparisons of A' to B', and C' to D'. To test this more formally, we ran the same regressions as reported in Table 2, but also included the nominal wage change as an

⁸ One should keep in mind, however, that these comparisons rely on within-subject comparisons. We are confident that qualitatively the same results would result in a between-subject design.

additional explanatory variable in each of the regressions. In Table 2, we report the p -values of these coefficients. As can be seen, there is not a single significant effect at conventional significance levels, while the nominal wage cut dummies remained highly significant. Thus, difficulty with disentangling nominal from real changes does not explain our results.⁹ Rather, only nominal wage cuts appear to evoke a strong emotional reaction, holding the real wage change constant.

It is also noteworthy that the effect of a nominal wage cut is quantitatively large. For example, the coefficient measuring the impact of a nominal wage cut is approximately twice as large as the coefficient measuring the impact of a change in the size of a real wage cut, from 4 percent to 2 percent.

C. The Consequences of Wage Cuts

Figure 2 displays results based on the two measures of loyalty to the employer. We find that nominal wage cuts undermine the loyalty towards the employer. Interestingly, the effect is much stronger for the tendency to look for a new job (possibly because refusing to do the unrewarded tasks also hurts fellow workers). Again, the results do not appear to be driven by confusion about nominal versus real changes in general. We report p -values in the table showing that the data do not support this hypothesis.

⁹ This contrasts with Shafir, Diamond and Tversky (1997) who find that money illusion affects judgment in general. However, the types of question they used asked individuals what they thought *others* would do in a given scenario, not what they themselves would do. Fehr and Tyran (2001) show that there is a tendency to overestimate the extent to which other individuals are prone to money illusion.

How do these results compare to the results of other studies? Bewley (1999) conducted interviews with 288 compensation managers and personnel officers. In these free-form interviews, almost 70 percent of managers spontaneously indicated that work morale would suffer in response to a nominal wage cut, 40 percent expressed fear of higher turnover and lower productivity in this situation, and 15 percent expected even more extreme reactions such as sabotage and theft. The results obtained in this study are thus consistent with Bewley's evidence. In a more recent study, Mas (2005) examines the impact of final-offer arbitrations between municipalities and police officer unions on subsequent law enforcement. The results he finds are striking: in municipalities where an arbitrator ruled against the police officers, arrest rates tend to decrease over the subsequent year.¹⁰ Furthermore, larger "losses" relative to the level desired by the police officers lead to a larger drop in the arrest rate. On the other hand, larger "gains" do not significantly increase arrest rates. The results in Mas (2005) are thus consistent with a model where individuals care about outcomes relative to a reference outcome. While the results does not speak to the point of whether nominal wage cuts lead to lower effort of workers, they show that outcomes below a salient reference point lead to lower effort.¹¹

¹⁰ Mas (2005) is careful to check whether districts where the arbitrator ruled against the police differed from districts where the arbitrator ruled in favor of the police. He finds no difference in arrest rates and other characteristics of the police force before the arbitration is announced.

¹¹ In the case of final-offer arbitration, there is a good reason why the current nominal wage should not be relevant. Before a wage bargain comes to the point of final-offer arbitration, the parties involved have gone through a lengthy process of unsuccessful negotiations. Thus, it is particularly likely that the cases that do reach this stage are ones in which the parties feel particularly strongly about their desired outcome.

Overall, the results summarized in this section provide a strong rationale for why firms might shy away from nominal wage cuts, especially when the desired nominal wage cut is small.

III. Evidence of Downward Nominal Wage Rigidity

The model in the previous section predicts that nominal wage cuts occur less often than if workers were emotionless. To examine the relevance of this prediction in the field, one needs data on individual wage observations. A mass point at zero percent wage changes, and a discontinuous drop in the density of the wage changes below zero would be evidence of downward nominal wage rigidity (DNWR), as it might arise from the behavioral forces described in the previous section.

One approach to testing for DNWR is thus to calculate a histogram of wage changes, and assess the extent of a spike at nominal zero and asymmetry around zero. Studies that take this approach using data on the U.S. labor market are McLaughlin (1994), Card and Hyslop (1996), Kahn (1997), and Lebow, Saks, and Wilson (2003). The results indicate some nominal wage rigidity, but not much. Indeed, in an average year during the sample periods in these studies, about 20 percent of the workforce received nominal wage cuts, even though labor productivity rose quite rapidly and inflation was quite high.

A possible concern with these studies is that they understate the extent of DNWR. Wage data is known to contain a large amount of measurement error because individuals do not remember exactly how much they earn (Bound and Krueger, 1991). Using polluted wage data can produce measurement error artifacts that look like nominal wage cuts. Indeed, when one examines the results from studies that use

data known to be free of measurement error, e.g. personnel files, there are almost no wage cuts (Baker, Gibbs, and Holmstrom, 1994; Wilson, 1999; Fehr and Goette, 2005). In most of these datasets, wage cuts are less than one percent of all wage change observations, while there is a massive pile-up at nominal zero. Similarly, studies that ask directly about nominal wage changes (instead of panel studies that ask about wages at two points in time), find wage change distributions that look more like the evidence from personnel files. Akerlof, Dickens, and Perry (1996) conducted such a survey in Washington DC in 1994. They found that less than one percent of the sample experienced wage cuts, with extraordinary circumstances in each case. In a phone survey that we conducted in the Fall of 2004 in the city of Zurich (N = 240), we find very similar results. We find a large number of individuals with unchanged nominal wages in every year, while only about one percent have experienced a cut in their nominal wage.

Given that evidence from personnel files, and surveys asking directly about wage changes, suggests that measurement error masked a considerable amount of DNWR in previous studies, more recent studies using survey data try to correct for measurement error using more complicated statistical models. Altonji and Devereux (2000) allow for measurement error in the data and find that, once one takes measurement error into account, there is a marked degree of wage rigidity even in panel data on wages. Gottschalk (2005) finds similar results, using a different identification strategy for measurement error. DNWR also appears to be a very robust phenomenon, which is not easily malleable in response to changes in economic conditions. For example, Fehr and Goette (2005) examine whether, after a long period of long inflation, individuals give up their resistance against wage cuts. They use data from Switzerland in the 90's, a period marked by a significant recession, but very low

inflation rates. If anywhere, this is the type of environment in which one would expect to see wage cuts. However, they find that after four years of virtual price stability, there are only very few wage cuts (around 5 percent), while approximately 60 percent of the workers received wage freezes. Given that studies discussed so far all differ in terms of method, the country considered, and the type of data used, it is difficult to compare results across studies. In an effort to overcome this problem, 13 country-teams have formed the International Wage Flexibility Project, and plan to use the same type of correction for measurement error on all data sets (see Dickens and Goette, 2005, for details). First results in Dickens et al. (2005) indicate that nominal wage rigidity is strong in most of the countries considered.

Thus, once one controls for measurement error, labor force surveys lead to the same finding as personnel files and direct questions about wage changes: nominal wage cuts are very rare, and even after a long period of low inflation, there is no measurable tendency for wage cuts to become more frequent. But to what extent do firms anticipate this and pay lower wages to begin with, because they know that to some extent, any wage increase is non-reversible? As worked out in detail in Elsby (2005), a fall in inflation gives firms an incentive to increase the compression of wage raises (because a given wage increase is more likely to be irreversible at low inflation rates). Indeed, in British and U.S. data, Elsby (2005) finds that the distribution of wage *raises* is more compressed when inflation is low. However, there are two points to keep in mind with this result: first, measurement error in the data can generate a pattern that looks exactly the same as "compression," and Elsby does not correct for measurement error. Secondly, his results imply that compression fails to offset the impact of DNWR, even though compression is overestimated and DNWR is underestimated due to measurement error.

IV. Implications

The evidence reviewed in the previous section implies that wages are higher than they would be in the absence of DNWR. It is tempting to interpret this as a positive effect of emotions on economic outcomes at the individual level: because firms anticipate retaliation in response to a nominal wage cut, workers are in a better bargaining position, and thus are to some extent shielded from wage cuts.

There is, however, a second channel through which emotions can affect labor market outcomes. Suppose, for instance, that inflation is zero. If, as the evidence suggests, DNWR increases average wages, profit-maximizing firms will try to pass on some the additional costs to their customers, by increasing product prices. This, in turn, tends to lead to lower market demand, which may cause firms to hire fewer workers. In this way, DNWR can create higher unemployment. This prediction is quite difficult to test, however, because it requires either estimating a fully specified general equilibrium model of an economy, or requires detailed estimates of what wages would have been in the absence of DNWR. Akerlof, Dickens, and Perry (1996) make a first attempt at building a model along these lines, and estimate their model using post-war data from the U.S.. The model explains the variation in inflation and unemployment remarkably well. More strikingly, the authors then conduct an out-of-sample prediction for inflation and unemployment during the Great Depression, where, at times, the inflation rate was negative. Given the unusual inflationary environment, this sample period provides a particularly interesting test of whether a model incorporating DNWR explains the path of macroeconomic aggregates. Their model, estimated on post-war data, predicts employment, output, and prices in the 30s very well. These results suggest that a large part of the drop in GDP during the Great

Depression must be attributed to the downward rigidity of nominal wages at that time. Fehr and Goette (2005) use their econometric model to calculate the extent to which DNWR affected wages in separate geographical labor markets and industries. They find a robust and positive correlation between unemployment and the wage sweep-up in a labor market. Again, this finding suggests that DNWR reduced employment considerably during the low-inflation episode in the 90s in Switzerland. It should be noted that not all studies find an association between the extent of DNWR and economic outcomes. As described earlier, Lebow, Saks and Wilson find evidence at the micro-level of DNWR. Yet, they find no correlation between unemployment and DNWR as estimated in their data. Similarly, Card and Hyslop (1996) find little evidence that DNWR plays an important role in hindering labor market adjustment: they examine whether in high-inflation years, wage changes are more responsive to unemployment than in low-inflation years, but find little evidence to support this hypothesis.¹² One should keep in mind, however, that both studies considered a period with rapid average nominal wage growth, and hence, DNWR did not have much of an impact on average wages anyway. The cross-country evidence in Dickens et al. (2005) is also supportive of the view that nominal wage rigidity has an impact on the real side of the economy. They estimate models similar to those in Fehr and Goette (2005), and find that DNWR is associated with higher unemployment.

In summary, answering the question of whether emotions lead to better or worse decisions faces a difficulty often overlooked. While emotions may benefit an

¹² Altonji and Devereux (2000) examine how individual-level labor market outcomes are correlated with whether an individual received a wage cut, wage freeze, or wage raise. They find no systematic pattern. This is not to be expected, however, since the theory makes no strong predictions regarding individual-level outcomes.

individual worker, by strengthening the worker's bargaining position and discouraging firms from making wage cuts, it is far from clear that emotions benefit workers overall. We have sketched the conditions under which this need not be the case: DNWR can lead to higher unemployment, and the existing evidence points in this direction. Hence, even in cases where emotions make individuals unambiguously better off, equilibrium effects can counteract, and perhaps fully reverse these benefits.