A Pay Change and Its Long-Term Consequences

Miriam Krueger, *Deutsche Bundesbank*

Guido Friebel, *Goethe University Frankfurt, Center for Economic and Policy Research, and Institute of Labor Economics*

In a professional services firm, top management unexpectedly adjusted the pay of consultants in some divisions to the pay in other divisions. In this quasi experiment, fixed wages increased and bonuses decreased, reducing pay for the high performers and increasing it for the low performers. Individual outputs and efforts decreased by 30%, and attrition and absenteeism increased. The effects were driven by those who were rationally expecting to lose from the pay change. Observing a period of more than 3 years, we show long-term negative reciprocity of those affected but no negative selection effects of new hires.

This paper is based on chapter 1 of Miriam Krueger's PhD thesis (2015). We are grateful for the comments of seminar participants at the University of California, Berkeley; Cambridge; Cologne; Essex; Goethe University Frankfurt; the Institute of Labor Economics, Bonn; Kellogg; Massachusetts Institute of Technology Organizational Economics; Rotterdam; Rotman, Toronto; Tilburg; Wharton; the Asian and Australasian Society of Labour Economics Conference in Singapore, 2019; the Relational Contracting Conference in Madrid, 2016; the German Economic Association meeting in Augsburg, 2016; the European Economic Association in Geneva, 2016; and the Colloquium on Personnel Economics in Cologne, 2014. We thank the partners in our study firm and Yvonne Klemm for excellent assistance with the data. This paper represents the authors' personal opinions and does not necessarily reflect the views of the Deutsche Bundesbank or its staff. Contact the corresponding author, Guido Friebel, at gfriebel@wiwi.uni-frankfurt.de. Information concerning access to the data used in this paper is available as supplemental material online.

Submitted November 12, 2020; Accepted October 19, 2021; Electronically published May 25, 2022.

Journal of Labor Economics, volume 40, number 3, July 2022.
 https://doi.org/10.1086/717728

I. Introduction

Finding the right compensation scheme for a firm's workers is an intricate challenge for management. In the early stages of a firm's operations, management lacks precise knowledge about how workers will react to fixed and variable wage components and how those reactions translate into profit.[1] Over time, management gains this knowledge and may be tempted to adjust the compensation scheme, to save on the costs of motivating and attracting productive workers. But adjustments, meant to reduce workers' rents, can be costly. Workers may consider such changes unfair (Akerlof 1980; Kahneman, Knetsch, and Thaler 1986; Akerlof and Yellen 1988, 1990) and react by shirking or leaving. Truman Bewley (1998, 475–76) summarized the belief about the potentially drastic consequences of pay changes as follows: "A firm would lose more money from the adverse effects of cutting pay than it would gain from lower wages and salaries. What restrained employers from cutting pay was the belief that doing so hurts morale and increases labor turnover."[2]

Despite the general agreement among economists that pay cuts can be prohibitively costly, firms do change their compensation schemes, but the real-world evidence on the effects is scarce. We investigate a unique setting that allows evaluation of the long-term effects of a pay change in which fixed wages were increased but bonuses decreased, thus resulting in a quasi experiment that reduced pay for the high performers and slightly increased it for the low performers, with an overall negative effect on pay. Workers' output in the affected divisions decreased by around 30%, and attrition increased. We document that effort decreased by the same order of magnitude as did absenteeism. Observing the effects over a period of more than 3 years, we show long-term negative reciprocity of those affected but find no negative selection effects of new hires.

The setting is as follows. Management of a global personnel search company decided to change the pay of some of its workers (in what follows, "consultants") in Germany. The company was operating under a number of different brand names; the respective divisions had initially operated in different markets (information technology [IT] specialists, accounting and finance, human resources [HR], engineering) and different sectors (e.g., banking, pharmaceutics), but increasingly their activities were converging. For these historical reasons, consultants' pay schemes had differed across divisions; more precisely, there was one group of divisions (group A, later also labeled the treatment group) with lower fixed pay and higher bonuses (on average)

[1] As modeled by the literature on specific knowledge (Jensen and Meckling 1992; Raith 2008).

[2] Campbell and Kamlani (1997) provide an overview of the theoretical literature, and Bewley (2007) provides an overview of the empirical work related to wage rigidity.

than the other divisions (group B, the control group).[3] Management decided to standardize compensation by adjusting the pay schemes of consultants in group A to match those for group B. Hence, the workers in the treatment group experienced a cut in bonuses and an increase in fixed wages, while compensation in the control divisions did not change. At a given effort and output, the change would result in a pay cut for the better-performing workers while the lower-performing workers would enjoy a pay increase, but the net effect over all workers would be negative.

The firm's case creates a quasi experiment to estimate the effects of a pay cut on worker effort and output. First, the change came as a surprise. The decision came from global top management and was not influenced by management in Germany. Also, it was announced only shortly before the pay change took effect. Second, we show that the trends of both treatment and control divisions were parallel before management's decision. Third, workers do not interact across divisions, minimizing the risk of spillovers. Fourth, workers were accustomed to volatile income streams. Fifth, they operated in an individualistic organizational culture in which we do not find evidence for spillovers. Thus, we believe that the experiment is not an extreme case and that the findings are likely to be lower- rather than upper-bound estimates for the effect of a pay cut on the behavior of workers.

We observe all offices in Germany before and after October 2009—when the compensation structure was harmonized across divisions—and over a period of almost 4 years, which makes this paper unique, not only for the study of pay cuts but also to gain insights into the long-run nature of negative reciprocity. The adjustment caused a permanent reduction in pay for the more productive employees but an increase for some of the less productive ones.

We find strong negative effects in employees' output—in the realm of 30% and more. The pay change also resulted in increased turnover of more productive employees and in more absenteeism. Furthermore, because the management information system (MIS) registered employees' activities, we can document the precise mechanisms of the output fall: the affected employees reduced effort along different dimensions of their activities, such as client meetings, vacancies, and candidates identified. We find no evidence that the performance of the historically less productive consultants increased.

These results are robust. There are sectoral differences between the two groups, but looking at IT specialists—the sector of the firm's activities that is most important for both treatment and control divisions—we find that differences disappear. Restricting the regressions to consultants who are active in hiring in IT-related fields, we observe similar regression results, providing an additional argument that our setting provides a useful quasi experiment. While the empirical strategy we use (discussed below) is robust to pretreatment

[3] In sec. II (in particular, fig. 1), we present the differences between the pay schemes of the two different groups of divisions.

differences that are not related to dynamics in outcome variables, we control for both pretreatment trends and sector-time fixed effects in our regressions.

Our results support arguments about the negative effects of a pay change in the literature (discussed below), but it also uncovers a surprise: we find no evidence that new entrants, after the pay change, are less productive than workers they replace. Hence, entrants are unaffected by potential long-term cultural effects of any mistrust that the pay change engendered. It also appears that the initial high bonuses were not necessary to attract high-quality workers; that is, we find no evidence for sorting of workers (Lazear 2000; Lazear and Shaw 2007). Indeed, we know that management's belief about the difference in the reward schemes contributing little to sorting of workers was an important factor in the decision to adjust the pay of the treated division (although the pay change was framed in a different way; see the next section). The absence of sorting effects can be explained by the fact that most hires had no experience in the job. Hence, they may have had little private information about their potential productivity.

There is overwhelming evidence that the introduction of well-designed pay-for-performance schemes increases firm profitability and workers' wages, but it also increases inequality between them (Prendergast 1999; Lemieux, MacLeod, and Parent 2009). Our paper, however, investigates how workers react to a reduction in the intensity of pay for performance. What we find in a real firm over a long horizon squares with a large literature in laboratory settings and short-term field experiments indicating that negative reciprocity tends to produce stronger reactions than positive reciprocity; some of these experiments are also looking at pay cuts and find adverse effects along a number of dimensions.[4] These papers have contributed much to understanding how workers react to a firm's activities that are perceived as unfair or unkind. We establish that these effects are large and that negative reciprocity in organizations is likely to be long lived. We observe people over a span of almost 4 years, and only after employees who experienced the pay change have left does the negative effect on output vanish. Our work thus fills a knowledge gap about the real-world effects of reciprocity in the workplace.

A number of different concepts of fairness could, in principle, play a role in explaining the reactions to a pay change. For instance, people may have fairness concerns relative to how others are treated, and they may feel better or worse if colleagues' wages are cut. The reference that seems to matter most in our case is, however, how one was treated in the past or what one believes one to be entitled to. This is reminiscent of studies showing that references and expectations about compensation matter for workers' willingness to act in the interest of their organization (Greenberg 1990; Mas 2006; Montizaan et al.

[4] Camerer and Weber (2013) review the earlier literature, mostly in labs. Recent lab-in-the-field experiments include Kube, Maréchal, and Puppe (2013), Cohn et al. (2014), Heinz et al. (2020).

2016). It should also be noted that the setting is quite similar to a relational contract (Gibbons and Roberts 2013) in which a principal may decide not to pay a promised bonus or renegotiate the contract and in which workers may carry out punishments off the equilibrium path. While this is consistent with the substantial reactions we document, we find no evidence that workers who do not lose because of the pay change carry out a punishment by reducing their performance, which is a theoretical possibility in relational contracting. We argue that this is due to the individualistic culture of the firm and, again, is indicative of fairness concerns about one's own treatment and not that of others.

Among the advantages of our setting are that the pay change was initiated by a headquarters from outside the country and individual panel data are observable on a monthly basis over several years. We observe the development of the effects on individuals and can disentangle posttreatment changes related to a decrease in productivity and those related to workforce changes, such as an increase in turnover and adverse selection. We are also in the unique situation that because output and employee activities are recorded, we have not only unambiguous individual profit but also effort measures for each employee. This personnel search firm also provides a perfect setting for a quasi-experimental study because employees work on their own and do comparable jobs. There are no risks of spillovers because employees do not interact across divisions.

Our results contrast with and complement two quasi-experimental papers on pay cuts in real organizations in different settings and with a different focus. Sandvik et al. (2021) find similar effects on attrition, no effort effects, and no evidence of a behavioral channel. These differences are likely to stem from the different contexts. In our setting, employees are given substantial discretion about their effort choices and have an average tenure of 18 months and career concerns,[5] while in call centers there is more control over employees' work and their horizons are shorter. Coviello, Deserranno, and Persico (2022) focus on exit or voice in Hirschmann's (1970) sense and find that some workers deliberately harm the firm after the pay cut (even if it harms them as well).[6]

We also investigate whether the reduction in performance is a rational reaction to a flatter compensation curve, in a way reminiscent of Lazear (2000) but

[5] Note that the presence of the MIS, in which activities and efforts are registered, does not mean that workers efforts are controlled. Ultimately, employees in our firm were retained and rewarded conditional on output, not input. Information about their activities was mainly used for training and feedback.

[6] We are aware of two nonexperimental papers: Greenberg (1990) looks at employee theft as a response to pay changes, and Lee and Rupp (2007) interpret delays of planes as pilots' reactions to pay cuts (but do not find much evidence of negative reactions). Although it does not examine a pay cut, the work of Mas (2006) is also related: he explores reactions to the outcomes of wage negotiations that do not meet police officers' reference points. Finally, Montizaan et al. (2016) look at the behavioral responses of workers on changes in pension rights.

with a different sign. A difference-in-differences regression excluding the consultants who experienced the pay cut yields insignificant results, making this explanation very unlikely. Jayaraman, Ray, and De Véricourt (2016) find that a pay raise is met by an initial reaction by employees that could be interpreted as reciprocity, but after a few months the workers' effort seems to be in line with a standard incentive model. We respond to their conclusion that more long-term studies are needed to investigate long-term reciprocity.

Despite the persistence of negative reciprocity that we document, it is not clear whether the pay change was, on balance, profitable. There is a large short-term loss, due to lower revenue on account of the pay change (in the realm of 24 months), but there is a long-term gain, because the firm hires equally productive workers at lower cost. A back-of-the-envelope calculation, however, suggests that the horizon would need to be more than 10 years for the pay change to be profitable, even at very optimistic assumptions about interest rates, firm reputation, and firm growth.[7]

We finally think that our paper provides in-depth evidence for one of the most important assumptions in many macro models, namely, downward wage rigidity. Going back to Kahn (1997), most of the literature (recently reviewed in Faia and Pezone 2019) has analyzed the institutional roots of downward rigidity. In many countries, though, white-collar workers do not enjoy as much protection as blue-collar workers do, so their pay—in particular the variable parts—can be adjusted downward. The workers we study are not unionized, and they are used to income volatility. Their fixed wage is even increased by the pay change, yet the effects we find are very large.

Our study complements a number of recent macro papers documenting wage cuts mainly in the context to macroeconomic downturns (Barattieri, Basu, and Gottschalk 2014; Sigurdsson and Sigurdardottir 2016; Grigsby, Hurst, and Yildirmaz 2021). These and other studies offer convincing empirical evidence on a large-scale level, but ours has the advantage of identifying causally the destructive effects of wage cuts on productivity, which we can measure precisely and over a long period of time. We believe that the personnel economics approach in our paper thus helps fill a gap in the literature and strongly supports assumptions of the macro literature, additionally because it looks at pay cuts through bonuses, which, for legal reasons, are easier to adjust downward.

II. The Firm, the Work, and the Workers

Our study firm is a personnel search firm. It operates in many countries and focuses on middle management and specialist vacancies in different industries. After entering Germany, the firm grew rapidly, opening around

[7] Moreover, there are important factors that we cannot price, such as (i) negative effects on the talent pool for team leaders and management positions and (ii) potential reputational losses for the firm.

18 offices in the biggest German cities; in each office, all workers belong to the same division. In total, there are four main divisions. Even though the divisions belong to the same parent company, offices are managed separately. Company headquarters is located outside Germany and supervises the German divisions; during the time of the pay change, it also maintained a few common support services for offices all over Europe.

Table 1 presents an overview of the firm's German operations, covering the entire span of the observation period from December 2008 to September 2012. The average division includes 58 employees in nonmanagement positions, who we will frequently refer to as "consultants"; the average office includes 25 consultants. We do not consider support staff and managers because they were not affected by the change and we have no performance data on them.

Consultants work individually with firms (the "clients") and job candidates to fill vacancies with matching workers. Upon successfully matching a candidate and a job, around 20%–30% of the annual income of the new jobholder is paid by the client to the consulting firm. The firm measures this as the "revenue" of an individual consultant and rewards the consultant accordingly (details in sec. III). We also use the number of deals (placements) made as an additional output measure. Within the divisions, some consultants fill permanent placements, and some place freelancers into projects (temporary placements). Around 68% of the consultants deal with permanent placements, and 32% with freelancers. While all of this has consequences for revenues, it is irrelevant for the purpose of this paper because the pay change affects the groups in similar ways and we can carry out fixed effects regressions. Consultants work for one division only, and there is no movement between the divisions at the consultant level.

Because the firm promotes only from within, consultants are normally hired after completion of their bachelor's degree, and most have no experience

Table 1
Overview: Consultants and Pay

	Mean	SD
Number of consultants per division	57.6	27.72
Number of consultants per office	25.41	14.52
Tenure in months	17.79	10.80
Percentage fixed pay per consultants	.63	.03
Male (percentage)	.53	.5
Year of birth	1980.48	13.65
German (percentage)	.85	.35
College degree (percentage)	.96	.20

NOTE.—This sample includes 572 employees (junior consultants, consultants, senior consultants) and covers a period of 46 months. For accounting reasons, a "month" can vary between 4 and 5 weeks. The unit of observation is an employee-month. "Tenure" shows the tenure in the last observed month in the sample. The full sample contains 8,936 employee-month observations.

in the job. Consultants' jobs are similar across the divisions, varying only by industry and regions. The divisions were set up to operate in different markets (e.g., pharmaceutical and medical companies vs. financial sector companies) and to specialize in different types of human resources (say, IT vs. HR). Increasingly, however, they began to operate in overlapping markets, although in separate offices and cities and with different brand names. This type of convergence in the job consultants carry out played a role in the decision of the firm's top management to standardize the compensation scheme across different divisions.

Management positions differ; they involve staffing, planning, training, and supervision of consultants. Typically, offices are supervised by regional managers. Within the offices, each team is supervised by a team manager who coaches the individual consultants by setting targets for output and activities. Promotion to a management position occurs after 30 months, on average; further career advances to positions of regional manager or in various support functions are possible. Consultants who leave the firm mostly stay in the personnel search business or related sales jobs. Average tenure is 18 months.

Output depends on the individual consultant's effort to fill a vacancy but is also determined by seasonal, sectoral, and regional variation. Effort entails different activities: meetings with clients as well as identifying new vacancies and new candidates. These activities are meticulously measured through an MIS. Consultants' output is recorded electronically through the MIS introduced to all German offices at the end of 2008. Since activities and output are regularly checked by the managers and the billing department, there are few measurement errors. Records are visible for the managers of each division. Consultants get to know their colleagues' performance in the same division because performance scores show the relative performance of employees within the division. Additionally, a score shows the relative performance of one's own office compared with the ones of other offices of the same division in Germany, Europe, and worldwide. Records of other divisions' employees are usually not visible except by upper management. In line with this highly transparent working environment and a competitive corporate culture, all offices are open spaced, similar to the workspace of financial traders.

III. Compensation and the Change in the Compensation Scheme

Consultants' pay consists of a fixed and a variable component, the latter representing, on average, one-third of total monthly compensation (as can be seen in table 1). Fixed pay on the nonmanagement level is, on average, around 2,500 euros per month. Consultants receive a commission conditioned on the pay the firm receives from clients, which we will refer to as "revenue." The precise computation of this commission is somewhat intricate, as it depends on the wage of the candidate the client hired, the timing of employment contracts between clients and candidates, and the distinction

between permanent and freelance positions. On average, around 10% of the revenue that accrues to the firm is directly paid to the responsible consultant, but the commission scheme is convex.

The firm's revenue from a deal depends on the hired candidates' annual wage and accrues 1 month after the new hire begins to work for the client. Consequently, if a consultant has placed more than one candidate and both happen to start working in the same month and given the convex nature of the commission, the consultant's commission is higher than if the starting date of the new hires were distributed evenly across a number of months. While a consultant's effort affects the number of candidates placed, they have little influence on start dates because they are not involved in the negotiations. The distinction between permanent and freelance contracts is also payoff relevant: while permanent contracts involve one payment only, freelance contracts involve a stream of monthly commissions over the duration of the freelancer/client relationship. Finally, there is a quarterly bonus, but it is paid out only to top performers and only very rarely (only 1.7% of the commission payments are quarterly bonuses).

Before the pay change in October 2009, divisions differed in their compensation structure, mostly for historical reasons, and, as pointed out before, the target groups/clients differed initially. Figure 1 compares the initial reward scheme of group A (the treatment group) with that of group B (the control group). Both groups include two divisions. Importantly, consultants in the control group saw no change, while those in the treatment group previously had slightly lower fixed wages but higher commission rates, with a notable jump at around 25,000 euros and higher pay beyond that number. The treatment adjusted their pay to that of group B. At a monthly revenue of zero euros, consultants in the control group would earn, on average, 118.5 euros more than consultants in the treatment group, while for higher monthly revenues, consultants in the treatment group would earn more (e.g., the pay difference for a monthly revenue of 55,000 euros is around 1,250 euros).

On October 1, 2009, the compensation scheme of the treatment group was adjusted to match the one of the control group. The pay change was first announced in August 2009 and became effective on October 1. As a consequence, the average expected pay of consultants in the treatment group—keeping revenue before and after the pay change constant—decreased by 3.5%. In figure 2, we provide the distributions of revenue (individual monthly performance) and pay in treatment and control groups before and after the pay change. One may expect a number of effects because of the pay change, most prominently, bunching in the treatment group around 25,000 euros of revenue before the pay change and more dispersion in worker output after the treatment around that mark. However, the data do not support bunching around 25,000 euros. Instead, it looks like there are more observations with zero revenue and less with high revenues in the treatment group after the pay change.

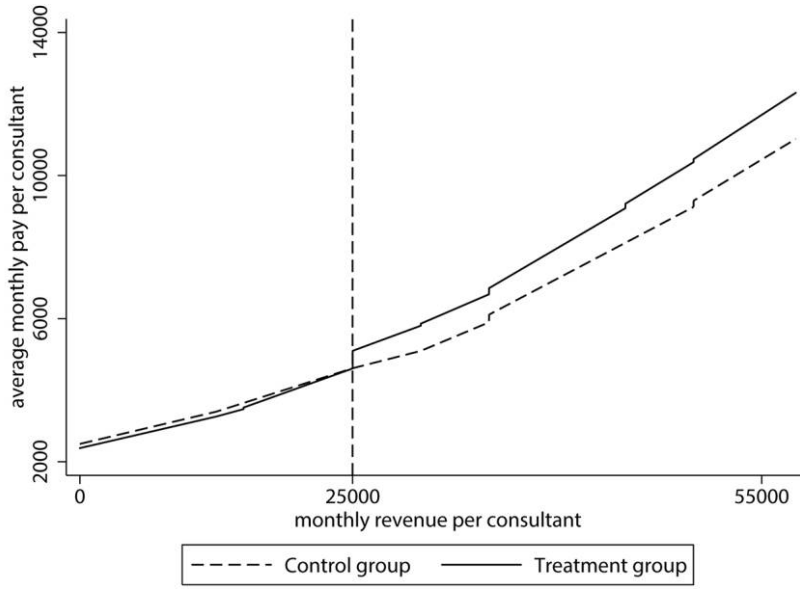

FIG. 1.—Compensation schemes for treatment and control groups before the pay change. This figure shows how performance (monthly generated revenue) relates to monthly pay for a consultant in the treatment and control groups before the pay change. Looking at a monthly revenue of zero euros, consultants in the control group would earn 118.5 euros more than consultants in the treatment group, while for higher monthly revenues, consultants in the treatment group would earn more (e.g., the pay difference for a monthly revenue of 55,000 euros is around 1,250 euros). The vertical line marks the threshold of 25,000 euros, beyond which the pay is higher in the treatment group compared with the control group before the pay change. After the pay change, the treatment group's pay scheme is identical to the control group's.

In an internal memo, top management communicated as the main reason for the change a strategic reorientation of the firm's operations and the need to harmonize HR policies across the four divisions: "The current compensation scheme focuses too much on short-term goals, rather than long-term career opportunities. . . . The new scheme puts all divisions in line with each other, making recruitment and mobility (between divisions) more effective."

Meetings were held in each office of the affected divisions, and employees thus learned the change could lead to a pay cut for consultants. However, career perspectives in the restructured firm and across divisions were also highlighted. Consultants in the treatment group were asked to sign amendments to their contracts. This created much discussion and frustration about the company (but not about other workers), in particular, among the top performers who felt deprived of the rewards for high performance. However, according to what we have learned, all consultants signed the new contracts. It may be surprising that nobody resisted the demand for the contract change

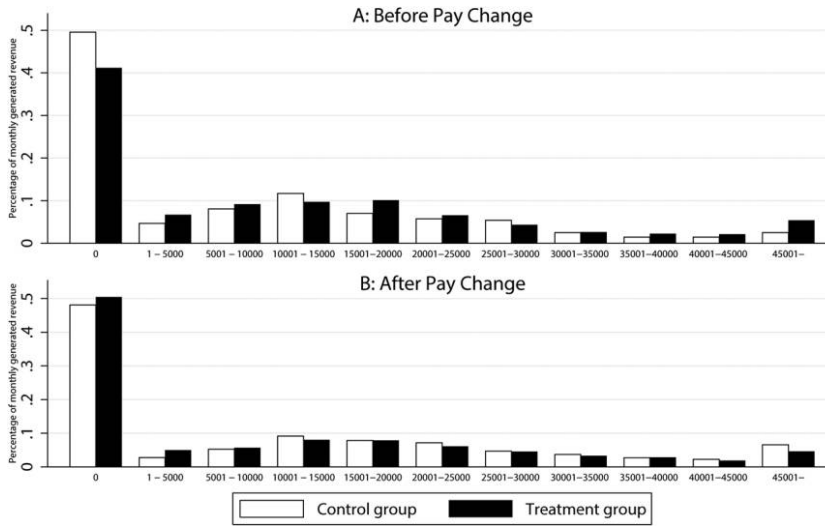

FIG. 2.—Dispersion of output before and after pay change, control versus treatment group. This figure plots the frequencies of monthly generated revenue before and after the pay change in steps of 5,000 euros.

because the pay stipulations were an explicit component of the initial contract. However, consultants may have expected adverse consequences from not signing and hence shied away from challenging the change in the firm or in court. Put differently, in a way similar to a relational contract, a promised bonus was renegotiated pointing to the incompleteness of the initial contract. In the Williamsonian sense, the source of this incompleteness would lie in incomplete enforcement. We cannot distinguish whether frustration was mainly triggered by this change in the rules of the game or the feeling of being deprived of a reward the consultants felt they had earned. Suffice it to say, the attempt of top management to highlight better promotion opportunities did little to mitigate this frustration, according to the anecdotal knowledge we have.

Consultants' expectations about their likely pay decrease would have depended on their expectations about future revenues, given their own performance records. Indeed, interviews in the firm revealed that when the new compensation scheme was announced, consultants in the treatment group calculated the pay change, depending on their past revenue. When we do the same, it becomes clear that roughly half of the consultants rationally expected cuts in pay, while the other half, because of the increase in fixed wages, would have stood to benefit. The largest effect, when holding constant the performance at the level prior to the pay change, would have been 5.4% for the most senior consultant level. Table 2 provides the expected pay change by pretreatment performance in the treatment group, which, by quintile,

Table 2
Expected Pay Difference Between Old and New Pay Scheme by Quintile of Pretreatment Performance (Consultant Monthly Averages)

Quintile	Mean Difference
1	−300.58
2	−94.45
3	−38.36
4	15.25
5	115.09

NOTE.—This table shows the average difference in pay between the old and the new pay scheme within the treatment group based on pretreatment performance. This sample only includes observations of consultants in the treatment group before the pay change. The rank (in quintiles) refers to the generated output (contract extensions for freelancers included) within the treatment group of the time period before the pay change. The first quintile represents the consultants with the highest performance.

covers a range from plus 115 euros for the lowest quintile to −300 euros for the highest quintile. While this does not seem dramatic, it turned out that because of the effort and output reduction we document below, consultants in the treatment group experienced a wage decrease of 19%, on average.

IV. Data and Research Design

Figure 3 depicts the structure of our data. In total, the data cover 572 consultants and 8,936 person/month observations. Our results are derived from (i) regressions using data on all employees and (ii) fixed effects regressions of the 128 employees who entered the firm before the pay change, dropped out of the sample after the pay change, and were working for either the treated divisions or the control division. In the graphic, this is represented by the arrow in the middle of the upper panel, while the arrow in the lower panel represents the 77 treated individuals. Consequently, there are 51 nontreated individuals.

Table 1 presents an overview of the personal characteristics of the consultants in treatment and control divisions and their locations. Table 3 presents statistics on the activities of the consultants as recorded by the firm's MIS. It is a unique feature of the data that the MIS provides details about consultants' output in terms of number of placements and total revenue as well as activities on the job. When collapsing the data on the consultant level, statistically significant differences exist concerning placements and revenue and concerning "candidates found." In all categories, the treatment group exhibits larger numbers. There are sectoral differences between the two groups (which we will control for in the regressions). Looking at the largest sector of the firm's activities, IT, we find that except for candidates found, the differences disappear. While the empirical strategy we use (discussed below) is robust to pretreatment differences that are not related to dynamics

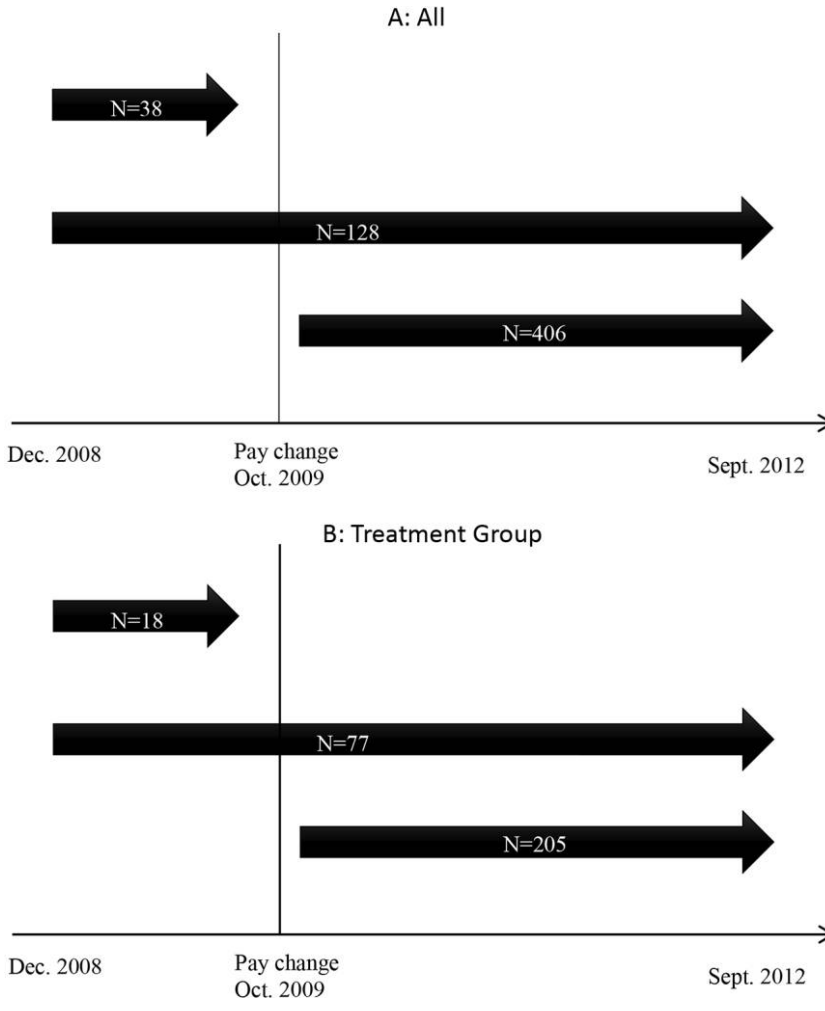

FIG. 3.—Data structure.

in outcome variables, we will later control for both pretreatment trends and sector-time fixed effects. Parallel trends are discussed in the next section.

We use a standard difference-in-differences design; in the robustness section, we also present an analysis that collapses the pre- and posttreatment periods to avoid potential problems of serial correlation (Bertrand, Duflo, and Mullainathan 2004). Our basic specification is

$$y_{ijt} = \alpha + \gamma_j + \lambda_t + \beta D_{jt} + \delta X_{ijt} + \varepsilon_{ijt}. \quad (1)$$

Here, γ_j is an indicator function that has a value of 1 for the treatment group, λ_t is an indicator function for posttreatment period, and D_{jt} is an indicator

Table 3
Activities

	Treatment Group		Control Group		Difference	
	Mean	SD	Mean	SD	Mean	SE
	Full Sample (December 2008–September 2009)					
Placements	.65	.48	.52	.41	.13	.07
Revenue	9,198.67	7,715.70	7,362.86	6,615.47	1,835.81	1,139.96
Vacancies found	4.30	2.61	4.00	3.84	.30	.50
Meetings scheduled	2.47	1.82	2.71	1.74	−.25	.28
Candidates found	7.25	5.30	4.01	4.55	3.24	.78
Absenteeism (days)	2.51	1.93	2.93	2.38	−.42	.34
Exit probability	.05	.16	.07	.17	−.02	.03
Number of consultants	95		71			
	Sector: IT (December 2008–September 2009)					
Placements	.65	.40	.53	.41	.12	.08
Revenue	8,681.55	6,463.49	7,120.80	6,040.16	1,560.74	1,291.28
Vacancies found	3.92	2.49	3.83	3.24	.09	.61
Meetings scheduled	2.84	1.72	2.66	1.69	.18	.35
Candidates found	6.46	4.01	3.62	4.25	2.84	.86
Absenteeism (days)	2.63	1.44	2.93	2.09	−.30	.38
Exit probability	.06	.17	.06	.12	.00	.03
Number of consultants	43		51			

NOTE.—This sample includes employees (junior consultants, consultants, senior consultants) in the pretreatment months (December 2008–September 2009). The data are collapsed on the consultant level. Sample IT includes only employees working in the IT sector.

for the change in the compensation scheme (i.e., the interaction of the two effects). To balance the pre- and posttreatment periods' durations and reflect employees' average tenure, we consider an event window of 20 months around the date of the pay change. Nevertheless, we also examine the entire 3-year posttreatment period to study the longer-term effects. Standard errors are clustered at the individual level to control for serial correlation problems (Bertrand, Duflo, and Mullainathan 2004).

We also run regressions with individual fixed effects, γ_i, and D_{it} represents the time-treatment effect on the individual level, as follows:

$$y_{ijt} = \alpha + \gamma_i + \lambda_t + \beta D_{it} + \delta X_{it} + \varepsilon_{it}. \quad (2)$$

Because this specification controls for all time-invariant person effects, comparison of the two regressions allows us to illuminate composition effects in the treatment group that are triggered by the exit of treated individuals and replacement by new recruits, similar to Lazear (2000).

V. Effects of the Pay Change on Revenue, Deals, and Retention

Figure 4 depicts the development of revenues and employee attrition over the observation period (December 2008 to September 2012) for all workers

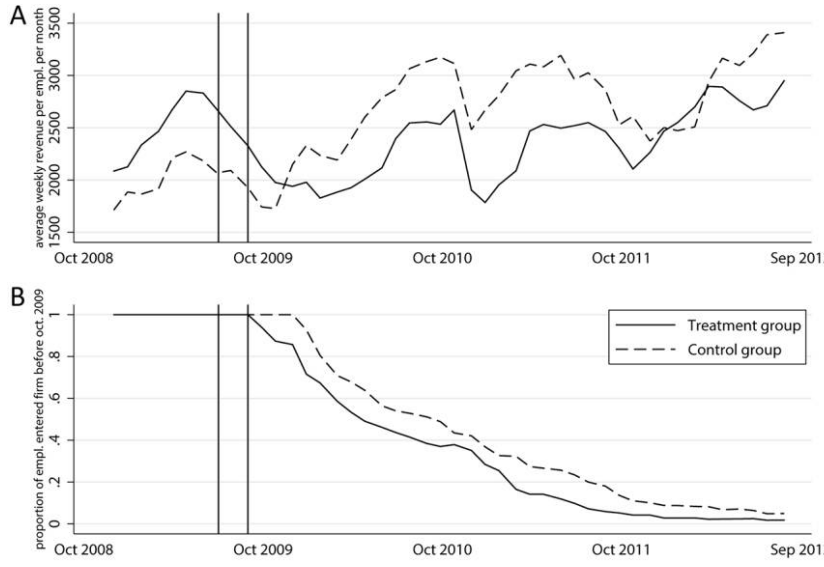

FIG. 4.—*A*, Revenues. *B*, Retention. *A* depicts 5-month moving averages of weekly generated revenue (without contract extensions) per month. Data cover the period from December 2008 to September 2012. *B* shows retention within the treatment and control groups. The vertical axis represents the proportion of employees who entered the firm before October 2009 and are still working as consultants in the firm at different points of time.

who entered the firm before the treatment and stayed at least until the moment of the pay change. Figure 5 does the same for the second outcome variable, number of deals.

The first vertical line in figures 4 and 5 indicates the date of the announcement of the pay change, while the second line indicates the date of its introduction. Pretreatment trends appear parallel in the control and treatment groups for revenues. (For attrition, however, they are identical by definition, because we look only at employees entering before and exiting after the pay change.)[8] The revenue trends diverge after the treatment. They remain parallel because of seasonal effects, but revenues are markedly higher in the control group (while the opposite is true before the treatment). Figure 4*B* complements this picture by showing that attrition was higher in the treatment than in the control group. When almost all employees of the two groups had left the firm or were promoted to being managers, differences in revenues seem to disappear. Looking at a figure 5, the same appears for number of deals. Figure 6 presents control and treatment group outcomes including a 95% confidence interval.

[8] Note, however, that in table 3, the exit probability (i.e., attrition of employees in both groups) is also identical when all employees are considered, including the ones leaving before the pay change.

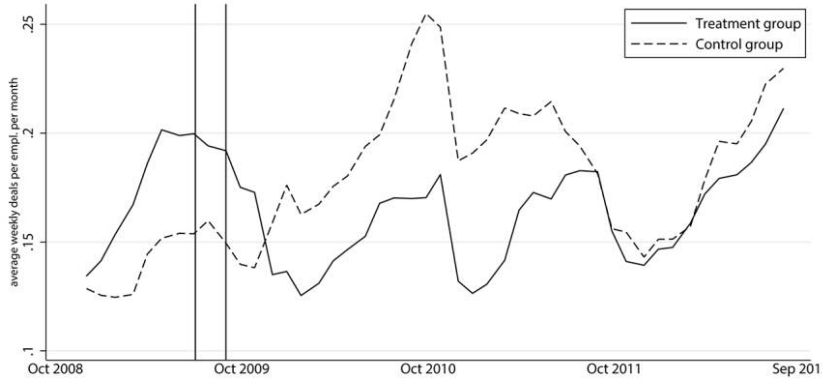

FIG. 5.—Deals. Lines show 5-month moving averages of weekly generated deals (without contract extensions) per month. Further explanations are found in the figure 4 legend.

Tables A3 and A4 (tables A1–A5 are available online) also summarize the results from a standard event study with a full set of mutually exclusive time categories interacted with an indicator for treatment group. We find insignificant coefficients on the prepolicy time categories.

Table 4 presents the results of the difference-in-differences regression (specification 1), using data for a total of 282 employees who either entered in a time window of 10 months before the treatment or 10 months afterward. We use these windows to have balanced pre- and posttreatment periods. We present the results for both revenue and deals (placements).[9]

Column 1 shows a strong effect of the after pay cut × treatment group interaction effect of an average of 3,600 euros (the pretreatment monthly average revenue is around 9,000 euros). Column 4 presents the same regression for number of deals. Column 2 (5, for deals) shows the results when tenure and tenure squared are introduced as control variables; while these variables are important in explaining revenues (the R^2 more than doubles), the point estimate of the after pay cut × treatment group variable stays at the same level. Adding separate trends for treatment and control groups also does not change the main results (not reported here, available on request). Because control and treatment groups are not fully balanced between sectors of activities, we also control for year × quarter × sector fixed effects (in cols. 3 and 6) that absorb the time-dependent influences of different markets. The point estimate is somewhat reduced. The second set of regressions uses deals as left-hand-side variables and produces similar results.

Table 5 presents fixed effects regressions on those employees who experienced the pay cut. Because consultants do not rotate over different offices

[9] In table 6 we provide regression results for the entire data set, and in sec. VII we present results by divisions, the level on which treatment occurs.

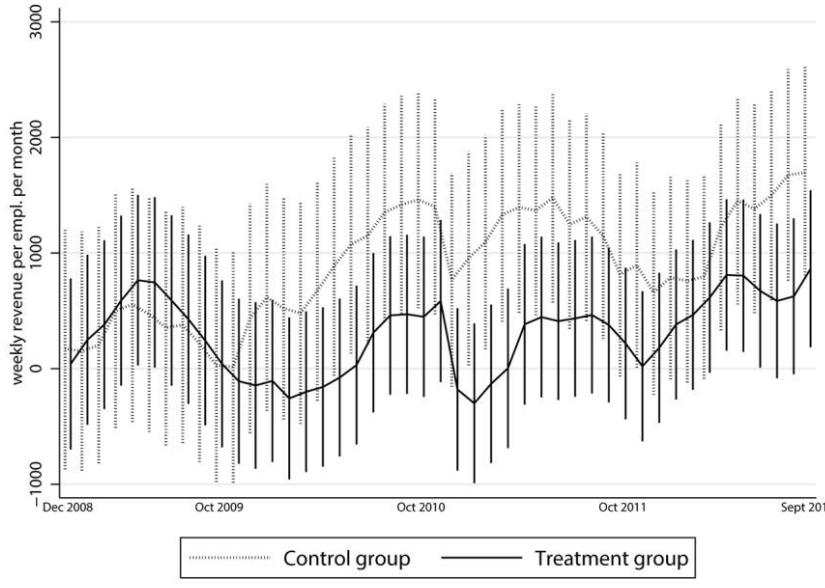

FIG. 6.—Revenue estimates over time. This figure shows regression coefficients and 95% confidence interval of time on 5-month moving averages of weekly generated revenues (without contract extensions) per month.

or sectors of activities, using person fixed effects precludes the use of city and sector fixed effects in table 5. Because we find results of the same order of magnitude for both revenue and deals in both regressions, sorting effects of the different bonus slopes do not seem to be important. There are two potentially related reasons for this observation. First, most recruits have no or very little experience in the job. Second, in another research project with the same firm (Friebel et al. 2019), we found that fresh recruits from another cohort entering the firm have very high levels of overconfidence. The average recruit believes that he or she will outperform 85% of the recruits in the same cohort. High confidence levels are negatively correlated with performance, indicating that people have little information about their productivity in the firm.

We now focus on the question of long-term consequences of the bonus cut. Table 6 presents the results of regressions in which we interact the treatment with three dummies representing different time spans of 1 year each after the treatment. Controlling for a battery of fixed effects, we find that for revenue the treatment effect is negative over 2 years, while the effect seemingly disappears in the third year. However, when looking only at consultants who experienced the pay change (in cols. 2 and 4) we find that revenues for the remaining consultants are lower than those for the entire sample. (There is no significant effect on number of deals in the second year.) We

Table 4
Difference-in-Differences Estimates of Employee Output

	Revenue/1,000 per Month per Employee			Deals per Month per Employee		
	(1)	(2)	(3)	(4)	(5)	(6)
After pay cut indicator	1.88	2.12*	2.36	.16	.17	.22
	(1.29)	(1.26)	(2.99)	(.13)	(.12)	(.38)
Treatment group indicator	1.96*	.90	.68	.14	.04	.03
	(1.15)	(1.08)	(1.14)	(.11)	(.11)	(.12)
After pay cut × treatment group	−3.58***	−3.12***	−2.58**	−.35***	−.31***	−.26**
	(1.24)	(1.18)	(1.28)	(.12)	(.11)	(.13)
Tenure (months)		1.01***	1.00***		.07***	.07***
		(.10)	(.10)		(.01)	(.01)
Tenure squared (months)		−.02***	−.02***		−.00***	−.00***
		(.00)	(.00)		(.00)	(.00)
City, sector, and demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Year × quarter fixed effects	Yes	Yes	No	Yes	Yes	No
Year × quarter × sector fixed effects	No	No	Yes	No	No	Yes
N	3,019	3,019	3,019	3,019	3,019	3,019
R^2	.03	.10	.11	.11	.19	.20

NOTE.—This sample includes 282 employees (junior consultants, consultants, senior consultants) 10 months before and 10 months after the implementation of the compensation change. All regressions include controls for month lengths, training days, and whether the employee deals with freelancers or permanent placements. "City" refers to the location of the office where a consultant is located, and "sector" refers to the main sector of activity. Dependent variables are revenue/1,000 and placements per month per employee (contract extensions for freelancers included). Standard errors in parentheses are clustered on the individual level.

* $p < .10$.
** $p < .05$.
*** $p < .01$.

interpret this as evidence for the effect at least being persistent over 2 years after the treatment and would carefully argue that there is, even after 3 years, a certain negative effect. Indeed, only a small number of individuals who experienced the treatment were still working as a consultant in the firm after 3 years (as depicted in figs. 3, 4).

Table 7 differentiates the regressions by "winners" and "losers" of the pay change. Workers could, at the outset of the new pay system, compute whether, at given historical performance, they would experience a pay cut. We label these the losers of the pay change. When running the regressions on output separately for these groups, it becomes clear that the negative effects are driven by those experiencing a pay cut, while we do not see any sign of increased performance of those who received an increase.

Figure 7 plots the tenure profile of treatment and control group consultants by pretreatment performance and shows that for the control group tenure is higher for high performers, while this is not the case for the treatment group. This motivated the regression whose results can be found in table 8,

Table 5
Difference-in-Differences Estimates of Output for Employees Who Entered Before October 2009

	Revenue/1,000 per Month per Employee		Deals per Month per Employee	
	(1)	(2)	(3)	(4)
After pay cut indicator	2.76**	1.08	.23*	.00
	(1.34)	(1.43)	(.13)	(.12)
After pay cut × treatment group	−3.69***	−3.59***	−.30**	−.29**
	(1.36)	(1.36)	(.13)	(.13)
Tenure (months)		2.42**		.30***
		(1.04)		(.11)
Tenure squared (months)		−.01***		−.00
		(.00)		(.00)
Individual fixed effects	Yes	Yes	Yes	Yes
Year × quarter fixed effects	Yes	Yes	Yes	Yes
N	2,166	2,166	2,166	2,166
R^2	.02	.03	.03	.04

NOTE.—This sample includes 128 employees (junior consultants, consultants, senior consultants) who entered the firm before October 2009 in a window of 10 months before and 10 months after the implementation of the compensation change. All regressions include controls for month lengths and training days. Dependent variables are revenue/1,000 and placements per month per employee (contract extensions for freelancers included). Standard errors in parentheses clustered on individual level.

* $p < .10$.
** $p < .05$.
*** $p < .01$.

in which the separation rate is the dependent variable. The interaction term shows that the treated consultants are more likely to leave.

An interesting question is which type of consultant is driving this result. Figure 7 plots the tenure profile by sales rank of consultants separately for the treatment and control groups. The first percentile represents the consultants with the highest performance. As expected, in the control group tenure increases in performance. However, this is not the case in the treatment group, indicating that the successful consultants are leaving from the treated divisions. For illustrative purposes, the graph also provides a simple linear fit of the observations in the two groups.

VI. Effect on Consultants' Activities

It is a particular advantage of our data that the MIS records the most important activities of the consultants. These are job vacancies found, meetings held, and candidates found. The system also recorded consultants' days of absence. Table 9 collects the results of a fixed effects regression on these four direct measures of consultants' behavior. The picture seems to be very much consistent with the regressions in section V: treated consultants have 1.4 (32%) fewer vacancies found. They have 0.74 (30%) fewer meetings with clients. They find two (27%) candidates fewer. And they roughly take 0.8 (32%) more days of absence. Notice that these effects are around one-third

Table 6
Difference-in-Differences Estimates of Employee Output (Development over 3 Years of Posttreatment Period)

	Revenue/1,000 per Month per Employee		Deals per Month per Employee	
	(1)	(2)	(3)	(4)
Treatment group indicator	.39		.02	
	(1.17)		(.12)	
Oct. 2009–Sept. 2010 × treatment group	−2.93**	−4.26***	−.29**	−.38***
	(1.35)	(1.39)	(.13)	(.13)
Oct. 2010–Sept. 2011 × treatment group	−2.91*	−9.33***	−.17	−.68***
	(1.72)	(2.56)	(.16)	(.23)
Oct. 2011–Sept. 2012 × treatment group	−1.50	−3.46	−.01	−.11
	(1.61)	(3.76)	(.14)	(.43)
Tenure (months)	1.07***	2.49***	.08***	.31***
	(.08)	(.89)	(.01)	(.09)
Tenure squared (months)	−.02***	−.01***	−.00***	−.00
	(.00)	(.00)	(.00)	(.00)
City, sector, and demographic controls	Yes	No	Yes	No
Year × quarter × sector fixed effects	Yes	No	Yes	No
Individual fixed effects	No	Yes	No	Yes
Year × quarter fixed effects	No	Yes	No	Yes
N	8,923	2,908	8,923	2,908
R^2	.10	.05	.19	.06

NOTE.—The sample in rows 1 and 3 includes 569 employees (junior consultants, consultants, senior consultants) 10 months before and 36 months after the implementation of the compensation change. All regressions include controls for month lengths, training days, and whether the employee deals with freelancers or permanent placements. "City" refers to the location of the office where a consultant is located, and "sector" refers to the main sector of activity. The sample in rows 2 and 4 includes 128 employees (junior consultants, consultants, senior consultants) who entered the firm before October 2009 in a window of 10 months before and 36 months after the implementation of the compensation change. All regressions include controls for month lengths and training days. Dependent variables are revenue/1,000 and placements per month per employee (contract extensions for freelancers included). Standard errors in parentheses are clustered on the individual level.

* $p < .10$.
** $p < .05$.
*** $p < .01$.

for each of the activities, which is commensurate to the effects on revenues and deals. We have run the same regressions on a 36-month window around the event and find quite similar results (table A1).

VII. Further Results and Robustness

In table 10, we present a number of additional regressions. As an important robustness check, in columns 1 and 5 we collapse each consultant's revenue and number of deals into one pretreatment and one posttreatment observation to deal with serial correlation concerns, as recommended by Bertrand, Duflo, and Mullainathan (2004). In columns 2 and 6, we exclude the period between announcement and implementation of the pay change to deal with the possibility that consultants may have reacted even before the pay change affected their wage.

Table 7
Difference-in-Differences Estimates: Winners versus Losers of the Pay Change

	Revenue/1,000 per Month per Employee				Deals per Month per Employee			
	Losers		Winners		Losers		Winners	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
After pay cut indicator	3.78**	1.84	1.40	.64	.29**	.05	.12	−.04
	(1.58)	(1.56)	(1.38)	(1.50)	(.14)	(.13)	(.12)	(.14)
After pay cut × treatment group	−8.59***	−8.49***	.07	−.37	−.72***	−.72***	.00	−.01
	(1.43)	(1.43)	(1.53)	(1.61)	(.14)	(.14)	(.17)	(.17)
Tenure (months)		2.64**		1.37		.30***		.22**
		(1.04)		(.96)		(.11)		(.10)
Tenure squared (months)		−.01*		−.01*		.00		.00
		(.00)		(.00)		(.00)		(.00)
Individual fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year × quarter fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	2,163	2,163	1,992	1,992	2,163	2,163	1,992	1,992
R^2	.04	.05	.07	.07	.05	.06	.08	.08

NOTE.—This sample includes 128 employees (junior consultants, consultants, senior consultants) who entered the firm before October 2009 in a window of 10 months before and 10 months after the implementation of the compensation change. The "loser" sample includes 94 consultants (51 consultants from the control group and 43 consultants from the treatment group), while the "winner" sample includes 85 consultants (51 consultants from the control group and 34 consultants from the treatment group). The loser sample includes all consultants from the treatment group who given their pretreatment performance would get less pay under the pay scheme implemented in October 2009. The winner sample includes all consultants from the treatment group who given their pretreatment performance would get the same amount or more pay under the pay scheme implemented in October 2009. All regressions include controls for month lengths and training days. Dependent variables are revenue/1,000 and placements per month per employee (contract extensions for freelancers included). Standard errors in parentheses clustered on individual level.

* $p < .10$.
** $p < .05$.
*** $p < .01$.

Columns 3 and 7 are particularly interesting for appreciating the validity of the data on activities. Excluding the observations of consultants leaving in the first 10 months after the pay change (i.e., until September 2010), we find little changes in the effect of the treatment on performance. This indicates that the effort response documented is not simply due to the fact that consultants who anticipate their quitting reduce their activities. Moreover, as shown in figure 6, it is in particular the most productive employees who leave first. The results hence indicate that the treatment effect is not entirely driven by the quitting behavior of the most productive employees.

Columns 4 and 8 are based on observations for consultants active in the IT sector only, one of the biggest sectors of activity in the firm. The point estimates are quite similar, lending support to our identification strategy: the reactions we observe do not seem to be driven by sector specificities of the

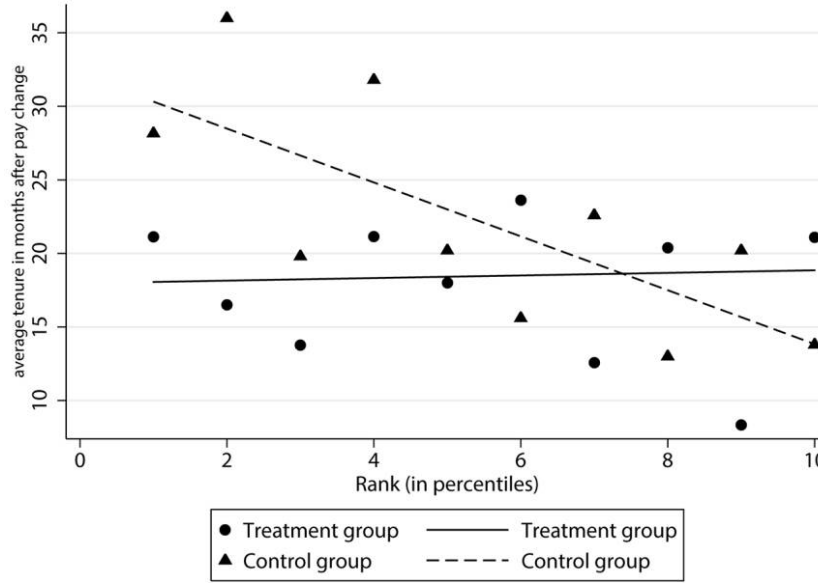

FIG. 7.—Average tenure by revenue rank. Average tenure after the pay change refers to the period in months from the pay change in October 2009 until the exit period or the last observed period of the data set (September 2012). The rank (in percentiles) refers to the generated revenue/1,000 (contract extensions for freelancers included) within the treatment or control group before the pay change. The first percentile represents the consultants with the highest performance.

client base or idiosyncratic shocks in sectors. Figure 8, the equivalent of figure 4 but for the IT sector only, provides additional support.

Finally, because treatment occurred on the divisional level, we have run the regressions separately for the two treated divisions (table A2). While the divisions are certainly not homogenous, the effect of the treatment in the difference-in-differences specification used is of similar magnitude.

VIII. Mechanisms

A possible interpretation of our results goes as follows. When the firm decides to reduce the bonus rate, consultants rationally adjust their behavior by reducing their effort, which translates into lower output. The magnitudes we discussed in the preceding sections do not seem to be consistent with such an interpretation, but the data allow some additional analyses. We carry out a difference-in-differences analysis on entrants in the treatment and control divisions before the pay change, in which the treatment group had higher incentives compared with entrants after the pay change, when all consultants worked with the same compensation scheme. Crucially, for this analysis

Table 8
Cox Proportional Hazard Model: Employee Separation Rates

	(1)	(2)
After pay cut indicator	.69**	.79**
	(.34)	(.36)
Treatment group indicator	−.38	−.17
	(.36)	(.36)
After pay cut × treatment group	.77*	.72*
	(.42)	(.42)
City, sector, and demographic controls	No	Yes
N	295	292

NOTE.—This sample includes 166 employees (junior consultants, consultants, senior consultants). "City" refers to the location of the office where a consultant is located, and "sector" refers to the main sector of activity. Observations are split into before and after pay cut time intervals. The results are reported as coefficients rather than hazard rates. Robust standard errors are in parentheses.
* $p < .10$.
** $p < .05$.

we exclude observations for the period after the pay change of those consultants who experienced the pay change.

Figure 9 illustrates which data from subgroups we use. The first group (in the upper left corner, $N = 38$) consists of consultants who entered and left the firm before the pay change. The second group (lower left corner, $N = 128$) consists of consultants who entered the firm before the pay change and left afterward. Of this group, 77 were affected by the pay change. We use the data on these 128 consultants only up to the moment the pay change was introduced. (Hence, we exclude the lighter arrow.) The third group consists of consultants who entered the firm only after the pay change, all of whom work under the same compensation contract. For this third group, we consider only those entering after the pay change.

If the reduction of the bonus per se had a substantial effect in explaining our main results, one should expect a strong negative coefficient of being in the treatment group in a difference-in-differences analysis of the type carried out before on the full data set. However, table 11 shows that there is no such effect. In particular, when controlling for sectors interacted with time, it is reduced remarkably and far from any conventional level of statistical significance. Note that in section VI, when comparing fixed effects regressions with ordinary least squares, we did not find evidence for sorting. Hence, it seems very unlikely that the strong effects we find would be caused by some sorting that would result in the individuals in the treated divisions having substantially different disutility functions than the other employees. While we cannot entirely exclude sorting effects, we see these results as supporting the interpretation of our main regression results that there is more happening than a simple reaction of effort supply to a changed incentive slope.

Table 9
Difference-in-Differences Estimates of Employee Activities for Employees Who Entered Before October 2009 (−10 to +10 Months Event Time Window)

	Vacancies Found		Meetings Scheduled		Candidates Found		Absenteeism (Days)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
After pay cut indicator	.76*	.52	.28	.36	.40	.60	.14	−.23
	(.42)	(.45)	(.35)	(.38)	(.59)	(.61)	(.45)	(.45)
After pay cut × treatment group	−1.37***	−1.36***	−.74**	−.72**	−1.97***	−1.97***	.77**	.80**
	(.47)	(.47)	(.30)	(.30)	(.63)	(.64)	(.37)	(.36)
Percentage change	−32	−32	−30	−29	−27	−27	−31	−32
Tenure (months)		.32**		−.05		−.25		.53**
		(.15)		(.16)		(.19)		(.23)
Tenure squared (months)		−.00		−.00**		−.00		−.00***
		(.00)		(.00)		(.00)		(.00)
Individual fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year × quarter fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	2,166	2,166	2,166	2,166	2,166	2,166	2,166	2,166
R^2	.07	.07	.06	.07	.11	.11	.03	.04

NOTE.—This sample includes 128 employees (junior consultants, consultants, senior consultants) who entered the firm before October 2009 in a window of 10 months before and 10 months after the implementation of the compensation change. All regressions include controls for month lengths and training days. Percentage changes serve interpretation purposes and are calculated on the basis of the numbers in row 1 of table 3 (e.g., 1.37/4.3 = 0.32). Dependent variables are meetings scheduled, vacancies, candidates found, and absenteeism (days) per month per employee. Standard errors in parentheses are clustered on individual level.

* $p < .10$.
** $p < .05$.
*** $p < .01$.

According to the fair wage-effort hypothesis (Akerlof and Yellen 1990), workers who perceive wage cuts as unfair react by withholding effort and leaving with a higher probability. What we find comports with this, and it also comports with the view of Hart and Moore (2008) and Halonen-Akatwijuka and Hart (2013) that contracts act as reference points, deviations from which are penalized by shading. Our results support the view brought forward by a large literature that shows, in laboratory settings, that negative reciprocity can have strong effects, but our setting offers the unique opportunity to show that this holds in the field over a long period.

We cannot exclude either that the reactions we document constitute a rational penalty in the framework of a relational contract between one worker

Table 10
Robustness of Event Study Estimates of Employee Output (−10 to +10 Months Event Time Window)

	Revenue/1,000 per Month per Employee			
	(1)	(2)	(3)	(4)
After pay cut indicator	−2.63	−3.43**	2.03	.20
	(3.57)	(1.65)	(1.69)	(1.88)
After pay cut × treatment group	−3.30**	−2.92*	−3.95**	−3.78**
	(1.34)	(1.57)	(1.54)	(1.67)
Tenure (months)	1.10**	1.06***	2.59**	3.46**
	(.44)	(.24)	(1.27)	(1.52)
Tenure squared (months)	−.01**	−.01***	−.01***	−.01*
	(.00)	(.00)	(.00)	(.00)
Collapse pre- and postperiod	Yes	No	No	No
Exclude September 2009/October 2009	No	Yes	No	No
Exclude employees left before September 2010	No	No	Yes	No
IT jobs only	No	No	No	Yes
Individual fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	No	Yes	No	No
Year × quarter fixed effects	No	No	Yes	Yes
Number of employees in sample	128	128	92	69
N	256	1,914	1,663	1,187
R^2	.17	.03	.04	.03
	Deals per Month per Employee			
	(5)	(6)	(7)	(8)
After pay cut indicator	.04	−.05	.06	−.10
	(.37)	(.15)	(.15)	(.16)
After pay cut × treatment group	−.26*	−.29*	−.32**	−.44**
	(.13)	(.15)	(.15)	(.18)
Tenure (months)	.04	.07***	.32**	.52***
	(.05)	(.02)	(.13)	(.17)
Tenure squared (months)	−.00	−.00	−.00	.00
	(.00)	(.00)	(.00)	(.00)
Collapse pre- and postperiod	Yes	No	No	No
Exclude September 2009/October 2009	No	Yes	No	No
Exclude employees left before September 2010	No	No	Yes	No
IT jobs only	No	No	No	Yes
Individual fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	No	Yes	No	No
Year × quarter fixed effects	No	No	Yes	Yes
Number of employees in sample	128	128	92	69
N	256	1,914	1,663	1,187
R^2	.13	.04	.05	.05

NOTE.—This sample includes employees (junior consultants, consultants, senior consultants) who entered the firm before October 2009 in a window of 10 months before and 10 months after the implementation of the compensation change. All regressions include controls for month lengths and training days. Dependent variables are revenue/1,000 and placements per month per employee (contract extensions for freelancers included). Standard errors in parentheses are clustered on the individual level.

* $p < .10$.
** $p < .05$.
*** $p < .01$.

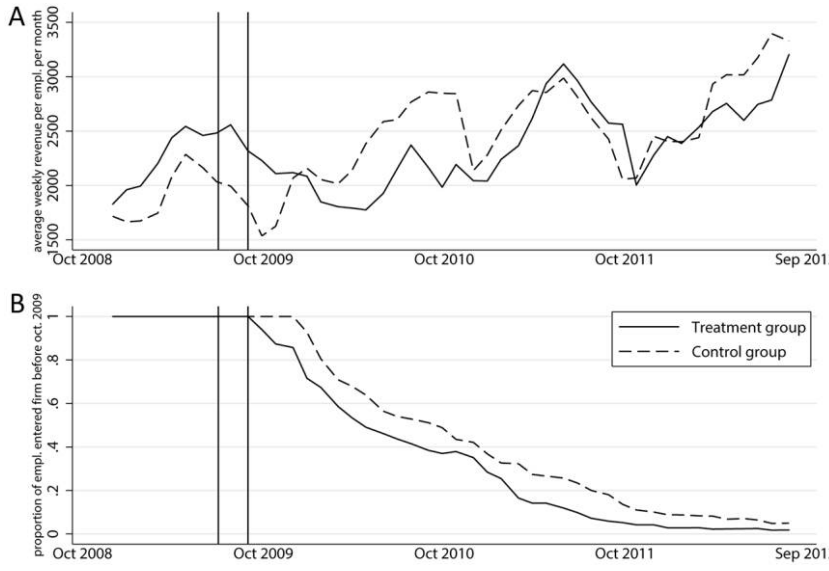

FIG. 8.—*A*, Revenues. *B*, Retention for IT only. *A* depicts 5-month moving averages of weekly generated revenue (without contract extensions) per month. Data cover the period from December 2008 to September 2012. *B* shows retention within the treatment and control groups. Both refer to consultants working with IT vacancies only. The vertical axis represents the proportion of employees who entered the firm before October 2009 and are still working as consultants in the firm at different points of time.

and the firm (for an extensive discussion, see Gibbons and Roberts 2013). In particular, higher attrition of the concerned workers is compatible with a grim-trigger strategy, in which the relationship is dissolved after the deviation of the principal. However, we do not find any evidence for multilateral

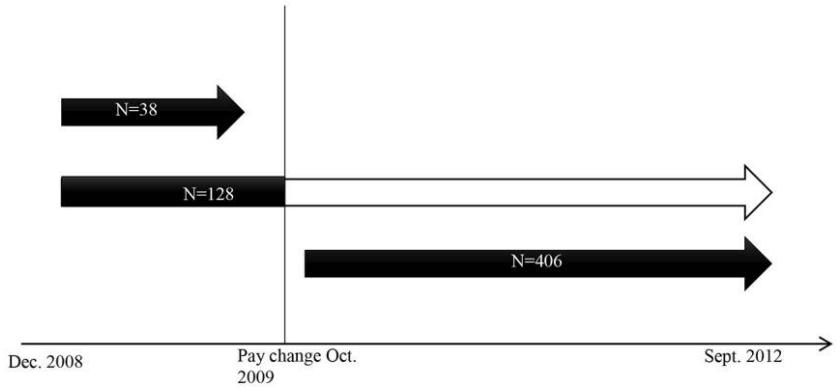

FIG. 9.—Data structure for incentive effect estimates.

Table 11
Difference-in-Differences Estimates: Effect of Pay Change on New Entrants

	Revenue/1,000 per Month per Employee		Deals per Month per Employee	
	(1)	(2)	(3)	(4)
After pay cut indicator	1.80	1.88	−.09	−.29
	(1.61)	(2.94)	(.15)	(.31)
Treatment group indicator	.92	.39	.10	.02
	(1.03)	(1.15)	(.11)	(.12)
After pay cut × treatment group	−2.04	−1.47	−.17	−.07
	(1.25)	(1.45)	(.12)	(.13)
Tenure (months)	1.57***	1.55***	.12***	.12***
	(.10)	(.10)	(.01)	(.01)
Tenure squared (months)	−.04***	−.04***	−.00***	−.00***
	(.00)	(.00)	(.00)	(.00)
City, sector, and demographic controls	Yes	Yes	Yes	Yes
Year × quarter fixed effects	Yes	No	Yes	No
Year × quarter × sector fixed effects	No	Yes	No	Yes
N	7,057	7,057	7,057	7,057
R^2	.11	.11	.17	.18

NOTE.—This sample includes 567 employees (junior consultants, consultants, senior consultants) 10 months before and 36 months after (only new entrants) the implementation of the compensation change. All regressions include controls for month lengths, training days, birth year, and whether the employee deals with freelancers or permanent placements. "City" refers to the location of the office where a consultant is located, and "sector" refers to the main sector of activity. Dependent variables are revenue/1,000 and placements per month per employee (contract extensions for freelancers included). Standard errors in parentheses are clustered on the individual level.

*** $p < .01$.

relational contracting in which pay deviations by the firm would not be penalized only by the affected worker but also by his or her colleagues who have not been harmed or by workers entering the firm afterward, who would provide lower effort because the firm has lost its reputation of maintaining bonuses. Indeed, we showed before that entrants into the treated divisions are, if anything, more productive than entrants in the nontreated divisions, showing that there is no collective penalty from the treated divisions; neither is there a cultural spillover. Table 11 also supports this interpretation: entrants who arrived after the pay change have no differential performance between treated and nontreated divisions.

IX. Concluding Remarks

The effects of the pay change are large: turnover increased in the treated group, and productivity dropped in the realm of one-third of the otherwise-to-be-expected sales in the treatment divisions. However, a priori we cannot exclude the possibility that the pay change was a profitable project for the firm, given the long-term perspective of owners. Not only were the costs of the pay change drastic, but they were quite short lived, because they

disappeared with the attrition of the consultants affected. A back-of-the-envelope calculation based on the wage savings of incumbent consultants and all consultants hired after the pay change reveals that the cost savings would outweigh the negative revenue effects only after roughly 10 years. This computation is based on the firm's employment growth of 10% per year and, to stack the deck in favor of the pay cut, an interest rate of only 1%. It is unlikely that management would have such a long horizon and such low interest rates, making it more likely that management underestimates the negative effects that we document.

Somewhat surprising is that the firm could recruit productive workers in the aftermath of the pay change. Consultants entering after the pay change were "cheaper" but as productive as those they replaced. This may not have come as a surprise for the company's management because, in the pretreatment period, the performance of consultants was quite similar in the two groups, despite the difference in bonuses. So management may have had correct expectations about some, but not all, parameters of the decision.

It seems puzzling that we find no evidence for negative sorting after the pay change, but most rookies know little about their suitability for a new job. In another, more recent study with the firm (Friebel et al. 2019), we found very high levels of overconfidence among fresh recruits. Upon entry, many consultants take a rather long period before concluding their first placement, and some leave even before such a placement happens. There thus seems to be some uncertainty among potential employees about their match with the job. Additionally, even after the pay change, the firm still offered quite competitive packages compared with the rest of the market.

The picture we have tried to paint is nuanced: the work culture in our firm is a highly individualistic one in which only the affected workers react. The firm is used to high turnover, and workers are used to volatile income. It is hence likely that what we found is not an extreme case of the likely effects that econometricians could find if data sets similar to the one we used were available. Potential costs of a pay change, beyond what we identified in our setting, would be the loss of managerial human capital and firm reputation. Taken together, the conjecture of a large literature that pay cuts are bad for firm performance seems to be well supported by our study.

References

Akerlof, George A. 1980. A theory of social custom, of which unemployment may be one consequence. *Quarterly Journal of Economics* 94: 749–75.

Akerlof, George A., and Janet L. Yellen. 1988. Fairness and unemployment. *American Economic Review* 78:44–49.

———. 1990. The fair wage-effort hypothesis and unemployment. *Quarterly Journal of Economics* 105:255–83.

Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk. 2014. Some evidence on the importance of sticky wages. *American Economic Journal: Macroeconomics* 6, no. 1:70–101.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics* 119, no. 1:249–75.

Bewley, Truman F. 1998. Why not cut pay? *European Economic Review* 42:459–90.

———. Fairness, reciprocity, and wage rigidity. In *Behavioral economics and its applications*, ed. Peter Diamond and Hannu Vartiainen, 157–88. Princeton, NJ: Princeton University Press.

Camerer, Colin, and Roberto Weber. 2013. Experimental organizational economics. In *The handbook of organizational economics*, ed. Robert Gibbons and John Roberts, 213–62. Princeton, NJ: Princeton University Press.

Campbell, Carl M., III, and Kunal S. Kamlani. 1997. The reasons for wage rigidity: Evidence from a survey of firms. *Quarterly Journal of Economics* 112:759–89.

Cohn, Alain, Ernst Fehr, Benedikt Herrmann, and Frédéric Schneider. 2014. Social comparison and effort provision: Evidence from a field experiment. *Journal of the European Economic Association* 12:877–98.

Coviello, Decio, Erika Deserranno, and Nicola Persico. 2022. Counterproductive worker behavior after a pay cut. *Journal of the European Economic Association* 20, no. 1:222–63.

Faia, Ester, and Vincenzo Pezone. 2019. Monetary policy and the cost of wage rigidity: Evidence from the stock market. SAFE Working Paper no. 242, Leibniz Institute for Financial Research SAFE, Frankfurt am Main.

Friebel, Guido, Anna Daelen, Matthias Heinz, and Nick Zubanov. 2019. Self-selection into personnel consulting. Unpublished manuscript.

Gibbons, Robert, and John Roberts, ed. 2013. *The handbook of organizational economics*. Princeton, NJ: Princeton University Press.

Greenberg, Jerald. 1990. Employee theft as a reaction to underpayment inequity: The hidden cost of pay cuts. *Journal of Applied Psychology* 75, no. 5:561–68.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz. 2021. Aggregate nominal wage adjustments: New evidence from administrative payroll data. *American Economic Review* 111, no. 2:428–71.

Halonen-Akatwijuka, Maija, and Oliver D. Hart. 2013. More is less: Why parties may deliberately write incomplete contracts. NBER Working Paper no. 19001, National Bureau of Economic Research, Cambridge, MA.

Hart, Oliver, and John Moore. 2008. Contracts as reference points. *Quarterly Journal of Economics* 123:1–48.

Heinz, Matthias, Sabrina Jeworrek, Vanessa Mertins, Heiner Schumacher, and Matthias Sutter. 2020. Measuring the indirect effects of adverse

employer behaviour on worker productivity: A field experiment. *Economic Journal* 130, no. 632:2546–68.
Hirschmann, Albert O. 1970. *Exit, voice and loyalty*. Cambridge, MA: Harvard University Press.
Jayaraman, Rajshri, Debraj Ray, and Francis De Véricourt. 2016. Anatomy of a contract change. *American Economic Review* 106, no. 2:316–58.
Jensen, Michael, and William Meckling. 1992. Specific and general knowledge and organizational structure. In *Contract economics*, ed. L. Werin and H. Wijkander, 251–74. Oxford: Blackwell.
Kahn, Shulamit. 1997. Evidence of nominal wage stickiness from microdata. *American Economic Review* 87:993–1008.
Kahneman, Daniel, Jack L. Knetsch, and Richard Thaler. 1986. Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review* 76:728–41.
Kube, Sebastian, Michel André Maréchal, and Clemens Puppe. 2013. Do wage cuts damage work morale? Evidence from a natural field experiment. *Journal of the European Economic Association* 11:853–70.
Lazear, Edward P. 2000. Performance pay and productivity. *American Economic Review* 90, no. 5:1346–61.
Lazear, Edward P., and Kathryn L. Shaw. 2007. Personnel economics: The economist's view of human resources. *Journal of Economic Perspectives* 21:91–114.
Lee, Darin, and Nicholas G. Rupp. 2007. Retracting a gift: How does employee effort respond to wage reductions? *Journal of Labor Economics* 25, no. 4:725–61.
Lemieux, Thomas, W. Bentley MacLeod, and Daniel Parent. 2009. Performance pay and wage inequality. *Quarterly Journal of Economics* 124, no. 1:1–49.
Mas, Alexandre. 2006. Pay, reference points, and police performance. *Quarterly Journal of Economics* 121:783–821.
Montizaan, Raymond, Andries de Grip, Frank Cörvers, and Thomas Dohmen. 2016. The impact of negatively reciprocal inclinations on worker behavior: Evidence from a retrenchment of pension rights. *Management Science* 62, no. 3:668–81.
Prendergast, Canice. 1999. The provision of incentives in firms. *Journal of Economic Literature* 37, no. 1:7–63.
Raith, Michael. 2008. Specific knowledge and performance measurement. *RAND Journal of Economics* 39, no. 4:1059–79.
Sandvik, Jason, Richard Saouma, Nathan Seegert, and Christopher Stanton. 2021. Employee responses to compensation changes: Evidence from a sales firm. *Management Science* 67, no. 12:7687–707.
Sigurdssona, Josef, and Rannveig Sigurdardottir. 2016. Time-dependent or state-dependent wage-setting? Evidence from periods of macroeconomic instability. *Journal of Monetary Economics* 78:50–66.