

EMPLOYMENT PROTECTION, BONUS PAY, AND LABOR MARKET
PERFORMANCE

Armin Falk¹, David Huffman², Bentley MacLeod³

March, 2007

Abstract

Employment protection legislation (EPL) refers to rules limiting the ability of an employer to fire a worker. EPL is often believed to have a strong negative impact on market performance, but field evidence suggests a weak or zero effect. This paper uses an experiment to study EPL, in which firms and workers can endogenously engage in repeated interactions, and firms are unable to enforce effort through explicit contracts. We exogenously vary the presence of an EPL institution in the market, and the form of labor contract available to firms. The first main result is that EPL has a strong negative impact on market efficiency in the case of efficiency wage contracts. In reality, however, firms typically have the option to offer deferred compensation, such as a bonus. The second main result is that EPL has a much milder effect on market efficiency when firms have the option to pay bonuses. We show that EPL harms efficiency in an efficiency wage world because it attacks a key incentive device: threat of firing. Given the option, it turns out that firms are able to credibly promise to reward effort with bonuses, so that the bonus substitutes for threat of firing as an incentive device. Thus, our findings suggest an explanation for the weak impact of EPL observed in the field: firms may be able to circumvent much of the negative effects of EPL by using bonus payments. In fact, we show that if firms are given the option, they endogenously make more use of bonuses in the presence of EPL. Given the salutary effect of the bonus option, it is natural to investigate whether it also improves market efficiency even in the absence of EPL. We find that bonuses perform better than wage contracts in early market periods, when most interactions are one-shot. In the long-run, however, when relationships have had time to develop and workers are faced with a meaningful firing threat, wage contracts are equally efficient.

Keywords: Employment protection, firing costs, incomplete contracts, bonus pay, experiment.

JEL codes:

¹ Institute for the Study of Labor (IZA) and University of Bonn.

² IZA

³ Columbia University

* Corresponding Author. Address: IZA, P.O. Box 7240, D-53072 Bonn, Germany.
Email: falk@iza.org; Tel.: +49-228-3894-112; Fax: +49-228-3894-180)

1 Introduction

Employment protection legislation (EPL) is a labor market institution that affects many developed nations (OECD, 1999). EPL refers generally to rules or legislation that limit the ability of an employer to dismiss a worker. The impact of EPL on labor market outcomes is one of the most intensely debated economic policy issues today. Empirical studies have used cross-country or cross-state differences in EPL, and have typically found either zero effect or a small negative effect (see, e.g., Nickell, 1997; Kugler and Saint-Paul, 2000; Autor et al., 2006). Even in the absence of EPL institutions, firing costs can arise, for example due to relationship-specific investments. Thus, the more general question is how an increase in the costs of firing changes market performance.

This paper approaches these questions from a contract enforcement perspective, because in many labor markets worker effort is unverifiable to third parties.¹ This lack of verifiability prevents the use of explicit contracts to enforce effort, and thus leaves firms to find other ways to mitigate worker shirking.² In such settings, understanding the impact of EPL requires an appreciation of how the institution interacts with the strategic response of workers and firms to the contract enforcement problem. There has been some theoretical work in this vein (Galdon-Sanchez and Güll, 2000; MacLeod, 2003 and 2006), but little empirical evidence that explicitly addresses how the impact of EPL depends on details of the contract enforcement problem.³ It is particularly useful to have evidence in the case of settings with repeated interactions, given that there are typically many possible equilibria and theory alone provides little guidance (see, e.g., Fudenberg and Maskin, 1986).

We study EPL in an experimental labor market, similar to the framework used in Brown, Falk, and Fehr (2004), in which effort is unverifiable and firms and workers can endogenously engage in repeated interactions. We exogenously vary whether or not there is an EPL institution in the market, which takes the following form: a firm can fire a worker during an initial probation period, but loses the ability to fire the worker if it

¹ See, e.g., Chevalier and Ellison (1997) and (1999); McMillan and Woodruff (1999); Banerjee and Duflo (2000); Hong, Kubik and Solomon (2000).

² For theoretical work on how markets function in the presence of contract enforcement problems see Gintis (1976); Klein and Leffler (1981); Shapiro and Stiglitz (1984); Bowles (1985); Bull (1987); Hart and Holmström (1987); MacLeod and Malcolmson (1989), (1993), and (1998); Baker, Gibbons and Murphy (1994); Dixit (2003); Levin (2003).

³ For other theoretical work on EPL, see, e.g., Lazear (1990), Kugler and Saint-Paul (2004), and Fella (2000).

chooses to continue the relationship beyond that point. We also exogenously vary the nature of the contract enforcement problem in the market, by varying which of two labor contracts is available to firms. One contract form, which has been studied extensively in the efficiency wage literature (Shapiro and Stiglitz, 1984), involves firms and workers agreeing on an enforceable, up-front wage. The other type of contract provides a more flexible technology for contract enforcement, giving the firm the option to offer a bonus in addition to the up-front wage, where the bonus is paid at the firm's discretion after observing worker performance. This design results in four exogenous treatments: T-Baseline, in which contracts involve only wages and there is no EPL; T-EPL, which is the same as T-Baseline but adds EPL; T-EPL-Bonus, which incorporates both EPL and the bonus option simultaneously; T-Bonus, which completes the design by adding the bonus option but not EPL.

The experiment allows us to investigate the impact of EPL at the aggregate and micro levels. Our main hypothesis is that EPL will have a strong negative impact on aggregate market efficiency (total economic surplus) in an efficiency wage world, but a much milder impact in the case that labor contracts include a bonus option. If a bonus option does in fact allow firms to circumvent most of the strategic disadvantage created by the strong form of EPL introduced in our experiment, this could potentially help explain the weak aggregate effect of EPL typically observed in field evidence. We use the experiment to lay bare the mechanisms at the micro level that explain how these aggregate effects of EPL arise. We hypothesize that EPL has a negative impact in the efficiency wage world because it attacks a key incentive device: threat of firing. This is likely to lead to low efficiency through increased worker shirking in long-term relationships. Given a bonus option, on the other hand, firms may be able to credibly promise to reward effort with a bonus. This could act as a substitute incentive device for threat of firing, and allow firms to eliminate some or all of worker shirking. EPL may have other important effects at the micro level, including changing the frequency of long-term relations in the market, changing worker behavior in initial, probation periods of relationships, and interfering with the ability of firms to screen out bad worker types. We address two other, related questions. One is whether firms make more use of a bonus option in a market with EPL, compared to a case where they have the bonus option but there is no EPL. This would imply that the types of contracts offered in the market respond endogenously to the presence of EPL.

The other question is whether the bonus option always improves market efficiency, even in the absence of EPL. There is some evidence on how wage and bonus contracts compare in terms of efficiency, in one-shot interactions (Fehr, Klein, and Schmidt, 2007). At this point, however, there is little evidence on how these two important contract forms perform in the long run, in a setting with repeated interactions such as the one we consider.

We begin the analysis by investigating T-Baseline. Efficiency is relatively high, reaching roughly 70 percent of the maximum in most periods. How does the market achieve this relatively good performance, despite the contract enforcement problem? High effort levels, which are a prerequisite for efficiency in our design, are mainly observed in long-term relationships, where firms and workers endogenously engage in a sequence of interactions with each other. Effort levels and efficiency are high because firms fire workers who do not perform well, and pay a relatively high (efficiency) wage so that the threat of firing is a strong incentive. In the final market period, where there is no potential for future rents, many workers reduce effort sharply to the minimum level, illustrating the importance of firing threat. Notably, some workers still put in high effort even in the final market period, in response to a relatively generous wage. This behavior is consistent with previous evidence that a subset of the population is motivated by reciprocity (Fehr, Kirchsteiger, and Riedl, 1999; for a survey see Fehr and Gaechter, 2000). These reciprocal types are attractive workers for firms, because they work hard in response to a fair wage, regardless of future considerations. In summary, performance in the baseline condition provides strong support for the disciplining version of the efficiency wage hypothesis (Shapiro and Stiglitz, 1984) but at the same time supports the fairness version of the hypothesis as well (Akerlof and Yellen, 1990).

Our first main result concerns the impact of EPL on aggregate outcomes in an efficiency wage world. Comparing T-Baseline and T-EPL, we find that EPL causes a substantial drop in efficiency. Furthermore, the difference is growing over time: in T-Baseline, efficiency improves strongly over initial market periods, before leveling off later in the game, whereas in T-EPL the trend is actually downward sloping. The time profile of efficiency in T-Baseline closely tracks the development of long-term relationships over time in the market. Long-term relationships also form in T-EPL, but as we discuss below, these do not have the same benefits in terms of high effort levels as in T-Baseline. Turning to the case where EPL is introduced in a market with the bonus option, we find a dramatic

difference: efficiency is much higher than in T-EPL, and is only slightly below efficiency in T-baseline. Thus, the impact of EPL depends strongly on the structure of the labor contract. Interestingly, the number of long-term relations is substantially lower in T-EPL-Bonus than in T-EPL or T-Baseline, suggesting a fundamentally different approach to contract enforcement.

For a deeper understanding of the aggregate effects of EPL, we turn to an investigation of firm and worker behavior at the micro level. We hypothesized that EPL would have a negative impact on efficiency in an efficiency wage world because it removes the threat of firing as an incentive device in long-term relationships. In fact, average effort is much lower in long-term relationships in T-EPL than in T-Baseline. Furthermore, within-worker effort is very likely to drop sharply between the first and second private offer in T-EPL.⁴ In one-shot interactions, where EPL has no direct impact on ability to fire, we find little difference in effort levels compared to T-Baseline.

We predicted that giving firms a bonus option could help firms provide an incentive for high effort, even in the absence of firing threat, to the extent that the promise to pay a bonus is credible. *A priori*, it is not obvious that the promise to pay a bonus will be credible. As discussed in MacLeod and Malcomson (1998), a setting with excess supply of workers makes it easier for a firm to cheat on bonuses, because the firm can move on to another worker each time (at least until qualified workers run out). We observe, however, that bonus payments are strongly increasing in worker effort levels. One factor contributing to credibility of bonuses could be EPL: once a firm cannot fire a worker, this makes a bonus promise more credible, because the firm cannot avoid worker retaliation. We also observe some firms paying a bonus in response to high effort even in the final market period, even though there are zero reputation or retaliation costs in this case. This suggests another force making bonus promises credible, namely the presence of some reciprocally-motivated firms. Turning to effort levels in long-term relationships, we find that effort levels are as high as in T-Baseline, a dramatic improvement relative to T-EPL and a clear indication that the bonus option is an effective device for deterring worker shirking. Similar to T-Baseline, there is little evidence of a within-worker reduction in effort when workers receive a second private offer.

⁴ This finding is consistent with Ichino and Riphan (2001), who find an increase in employee absenteeism in Italy, once workers are retained beyond the institutional period for EPL.

EPL also has a strong impact on several other dimensions of market performance. The distribution of relationship lengths becomes more bi-modal in T-EPL than in T-Baseline: there are many more one-shot interactions, reflecting reluctance of firms to enter relationships, and then a mass of very long-lived relationships, reflecting the fact that relationships that do form are very stable because workers almost never quit. In T-EPL-Bonus, worker shirking was no more pronounced than in T-Baseline, and yet firms are even less likely to enter relationships than in either T-EPL or T-Baseline. This reflects the fact that firms solve the contract enforcement problem in a fundamentally different way in T-EPL-Bonus, relying on the bonus option rather than relationships as an incentive device.

In an efficiency wage setting EPL also changes the behavior of workers in the initial, probation periods of a relationship. Workers in T-EPL put in significantly higher effort in the initial contract with a firm, compared to workers in T-Baseline, despite the tendency to put in much lower effort once they are rehired and EPL prevents dismissal. This appears to be strategy for enticing firms into long-term relationships, motivated by the strong incentive to achieve EPL status. Effort variance is also compressed in the probation period in T-EPL, and then increases substantially once workers are in the relationship protected by EPL. In T-Baseline, by contrast, effort variance is lower, and decreases even more after the probation period. In T-EPL-Bonus, both effort levels and effort variance are very close to T-Baseline, consistent with less incentive for workers to enter long-term relationships, relative to T-EPL .

EPL prevents firms from screening workers by interacting with a worker multiple periods in a row. Thus, one might expect that firms would be less successful at selecting only reciprocal types for long-term relationships. The problem is exacerbated by the tendency for selfish workers to imitate reciprocal workers in the probation period. Worker behavior in the final market period provides an indication for the proportion of reciprocal versus selfish workers in relationships. Indeed, we find that there are more workers who shirk completely in the final period of long-term relationships in T-EPL, compared to in T-Baseline. Thus, EPL does appear to lead to adverse selection of selfish types into long-term relationships, a type of adverse selection that has not typically been discussed in the literature on EPL.⁵ In T-EPL-Bonus, relatively fewer workers put in effort of 1 in

⁵ See Kugler and Saint-Paul (2004) for a model with firing costs and adverse selection on ability.

the final period, It is ambiguous whether this indicates relatively few selfish types, because the firm is the second mover in T-EPL-Bonus, and thus selfish types have an incentive to imitate reciprocal types even in the final period.

Given the salutary effects of the bonus option in the presence of EPL, one might expect firms to make more use of a bonus option in response to the presence of an EPL institution. We compare T-EPL-Bonus to T-Bonus, where there is a bonus option but no EPL. We find that firms do offer substantially higher bonuses in T-EPL-Bonus than in T-Bonus. The difference is most pronounced in long-term relationships, which makes sense given that this is where EPL removes threat of firing as an incentive device altogether. These findings complement recent evidence that performance pay is more likely for jobs with long tenure and a high degree of relationship-specific investment (Lemieux et al., 2006). Our results suggest that causality may run from relationship-specific investments to the use of bonus pay, rather than the other way around, something that is more difficult to determine in the field data. These findings also emphasizes the potential for firms to respond to EPL by making more use of bonus pay, thereby circumventing much of the negative impact on efficiency.

A final interesting question is whether the bonus option is beneficial *per se*, even in the absence of EPL. We observe that market efficiency in T-Bonus is significantly higher than in T-Baseline, for early market periods. In the long-run, however, T-Baseline catches up, due to the build-up of high-performing long-term relationships in the market. Thus, our results for the short-run, where most interactions are one-shot, are consistent with Fehr, Klein, and Schmidt (2007), who find that bonus contracts outperform wage contracts in a one-shot interaction setting. Because our setting includes the potential for long-term relationships, however, we are able to establish that in the long run there is no longer a compelling advantage of the bonus contract over the less flexible wage contract. Given sufficient time, the formation of stable long-term relationships allow a market with wage contracts to perform as well as a market where firms use bonus contracts. We also show that the bonus option has another important impact on market performance, leading to substantially fewer long-term relationships in the market as firms rely on the bonus rather than relational incentives.

The remainder of the paper is organized as follows. Section 2 describes our experimental design. In Section 3 we establish a baseline for market performance in an efficiency

wage setting, and show the key role of firing threat in sustaining market efficiency. Section 4 presents our results on how the impact of EPL depends on the form of the contract enforcement problem. Section 5 investigates the endogeneity of bonus pay with respect to EPL, and the impact of a bonus option on market performance in the absence of EPL. Section 6 concludes.

2 An Experimental Approach

The ideal data set for investigating the impact of EPL would involve two markets that are exactly the same, except for an exogenous implementation of EPL in one case. One reason why experiments are a useful complement to field evidence for studying labor market institutions is the possibility for truly exogenous variation of this type. In the field, by contrast, there is typically the possibility that the EPL institution was adopted endogenously in response to labor market conditions, making causal inferences more difficult (Lazear, 1990; Besley and Case, 2000). Investigating how the impact of EPL varies with the structure of the labor contract also requires exogenous variation in the presence or absence of a bonus option. In the field it is usually not possible to find a situation where there is exogenous variation of this type, holding all else constant, and it may even be non-trivial to establish whether a job involves performance pay or not, due to data limitations (Lemieux et al., 2006). Experiments also have the advantage that it is possible to observe the impact of an institution at the micro level, on the behavior of individual workers and firms. This is useful for investigating the mechanisms underlying changes at the aggregate level. Other advantages include the possibility to know the benchmark of maximum efficiency in the market, to observe all contract offers in the market, even those that are not accepted, and to observe crucial variables such as worker effort levels, all with zero measurement error.

2.1 Design

In the experiment there were 18 trading periods. In each period a firm could hire at most one worker, and a worker could have at most one job. An individual period involved two or three phases, depending on the treatment. The first phase was always a market phase, in which the firms made contract offers and workers could only accept or reject. Firms could

make as many contract offers as they wanted during the time limit of three minutes; if one of a firm's contracts was accepted, all of the other offers by that firm were immediately removed from the market.⁶ In the case that a firm and a worker agreed on a contract, they entered a second phase in which the worker could decide how much effort, e to exert. In treatments where the contract offer could include an offered bonus, there was a third phase in which the firm was informed about the worker's effort choice and could decide how much of a bonus, b , to pay. Importantly, neither the worker's effort level or the firm's bonus payment were restricted by the initial contract agreement, whereas a wage specified in the agreement was binding. After the second (third) phases, the firm and worker were informed about their profits and earnings, respectively, and then a new period began.

Contract offers consisted of a wage, w , a desired effort level, \tilde{e} , and in some treatments an offered bonus, \tilde{b} . The offer also included the firm's ID number. Firms could make two types of contract offers during the market phase: public offers or private offers. Public offers were observed by all workers, and thus were available to any worker. Private offers were observed only by a worker specified by the firm, and thus were available only to that particular worker. In the case that a firm made a private offer, the firm specified a worker's ID number, in addition to the contract terms. Worker and firm ID numbers remained constant over the entire 18 periods, so it was possible for a firm to intentionally make a private offer to the same worker over multiple periods, and for a worker to recognize offers coming from a specific firm. This design made it possible for a firm and worker to endogenously form a long-term relationship, by choosing to repeatedly engage in private-offer contracts with each other over multiple periods. Public offers were a way for firms to engage in a spot market for labor rather than engaging in long-term relationships. During the market phase, firms were kept constantly informed about which workers had already accepted a contract, so as to avoid having firms make a private offer to a worker that was no longer available.

In treatments with EPL, the firm lost the ability to fire a worker after making a second private offer in a row to the same worker. This design captures a common feature of EPL institutions, which is a specified probation period during which the firm is still able to fire the worker. Having chosen to hire the worker again, after the initial private offer,

⁶ If all firms had contracts, there was no potential for further trades. Thus, the market phase was designed to end automatically after three minutes, or after the last firm had a contract offer accepted, whichever came first.

EPL took effect and the firm had to make an offer to that same worker at the beginning of each subsequent period until the end of the game or until the worker decided to quit. Firms in EPL contracts made their offers in a special phase before the market phase. Importantly, the wage offer had to be at least as high as in the period when EPL was initiated. Some rigidity of the wage is required for an EPL institution to work, because otherwise a firm could effectively fire a worker by reducing the wage to zero. Firms were still able to increase the wage during subsequent periods, and to reduce the wage back to the initial level. After firms had made their EPL offers, the market period began and workers on EPL could see the standing offer from the firm, in addition to the other market activity. At any time, the worker could accept the standing offer, in which case the firm was informed. Alternatively, the worker could reject the offer by accepting another contract in the market. As soon as the worker rejected the standing offer, the firm was informed, and allowed to make offers during the remainder of the market phase.

In our design we abstract away from several issues sometimes discussed in the literature on EPL. The effort cost function for workers, described below, is the same across all individuals so there are no differences in ability. This allows us to focus on the moral hazard problem in terms of effort and bonus payment, without the complication of adverse selection in terms of ability. We also do not implement cyclical shocks to market conditions, or worker redundancies (multiple workers at one firm). This simplifies an already complex inter-temporal choice environment, and makes it possible to first understand the impact of EPL on the strategic behavior surrounding the contract enforcement problem. Ultimately, this behavior will be the driver of market performance in the case of interesting extensions involving cyclical shocks, or worker redundancies.

We implemented four exogenous treatments. In T-Baseline, contracts were wage-only. There was no EPL institution in the market, so firms could engage in as many private offers in a row with a worker as they wanted, while always having the option to fire the worker, i.e., not make the worker a private offer in the next period. This treatment is very similar to the ICF treatment Brown, Falk, and Fehr (2004), except that their design involved only 15 trading periods rather than 18. In T-EPL, contracts were still wage-only, but we introduced our EPL institution. In T-EPL-Bonus, the EPL institution was in effect, but firms had the option to offer a bonus, in addition to or instead of a wage. In T-Bonus, there was no EPL, but firms had the option to offer bonuses.

2.2 Parameters, Information Conditions, Procedure, and Subject Pool

All market sessions lasted 18 periods, and had 7 firms and 10 workers. The material payoff to a firm was given by the function

$$\pi_f = \begin{cases} 10 \cdot e - w - b & \text{if a contract offer was was accepted} \\ 0 & \text{if no contract offer was accepted} \end{cases}$$

and the payoff function for a worker was given by

$$\pi_w = \begin{cases} w + b - c(e) & \text{if a contract offer was was accepted} \\ 5 & \text{if no contract offer was accepted} \end{cases}$$

where $c(e)$ was a cost of effort function, and 5 was the unemployment benefit in the case that a worker did not engage in a trade. The wage, w , the offered bonus, \tilde{b} , and the bonus actually paid, b , could each take on an integer value 0, 1, 2, ...100. The desired effort level, and the actual effort level chosen by the worker could take on integer values 1, 2, ..., 10. The effort cost function is shown in Table 1.

[Table 1 about here]

The cost function is increasing and convex. Because the marginal cost of effort is at most 3, while the marginal benefit to a firm is always 10, the efficient effort level is 10.

Payoff functions for workers and firms, including the effort cost function, were common knowledge. Participants were aware that the market would last 18 periods. Reputations could form bi-laterally: firms learned about the effort choices of workers that they traded with, but did not observe the effort choices, or firm bonus decisions, in interactions that they were not a part of; workers learned about the bonus decisions of firms that they encountered, but not about worker effort choices or firm bonus decisions in other market interactions. Firms observed all public offers on the market during the market phase. Workers were informed not only about private offers they had received, but also about all public offers on the market.

The experiment was computerized using Z-Tree software (Fischbacher, 1999). There was a practice period before the experiment began, which consisted only of a market phase,

in which subjects had the experience of making and accepting offers. After the practice period, the first period of the paid experiment began. At the end of each period, a subject's period profits were summarized, along with the profits of the trading partner in the case of a trade. Subjects were also reminded of the partner's ID number, the terms of the initial contract, the actual effort choice, and the actual bonus paid. Subjects could record this information on a separate sheet of paper, ensuring that subjects were fully informed about their own trading history over the course of the experiment.

There were 408 participants in the experiment. We conducted six market sessions for each of the four treatments, for a total of twenty-four sessions. Subjects were students at the University of Bonn, from various fields of study. No subject participated in more than one session. On average, a session lasted roughly 100 minutes, and a subject earned 25 Euros (32 USD).

3 Baseline Market Performance and the Role of Firing Threat

In this section we describe how the labor market functions in T-Baseline, where there is no EPL and no bonus option. Our aim is to establish a benchmark for aggregate market efficiency, and to highlight key mechanisms at the micro level that explain the relatively good performance of the market. The section is deliberately brief, given that the results are largely a replication of Brown, Falk, and Fehr (2004), who conducted a very similar treatment except that their design involved 15 rather than 18 trading periods.

The top panel of Figure 1 shows that average efficiency in T-Baseline is between 60 and 70 percent of the maximum in most periods, and is rising over time until a sharp drop in the final period.⁷ Efficiency is increasing monotonically in effort levels in our design, so high efficiency is equivalent to high worker effort. In fact, the modal effort level across all periods is 10 in T-Baseline, despite the fact that the sub-game perfect equilibrium for the stage game is an effort level of 1.

How does the market achieve relatively high aggregate efficiency, or equivalently, high worker effort levels? A first observation is that high effort levels are mainly ob-

⁷ The figure reports the ratio of total surplus generated in the market relative to maximum possible surplus. Total surplus is firm revenues minus worker effort costs and opportunity costs in terms of unemployment benefits. With 7 firms in each session, maximum surplus in a period is equal to $7 \cdot (10 * 10 - 18 - 5) = 77 \cdot 7 = 539$.

served in long-term relationships. Long-term relationships form when a firm and worker endogenously choose to engage in an unbroken sequence of private offer contracts with each other. Long-term relationships are prevalent in the market. Of all contracts, 67 percent are formed using a private offer, and out of all private offer contracts roughly 64 percent end up being part of a relationship lasting at least 5 periods. For private offer contracts, the correlation between effort level and the ultimate length of the relationship is 0.64 (Spearman; $p < 0.001$). The bottom panel of Figure 1 suggests why long-term relationships are characterized by high effort levels and high efficiency: firms only renew a private contract if the worker performed well in the previous period. For low effort levels, the probability that the firm terminates the relationships is close to 1, and for the maximum effort level the probability of termination is close to zero.

[Figure 1 about here]

The threat of firing creates a material incentive for workers to perform if continuing the relationship offers a rent relative to the alternative. In fact, averaging over all market periods, worker earnings are roughly 57 percent higher in contracts that are part of a long-term relationship, compared to in one-shot public-offer or private-offer interactions.⁸ Workers earn rents in long-term relationships, despite putting in relatively high effort, because firms pay substantially higher wages in long-term relationships than in one-shot interactions (mean wages are 50 rather than 30). Long-term relationships are also profitable for firms, explaining why these bilateral “trading islands” emerge in equilibrium. Average firm profits are roughly 45 percent larger in long-term relationships than in one-shot interactions.

Worker behavior in the final market period confirms the importance of future rents, and threat of firing, for motivating high effort levels. 60 percent of workers in long-term relationships reduce effort when they reach period 18, and 24 percent of workers choose an effort level of 1, despite choosing effort of 10 in period 17. Firms reduce wages somewhat in the final period, perhaps anticipating this behavior, but even controlling for the lower

⁸ The attractiveness of relationships can alternatively be seen in a forward-looking calculation of rents for each market period. We calculate the average current and future earnings, from t to $T = 18$, of all workers in a private offer contract in market period t . We compare this to the average current and future earnings of all workers currently in a public-offer contract. The difference between these two values is large and positive in every period. Results available upon request

wage in period 18, the (regression adjusted) average change in effort is -2.57, significant at the five percent level.⁹

Although it is true that on average worker effort levels drop when the threat of firing is removed, this conceals an important type of heterogeneity in worker motives. In particular, some workers are willing to put in high effort even in the absence of future rents. 10 percent of workers in long-term relationships put in an effort level of 10 in period 18. A clue as to the source of this behavior comes from the wage-effort relation: worker effort levels increase strongly in response to the wage, even in the final period, and effort levels of 10 are realized in those contracts with relatively high wages. This non-strategic effort provision is consistent with evidence that some workers in the population have reciprocal inclinations, such that they reward kind actions, or fair wages, with high effort (Fehr, Kirchsteiger, and Riedl, 1999; for a review see Fehr and Gaechter, 2000). It suggests that some of the effort provision throughout the game may be driven by fairness, rather than strategic considerations, although there is clearly a substantial fraction of workers who are strategic.¹⁰ Reciprocal workers are attractive types for firms, given that they work hard in response to a fair wage, independent of future considerations. Later in the analysis we investigate how EPL affects the selection of selfish versus reciprocal workers into long-term relationships.

In summary, T-Baseline establishes a benchmark for market efficiency, at 60 to 70 percent of the maximum. It also shows that rents, and threat of firing, are crucial for eliciting worker effort, although a subset of workers respond to high wages with high effort even in the absence of future rents. In the next section, we investigate how EPL changes market performance relative to this baseline, at the aggregate level and in terms of the behavior of workers and firms.

⁹ This result is based on an OLS regression of effort on the wage and a constant, for all contracts in period 18. Standard errors are robust, and adjusted for clustering on session.

¹⁰ In an appendix, Brown, Falk, and Fehr (2004) argue that the presence of some reciprocal types is important for sustaining cooperation of selfish workers in early market periods. Intuitively, the presence of reciprocal workers allows firms to pay high wages in the final period, thereby providing a rent for selfish workers. This gives selfish workers an incentive to perform in the second to last period. Without rents in the final period, performance of selfish workers would be unsustainable in all market periods, by backwards induction arguments.

4 The Impact of EPL on Market Performance

We start our analysis by discussing the impact of EPL on market efficiency, firm profits, and worker earnings, in Section 4.1. We then turn to an investigation of how EPL affects performance of firms and workers at the micro level, with and without bonus pay. Section 4.2 considers the impact of EPL on worker shirking in long-term and one-shot interactions. In Section 4.3 we investigate firm willingness to enter relationships and the frequency of long-term relationships in the market. Section 4.4 explores worker behavior in the initial probation period of relationships, and Section 4.5 investigates which types of workers, in terms of selfish or reciprocally inclined individuals, end up being in long-term relationships.

4.1 Aggregate Market Efficiency

A first indication that EPL dramatically changes market performance comes from the upper panel of Figure 2. The figure shows that the presence of EPL in an efficiency wage world leads to substantially lower efficiency. On average efficiency in T-EPL is about 12 percentage points lower than in T-Baseline. The difference is also growing over time: starting around period 6, efficiency in T-Baseline begins to climb steeply and leaves efficiency in T-EPL behind. From period 10 through 17, for example, efficiency in T-Baseline is about 17 percentage points higher than in T-EPL, and the difference is statistically significant using a conservative nonparametric test that treats the session average as one independent observation (ranksum; $p < 0.05$). Thus, the costs of introducing EPL are higher in the long run. Notably, the increase in efficiency in T-Baseline tracks closely the development of long-term relationships in the market: the number of long-term relationships increases strongly over the game, and then remains roughly constant after about period 11, at the same time that efficiency levels off. In T-EPL, the fraction of interactions in long-term relationships is also increasing over time, although not as quickly, but as we will see below, long-term relations function very differently in T-EPL and do not have the same benefits in terms of high worker effort levels and high efficiency.

[Figure 2 about here]

From the lower panel of Figure 2 it is clear that the impact of EPL changes fundamentally when the employment contract includes the bonus option. Efficiency is substan-

tially higher in T-EPL-Bonus than in T-EPL, a positive and significant difference of 11 percentage points on average, which is growing over time as efficiency in T-EPL trends slowly downward. The difference is almost significant at the 5 percent level using the most conservative nonparametric test (ranksum; $p < 0.06$). On the other hand, efficiency in T-EPL-Bonus is only 1 percentage point lower than average efficiency in T-Baseline, and is not statistically different (ranksum; $p < 0.87$). Efficiency is actually higher in T-EPL-Bonus than in T-Baseline for early market periods, although T-Baseline begins to outstrip T-EPL-Bonus around period 6. The long-run trend may be for efficiency in T-Baseline to surpass efficiency in T-EPL-Bonus, but the costs of EPL are clearly much smaller than when the contract form includes only wages. Interestingly, as discussed in more detail below, the number of long-term relationships in T-EPL-Bonus is consistently smaller than in either T-Baseline or T-EPL. This suggests a very different approach to solving the contract enforcement problem, which nevertheless sufficiently ameliorates the negative effects of EPL institution so as to come close to T-Baseline in terms of efficiency.

Given our prediction that EPL harms market efficiency by weakening the strategic position of firms who choose to enter long-term relationships, we might expect that EPL harms workers more than firms in T-EPL. This is in fact the case. Workers earn about 8.5 percent less than in T-EPL than in T-Baseline, but firms suffer relatively more, earning average profits that are 27 percent lower than in T-Baseline. To the extent that the bonus option gives firms an additional incentive device, which can be used to effectively elicit effort, we would expect firms to be less harmed by EPL in T-EPL-Bonus than in T-EPL. Indeed, firms are substantially better off in T-EPL-Bonus than in T-EPL, earning 3 percent higher profits than in T-Baseline. Workers, on the other hand, are still worse off in the presence of EPL, despite the presence of the bonus option, earning 7 percent less in T-EPL-Bonus than in T-Baseline.

In summary, the results in this section show that EPL has a negative impact on market efficiency in an efficiency wage world, and the difference is growing over time. The negative impact of EPL is much milder, however, if the contract form in the market includes the option of bonus pay.

4.2 Worker Shirking

For a deeper understanding of how EPL changes the functioning of the market, we turn to an investigation of behavior of firms and workers at the micro level. Given the importance of firing threat for sustaining high effort levels and efficiency in the baseline treatment, we expect that the negative impact of EPL on efficiency works through a greater propensity for workers to shirk, once they enter long-term relationships and the threat of firing is removed. On the other hand, the salutary effect of bonus pay at the aggregate level suggests that it may help deter worker shirking, even in the presence of EPL.

The top panel of Figure 3 shows that EPL does in fact cause a dramatic drop in worker effort in long-term relationships, in an efficiency wage setting.¹¹ Average effort is 5.5 in long-term relationships under EPL, in contrast to 8.2 in long-term relationships in T-Baseline (ranksum; $p < 0.001$). Notably, whereas long-term relationships frequently break-up in T-Baseline, due to firms firing poor performers, EPL relationships essentially never breakup in T-EPL, despite low effort levels. This is because firms cannot fire workers, and workers essentially never quit EPL relationships, doing so in only 3 out of 283 EPL contracts.. This is understandable given that worker earnings are about 63 percent higher on average in EPL relationships than in one-shot interactions, in T-EPL. The bottom panel of Figure 3, on the other hand, shows little difference in effort levels in one-shot interactions. This is expected given that in one-shot interactions EPL has no impact on firms' ability to fire the worker. If anything, worker effort is actually higher in T-EPL than in T-Baseline. Below, we investigate one possible explanation for this latter result: a tendency for workers in T-EPL to perform especially well during the probation period of a relationship, in an attempt to entice firms into long-term relationships.

[Figure 3 about here]

The strongest evidence that workers respond strategically to the reduced firing threat under EPL would be a within-worker drop in effort, occurring precisely when the EPL barrier to firing takes effect: in the second private-offer contract . Indeed, we find that the unconditional probability that a worker reduces effort in the second period of a relationship relative to the first is more than fifty percent (0.54) in T-EPL. In T-Baseline, by contrast,

¹¹ In the figure, long-term relationships are defined as relationships consisting of at least two private offer contracts.

the probability is half as large (0.23). The average difference between effort in the first and second private offer contract is negative in T-EPL, equal to -1.57, whereas it is positive in T-Baseline, equal to 0.11.

Columns (1) and (2) in Table 2 tell a similar story in a regression framework. The table reports Probit regressions, estimated only for private offer contracts that take place in the second period of a relationship. The dependent variable is equal to 1 if a worker reduced effort relative to the first period and zero otherwise. Coefficients are marginal effects; standard errors are robust, adjusted for clustering on session. The coefficient on the dummy variable for T-EPL in Column (1) is positive and highly significant, showing that workers are much more likely to reduce effort in the second relationship period in T-EPL, compared to T-Baseline. Column (2) shows that the effect is still strong controlling for other potentially relevant variables, such as current and previous period wages, and the worker's effort level in the previous period. Thus, the lower effort levels in long-term relationships in T-EPL reflect a moral hazard problem such that workers tend to shirk once the threat of firing is removed.

[Table 2 about here]

We have shown that EPL substantially lowers efficiency in a market that otherwise performs relatively well through a combination of efficiency wages and firing threat. In reality, however, many firms use more than just wages and dismissal threat to motivate workers in an ongoing relationship. The prime example is deferred compensation, paid contingent on performance. We hypothesized that the detrimental impact of EPL might be mitigated in the later setting, and this is what we observed in our analysis of aggregate outcomes. We predicted that the bonus option could help undo the negative effects of EPL if firms are able to credibly promise to reward effort with a bonus. In this case, the bonus option provides an incentive for workers to perform, and could thus substitute for threat of firing.

Whether or not the promise to pay a bonus will be perceived as credible by workers is not obvious *a priori*. Figure 4 shows, however, that actual bonus payments are strongly increasing in worker effort levels, implying that bonus promises are credible. We also regressed actual bonus payments in period t on worker effort levels in t (interval regression), and found that the relationship between effort and bonus payments is also highly

statistically significant ($p < 0.001$; robust s.e., clustering on session). There is a positive and significant relationship between bonus payments and effort even in the final market period, despite the lack of any reputation or other strategic motives to pay. This later finding suggests the presence of some firms with reciprocal inclinations. In fact, among all firms in T-EPL-Bonus, we observe that 29 percent pay a bonus greater than zero in period 18.

[Figure 4 about here]

The bonus option does appear to solve the problem of worker shirking in long-term relationships. The top panel of Figure 3 shows that in T-EPL-Bonus worker effort levels in long-term relationships are much higher than in T-EPL, a statistically significant difference (ranksum; $p < 0.001$). In fact, worker shirking is no more pronounced than in T-Baseline (ranksum; $p < 0.87$). The bottom panel of Figure 3 shows little difference in effort levels in one-shot interactions, compared to T-Baseline or T-EPL. Column (3) of Table 2 shows that the probability of a drop in worker effort, going from the first to the second period of a relationship, is significantly lower in T-EPL-bonus than in T-EPL. Column (4) shows that this is true conditioning on the wage and other controls. In Columns (5) and (6) we see the probability of a drop in effort level is not statistically different from T-Baseline. Overall, these findings suggest that the bonus option contributes to market efficiency in the presence of EPL by helping firms to deter worker shirking. The bonus option appears to matter most in long-term relationships, which is precisely where EPL removes the threat of firing as an alternative incentive device.

4.3 Frequency and Distribution of Long-Term Relationships

By interfering with the effectiveness of long-term relationships as an incentive device, EPL may also have a strong influence on the frequency and duration of long-term relationships in the market. This is a potentially important ramification, given that in the real-world long-term relationships entail benefits such as greater relationship-specific investments. A decrease in the prevalence of relationships due to EPL could also be another channel leading to lower aggregate efficiency, because a reluctance of firms to initial relationships reduces the incentives for workers to perform even in public-offer or first private-offer contracts. On the other hand, as we discuss in the next sub-section, there are potentially

offsetting motives that might cause workers to perform especially well in one-shot interactions in T-EPL. However, this this strategic behavior makes it more likely that firms end up with selfish workers in long-term relationships, which again tends to lower efficiency.

Given that EPL fosters worker shirking in long-term relationships, we might expect firms to be more reluctant to initiate relationships in T-EPL compared to T-Baseline. A first indication that firms do avoid relationships in T-EPL comes from the prevalence of one-shot interactions: 58 percent of all contracts were in one-shot interactions in T-EPL, compared to 48 percent in T-Baseline. Even more tellingly, we observe 23 percent of all firms in T-EPL pursuing a strategy of strict relationship-avoidance, never making two private offers in a row to the same worker, during the entire game. This contrasts with only 9 percent of firms exhibiting this strategy in T-Baseline. On the other hand, although there are more one-shot relationships in T-EPL, the average length of long-term relationships is actually slightly longer in T-EPL than in T-Baseline, 6 periods versus 5.89. This reflects the greater stability of EPL relationships relative to long-term relationships in T-Baseline. Conditional on entering a relationship, the relationship almost always lasts the rest of the game in T-EPL, whereas frequent dismissals make relationships less stable in T-Baseline. Thus, as can also be seen from the cumulative distributions of relationship length plotted in Figure 5, EPL changes the market by fostering a more bi-modal distribution of relationship lengths. There are substantially more interactions that last only one period, and then a mass of very long-lived relationships, with few relationships of intermediate lengths.

[Figure 5 about here]

One might expect that firms would be more willing to enter relationships in T-EPL-Bonus than in T-EPL, because EPL has a milder effect and worker shirking in long-term relationships is much less pronounced. On the other hand, we have seen that the bonus option is a potential substitute for relationships as an incentive device, and thus it could be that firms are actually less likely to engage in repeated interactions in T-EPL-Bonus. Looking at the frequency of one-shot interactions, we find a slightly higher proportion in T-EPL-Bonus than in T-EPL: 60 percent, versus 58 percent. The fraction of firms who strictly avoid relationships is 20 percent, close to the proportion observe in T-EPL and substantially larger than in T-Baseline. The average length of a relationship in T-EPL-

Bonus is somewhat shorter than in either T-EPL or T-Baseline, 4.84 as opposed to 6 or 6.99, despite the fact that workers quitting EPL is just as rare as in T-EPL. As shown in Figure 5, the cumulative distribution function for T-EPL-Bonus is similar to T-Baseline but shifted upward, indicating fewer long-term relationships than in either of the other treatments. Taken together, these results suggest a fundamentally different approach to contract enforcement in T-EPL-Bonus, such that firms rely on the bonus option and engage in fewer long-term relationships. An interesting question that we have not answered is how much firms would rely on relationships given the bonus option, and an absence of EPL. We address this question in the final results section of the paper, which analyzes the fourth treatment T-Bonus.

4.4 Worker Behavior in the Probation Period

EPL also has the potential to change the behavior of workers in the market, in the initial or probation periods of relationships. We have seen a greater reluctance of firms to enter relationships in T-EPL, compared to T-Baseline, and we argued above that this goes in the direction of reducing the incentive for workers to perform in one-shot interactions, because there is a lower probability of contract renewal. On the other hand, to the extent that workers see a greater incentive to enter long-term relationships because of the potential to exploit the lack of firing threat, it could be that workers actually perform even better in one-shot interactions than in T-Baseline. If this latter effect dominates, it would have implications for the variance of worker effort in T-EPL, as well as the level, leading to a more compressed distribution of effort levels in probation periods of relationships. This sub-section investigates these empirical questions.

Table 3 reports regressions using all contracts that take place in the first private offer between a worker and a firm. The dependent variable is the worker's effort choice. Because the dependent variable is elicited in intervals, values are left and right censored, so we use interval regression. Coefficients are marginal effects. Robust standard errors, clustering on session, are reported in brackets. Column (1) shows that effort levels are actually significantly higher in T-EPL, compared to T-Baseline, in the initial, probation period of a relationship. This difference remains similarly strong and significant in Column (2) where we add controls for the wage and other potentially relevant variables. The results are also robust if we look at effort levels in all one-shot interactions, including public-offer

contracts, which could also play some role as a probation period if a firm first tries out workers using public offers. Thus, the incentive to entice firms into long-term relationships appears to dominate, leading to a "probation period effect" such that workers who plan to shirk behave as non-shirkers until the threat of firing has been removed and they can exploit the firm.

[Table 3 about here]

In addition to the level, variance in worker effort is an important dimension of performance that may be influenced by EPL. In Figure 6 we show the distributions of worker effort levels, separately for the probation period and later periods of long-term relationships. The probation period category includes all private-offer contracts that ultimately lead to a longer relationship. Restricting the analysis to individuals who are in both the probation and later relationship period categories helps make clear how individual behavior changes before and after the second private offer. If we instead include all one-shot interactions in the probation period category, however, the same qualitative results emerge. The important qualitative results are also preserved if we instead look at distribution of effort per dollar of compensation, which indicates that the results are not driven by differences in the distribution of offered compensation across treatments.

The top-left panel of the figure shows a strong difference in the variance of, going from the probation period to later relationship periods. The height of the box in the plot, which shows the interquartile range (25th to 75th percentile) of the distribution, is relatively small in the probation period and the median (shown by the horizontal line in the box) is relatively high. The height of the box increases substantially once the second private offer is received, and the mean effort level drops sharply, showing that variance in effort is much higher and the median level is much lower in long-term relationships compared to the probation period. The top-right panel of the figure shows that In T-Baseline, by contrast, the distribution of effort actually becomes more compressed moving from the probation period to later periods. The median also increases relative to the probation period. The median effort in the probation period in T-EPL is higher than in T-Baseline, providing another indication that workers in T-EPL perform even better in probation periods than workers in T-Baseline. Finally, the lower panel shows that distributions in T-EPL-Bonus are similar to those for T-Baseline. The bonus option

eliminates the strong increase in variance in long-term relationships observed in T-EPL, leading instead to a more compressed distribution similar to T-Baseline.¹²

In summary, EPL has important implications for the level and variance of worker effort, in the initial period of relationships and in later periods. EPL leads to compressed variance and high effort levels in initial relationship periods in an efficiency wage world, as workers try to entice firms into long-term relationships. This is followed by a dramatic increase in variance and drop in median effort once the relationship is continued and threat of firing is removed. Notably, however, not all workers reduce effort to 1 in long-term relationships in T-EPL. Workers at the 75th percentile put in effort of 9, and some workers even put in the maximum level of 10. This suggests the presence of some reciprocal types in long-term relationships, and that selfish types imitate reciprocal ones during the probation period. In the next sub-section we compare the proportion of reciprocal and selfish workers who end up in long-term relationships, depending on the presence or absence of EPL.

4.5 Selection of Worker Types into Long-Term Relationships

The presence of EPL in the market makes it more difficult for a firm to screen workers, because it is no longer possible to try the same worker multiple periods in a row. Firms can adopt strategies where they interact repeatedly with a worker, with a break of one or more periods in-between each time, but this requires the cooperation of the worker and tends to slow the learning process *ceteris paribus*. As a result of this distortion in the market, one might expect firms to have more trouble identifying reciprocal workers and screening out selfish types who tend to shirk. Furthermore, the tendency for selfish workers to exaggerate their willingness to perform in the probation period makes it all the more difficult for firms to separate types in T-EPL.

In order to assess the proportion of selfish versus reciprocal types who end up getting through the screening process, and into long-term relationships, we look at worker behavior in the final market period. In period 18, selfish workers have no incentive not to reveal

¹² Notably, effort levels are higher in T-EPL bonus than in T-EPL or T-Baseline, but this reflects higher offered compensation. The ratio of effort to compensation is very similar in T-EPL-Bonus to T-Baseline, and is lower than in T-EPL. In particular, effort per unit of compensation is lower in the probation period than in T-EPL, showing that the bonus option eliminates the probation period effect, i.e., greater willingness of workers to perform for the same compensation.

themselves by choosing an effort level of 1. Our findings are suggestive of more selfish types ending up in long-term relationships in T-EPL, compared to T-Baseline. In the final market period, 44 percent of workers in long-term relationships choose an effort level of 1 in T-EPL, compared to 28 percent in T-Baseline. In order to control for potential differences in final-period wages across treatments, we also regressed effort levels in the final market period on a treatment dummy and the wage (interval regression). The sample included all private-offer contracts that were part of a long-term relationship going into the final period. We find that final-period effort levels are significantly lower in long-term relationships in T-EPL, compared to in long-term relationships in T-Baseline, controlling for the wage ($p < 0.05$; robust s.e., clustering on session). We also estimated a similar regression comparing effort levels in T-EPL-Bonus to T-Baseline. In this case, there is not a significant difference in effort levels, controlling for offered compensation ($p < 0.27$; robust s.e., clustering on session). This latter result may not fully reflect the number of selfish types in T-EPL-Bonus, however, because the presence of reciprocal firms who pay a bonus even in the final period gives even selfish types an incentive to perform. Again, the benefit of the bonus option for firms is that it substitutes for threat of firing as an incentive device, and unlike threat of firing can provide an incentive even in the final period.

In summary, there is some evidence of adverse selection of selfish worker types into long-term relationships in a market with EPL. This is consistent with a greater difficulty of screening workers, due the limited opportunities to try a worker multiple times, and due to selfish workers' strategic imitation of good behavior in the probation period.

5 Bonus Pay in the absence of EPL

In this section we turn to our fourth treatment, T-Bonus, which incorporates the bonus option but does not include EPL. This treatment allows us to make two additional comparisons. First, we investigate how much firms make use of a bonus option depending on whether or not EPL is present in the market. Second, we investigate how the bonus option changes market performance relative to T-Baseline, in the absence of EPL.

5.1 Endogeneity of the Bonus with Respect to EPL

We have shown that the bonus option helps firms compensate for the lack of firing threat in T-EPL-Bonus, deterring worker shirking and preventing the loss of profits that occurs when EPL is introduced without the bonus option. In T-Bonus, on the other hand, firms do not face the same restrictions on the use of firing threat as an incentive device. Thus, one might predict that firms would make relatively more use of the bonus option in T-EPL-Bonus, compared to T-Bonus. Furthermore, it would make sense for the difference to be appear in long-term relationships, given that this is where EPL removes the threat of firing in T-EPL-Bonus. Figure 7 shows the impact of EPL on contract terms in long-term relationships, when bonuses are an option. The average wage is similar across the two treatments, but offered bonuses are substantially higher in T-EPL-Bonus than in T-EPL. Thus, firms offer higher total compensation in long-term relationships in T-EPL-Bonus, but rather than being reflected in wages, this difference is largely in the form of higher offered bonuses. In one-shot interactions, by contrast, average wages are similar across treatments and offered bonuses are actually higher in T-Bonus than in T-EPL-Bonus.

[Figure 7 about here]

Table 4 tells a similar story, reporting regressions of the offered bonus on a treatment dummy for T-EPL-Bonus, the wage, and other controls (interval regression). Offered bonuses are significantly higher in T-EPL-Bonus than in T-Bonus, controlling for the wage, consistent with the hypothesis that firms offer more bonus pay in long-term relationships if the market is affected by EPL.

[Table 4 about here]

In summary, EPL causes a systematic change in contract terms, reflected mainly in offered bonuses rather than wages. This illustrates the potential for contract forms in the market to respond endogenously a change in institutions, and it suggests that firms recognize the value of the bonus option as a way of circumventing the strategic disadvantages introduced by EPL.

5.2 Impact of the Bonus Option on Market Performance

We have shown that a bonus option improves market efficiency substantially in the presence of EPL. The market also functioned quite differently in the presence of EPL and the bonus option, in the sense that firms were less likely to form long-term relationships. To complete the picture of how the bonus option affects market performance, it is thus interesting to investigate whether adding a bonus option is always good for market efficiency, even in the absence of EPL, and how it affects relationship lengths in the market.

There is some previous evidence on how wage contracts perform relative to contracts with a bonus option, in one-shot interactions. Fehr, Klein, and Schmidt (2007) conducted experiments with the same payoff functions, and a similar effort cost function, to our design. Principals and Agents participated in ten periods, each time matched with a different anonymous partner. A central finding of this study is that a contract with a bonus option, in addition to an up-front wage, is substantially more efficient than a less-flexible wage-only contract. Comparing performance of the two contract forms in our setting, with repeated interactions, has the potential to add to this line of research.

Figure 8 compares efficiency in T-Bonus to efficiency in T-Baseline. In early market periods, efficiency is actually higher in T-Bonus than in T-Baseline, although the difference is not statistically significant (ranksum; $p < 0.15$). By period 7, however, efficiency in T-Baseline catches up, and is at least as high as in T-Bonus for the remainder of the game. Thus, our findings for early market periods, where most interactions are one-shot, go in the same direction as Fehr, Klein, and Schmidt (2007), suggesting that bonus contracts might be more efficient. In the long-run, however, once relationships have had time to develop and the threat of firing takes effect, wage contracts perform as well as bonus contracts. The equivalence of wage and bonus contracts also provides an explanation for why one contract form need not drive the other out of the market, due to efficient-market arguments.

[Figure 8 about here]

Turning to an investigation of relationship lengths, Figure 9 shows the cumulative density function for ultimate relationship lengths in T-Bonus, compared to the other treatments. The striking feature of the figure is the smaller number of long-term relationships

in T-Bonus compared. For example, 80 percent of all relationships are less than 4 periods long in T-Bonus. The next highest fraction, which is in T-EPL-Bonus, is only 64 percent. Thus, the pattern of firms substituting bonus incentives for relationships that we observed in T-EPL-Bonus is even more pronounced in the absence of the EPL institution. In a setting with the potential for long-term relationships, an important consequence of the bonus option is that it fewer of these relationships to occur.

[Figure 9 about here]

In summary, these findings enrich understanding of how wage and bonus contracts perform, , suggesting that the answer depends on the presence or absence of repeated interactions. They also show that T-Baseline and T-Bonus provide similar benchmark efficiency levels for the purposes of comparison to T-EPL and T-EPL-Bonus. Introducing EPL has a strong negative effect relative to these benchmarks in an efficiency wage world, but a much milder effect if introduced in an environment with a bonus option. This result highlights another relevant dimension for comparing performance of different contract forms, namely robustness to the introduction of EPL or firing costs. One implication of our study is that the level of efficiency provided by contracts with a bonus option is more robust in this sense.

6 Conclusion

Employment protection legislation (EPL) refers to rules limiting the ability of an employer to fire a worker. EPL is often believed to have a strong negative impact on market performance, but field evidence suggests a weak or zero effect. This paper uses an experiment to study EPL, in which firms and workers can endogenously engage in repeated interactions, and firms are unable to enforce effort through explicit contracts. We exogenously vary the presence of an EPL institution in the market, and the form of labor contract available to firms. The first main result is that EPL has a strong negative impact on market efficiency in the case of efficiency wage contracts. In reality, however, firms typically have the option to offer deferred compensation, such as a bonus. The second main result is that EPL has a much milder effect on market efficiency when firms have the option to pay bonuses. We show that EPL harms efficiency in an efficiency wage world because it attacks a key incen-

tive device: threat of firing. Given the option, it turns out that firms are able to credibly promise to reward effort with bonuses, so that the bonus substitutes for threat of firing as an incentive device. Thus, our findings suggest an explanation for the weak impact of EPL observed in the field: firms may be able to circumvent much of the negative effects of EPL by using bonus payments. In fact, we show that if firms are given the option, they endogenously make more use of bonuses in the presence of EPL. Given the salutary effect of the bonus option, it is natural to investigate whether it also improves market efficiency even in the absence of EPL. We find that bonuses perform better than wage contracts in early market periods, when most interactions are one-shot. In the long-run, however, when relationships have had time to develop and workers are faced with a meaningful firing threat, wage contracts are equally efficient.

References

- AKERLOF, G., AND J. YELLEN (1990): "The Fair Wage-Effort Hypothesis and Unemployment," *Quarterly Journal of Economics*, 105, 255 – 283.
- AUTOR, D., J. DONOHUE, AND S. SCHWAB (2006): "The Costs of Wrongful Discharge Laws," *Review of Economics and Statistics*, 88(2), 211 – 231.
- BANERJEE, A., AND E. DUFLO (2000): "Reputation Effects and the Limits of Contracting: A Story of the Indian Software Industry," *Quarterly Journal of Economics*, CXV, 989–1017.
- BESLEY, T., AND A. CASE (2000): "Unnatural Experiments? Estimating the Incidence of Endogenous Policies," *Economic Journal*, 110(467), F672–F694.
- BROWN, M., A. FALK, AND E. FEHR (2004): "Relational Contracts and the Nature of Market Interactions," *Econometrica*, 72(3), 747–780.
- CHEVALIER, J., AND G. ELLISON (1997): "Risk Taking by Mutual Funds as a Response to Incentives," *Journal of Political Economy*, 105, 1167–1200.
- (1999): "Career Concerns of Mutual Fund Managers," *Quarterly Journal of Economics*, 114, 389–432.
- FEHR, E., A. KLEIN, AND K. SCHMIDT (2007): "Fairness and Contract Design," *Econometrica*, 75(1), 121–154.
- FELLA, G. (2000): "Efficiency Wage and Redundancy Pay," *European Economic Review*, 44(8), 1473–1490.
- FUDENBERG, D., AND E. MASKIN (1986): "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54(3), 533–554.
- HONG, H., J. KUBICK, AND A. SOLOMON (2000): "Security Analysts' Career Concerns and Herding of Earnings Forecasts," *RAND Journal of Economics*, 31, 121–144.
- ICHINO, A., AND R. RIPHAN (2001): "The Effect of Unemployment Protection on Worker Effort: A Comparison of Absenteeism During and After Probation," IZA Discussion Paper No. 385.
- KUGLER, A., AND G. SAINT-PAULD (2004): "How Do Firing Costs Affect Worker Flows in a World With Adverse Selection?," *Journal of Labor Economics*, 22(3), 553–584.
- LAZEAR, E. (1990): "Employment Security Provisions and Employment," *Quarterly Journal of Economics*, 105(3), 699–726.
- LEMIEUX, T., B. MACLEOD, AND D. PARENT (2006): "Performance Pay and Wage Inequality," Unpublished Manuscript.
- MACLEOD, B. (2003): "Optimal Contracting with Subjective Evaluation," *American Economic Review*, 93(1), 216–240.
- (2006): "Regulation or Markets? The Case of Employment Contracts," *CEifo Economic Studies*, 51(1), 1–46.

NICKELL, S. (1997): "Unemployment and Labor Market Rigidities," *Journal of Economic Perspectives*, 11(3), 55–74.

OECD (1999): *Employment Outlook* Paris, France: OECD.

SHAPIRO, C., AND J. STIGLITZ (1984): "Equilibrium Unemployment as a Worker Discipline Device," *American Economic Review*, 74(3), 433–444.

Tables

Table 1: Effort Cost Schedule

Effort	1	2	3	4	5	6	7	8	9	10
Cost	0	1	2	4	6	8	10	12	15	18

Table 2: Probability that Worker Reduces Effort after First (Probation) Period of Relationship

Dependent Variable:	1 if effort drops in second period					
	T-EPL vs. T-Baseline (1)	T-EPL-Bonus vs. T-Baseline (2)	T-EPL-Bonus vs. T-EPL (3)	T-EPL-Bonus vs. T-EPL (4)	T-EPL-Bonus vs. T-Baseline (5)	T-EPL-Bonus vs. T-Baseline (6)
T-EPL	0.304** [0.121]	0.262** [0.084]				
T-EPL-Bonus			-0.326** [0.129]	-0.423*** [0.134]	-0.020 [0.105]	-0.077 [0.090]
Wage		-0.031*** [0.011]				
Previous wage		0.009** [0.010]				
Previous effort		0.205*** [0.055]		0.143 [0.132]		0.068 [0.049]
Desired effort		-0.038 [0.052]		0.026 [0.066]		-0.032 [0.075]
Previous desired effort		0.013 [0.038]		-0.174 [0.094]		-0.041 [0.045]
Period		0.021** [0.007]		0.004 [0.013]		-0.006 [0.009]
Total offered compensation				-0.050*** [0.019]		-0.024** [0.011]
Previous offered compensation				0.048*** [0.018]		0.021* [0.012]
Log pseudo-likelihood	-49.68	-31.06	-36.39	-27.22	-47.40	-37.29
Observations	84	84	61	61	89	89

Coefficients are marginal effects from Probit regressions. The sample includes only private offer contracts that take place in the second period of a long-term relationship between firm and worker. The dependent variable in Columns (1) to (6) is equal to 1 if worker effort is lower in the current (second) period than in the previous (first) period of the relationship, and zero otherwise. The omitted category in Columns (1) and (2) is T-Baseline, the omitted category in Columns (3) and (4) is T-EPL, and the omitted category in Columns (5) and (6) is T-Baseline. Robust standard errors in brackets, clustering on session; ***, **, * indicate significance at 1-, 5-, and 10-percent level, respectively.

Table 3: Worker Effort in Initial (Probation) Period of Relationships

Dependent Variable:	Effort					
	T-EPL vs. T-Baseline (1)	T-EPL-Bonus vs. T-Baseline (2)	T-EPL-Bonus vs. T-EPL (3)	T-EPL-Bonus vs. T-EPL (4)	T-EPL-Bonus vs. T-Baseline (5)	T-Baseline (6)
T-EPL	1.985*** [0.487]	1.672*** [0.388]				
T-EPL-Bonus			-0.365 [0.462]	-1.353** [0.580]	1.684*** [0.529]	0.130 [0.529]
Wage		0.125*** [0.014]				
Desired effort		0.198*** [0.071]		0.735*** [0.158]		0.291** [0.142]
Period		-0.087*** [0.025]		-0.082*** [0.031]		-0.111*** [0.035]
Total offered compensation				0.037 [0.029]		0.087*** [0.030]
Constant	4.118*** [0.376]	-0.628 [0.747]	6.115*** [0.307]	-0.009 [0.638]	4.073*** [0.389]	-0.014 [0.791]
hline						
Log pseudo-likelihood	-1306	-1146	-649	-593	-1224	-1130
Observations	583	583	288	288	545	545

Coefficients are marginal effects from Probit regressions. The sample includes all private offer contracts that are not yet part of a long-term relationship, i.e., probation period contracts. The dependent variable in Columns (1) to (6) is the worker's chosen effort level. The omitted category in Columns (1) and (2) is T-Baseline, the omitted category in Columns (3) and (4) is T-EPL, and the omitted category in Columns (5) and (6) is T-Baseline. Robust standard errors in brackets, clustering on session; ***, **, * indicate significance at 1-, 5-, and 10-percent level, respectively.

Table 4: Impact of EPL on use of Bonus Pay

Dependent Variable:	Offered bonus		
	T-EPL-Bonus vs. T-Bonus		
	(1)	(2)	(3)
T-EPL-Bonus	5.340*	6.529**	6.508**
	[2.953]	[2.910]	[2.930]
Wage	-0.680***	-0.708***	0.707***
	[0.081]	[0.072]	[0.073]
Desired Effort		2.941***	2.910***
		[0.571]	[0.618]
Period			0.050
			[0.132]
Constant	42.460***	15.694***	15.323***
	[2.770]	[5.212]	[5.179]
Observations	455	455	455

Interval regression coefficient estimates. The omitted category in all columns is T-Bonus. Robust standard errors in brackets, adjusted for clustering on session; ***, **, * indicate significance at 1-, 5-, and 10-percent level, respectively.

Figures

Figure 1: Baseline Market Performance and the Threat of Firing

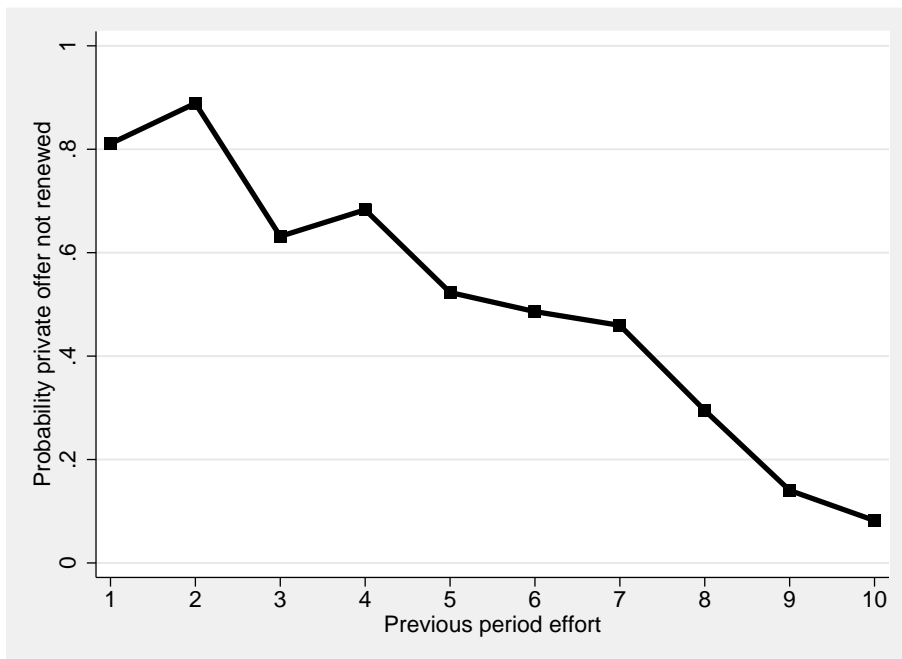



Figure 2: The Impact of EPL on Market Efficiency

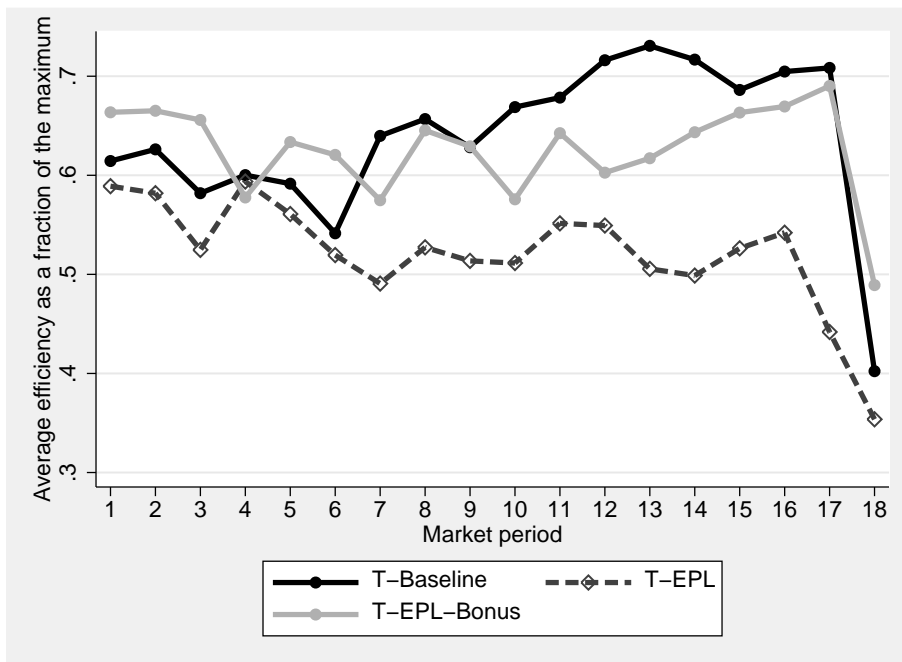
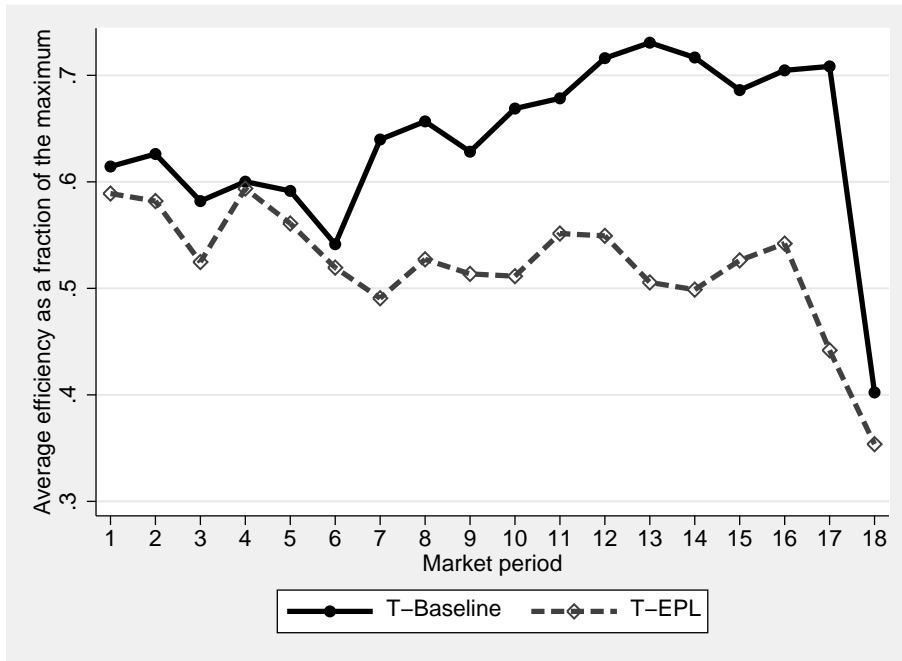


Figure 3: The Impact of EPL on Effort Levels in Long-Term and One-Shot Relationships

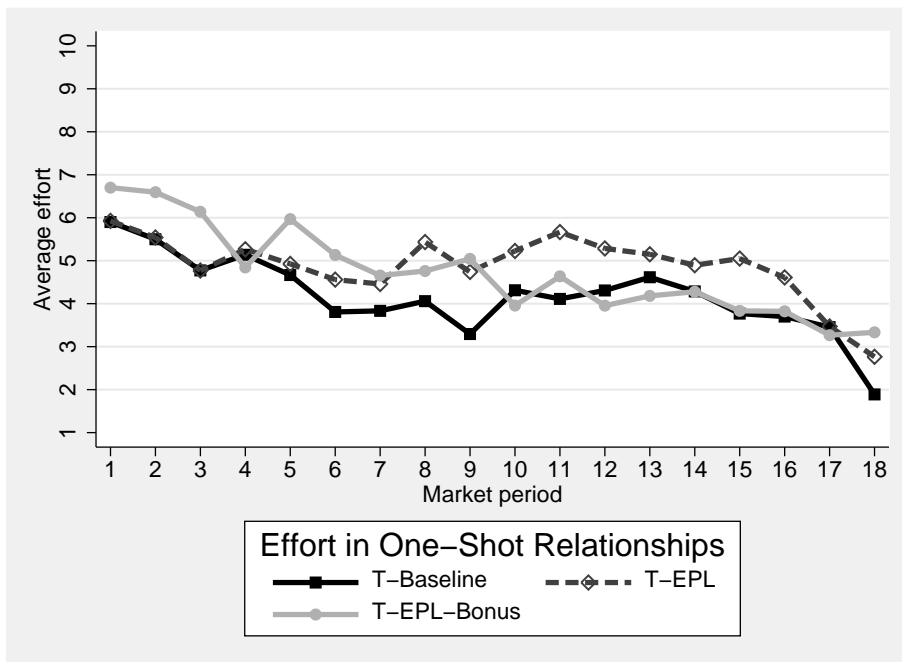
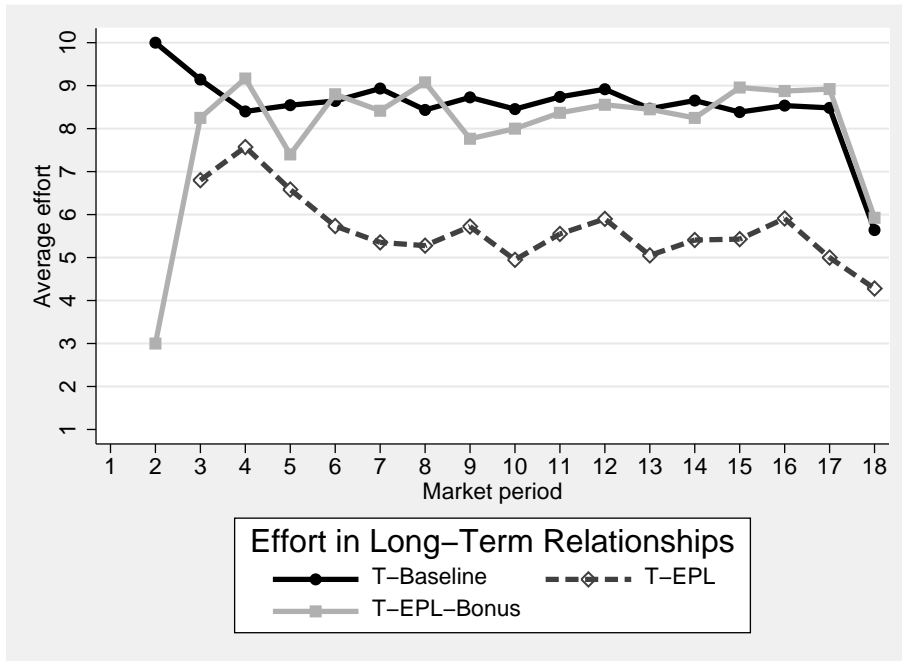


Figure 4: The Effort-Bonus Relation




Figure 5: Cumulative Distribution of Relationship Lengths

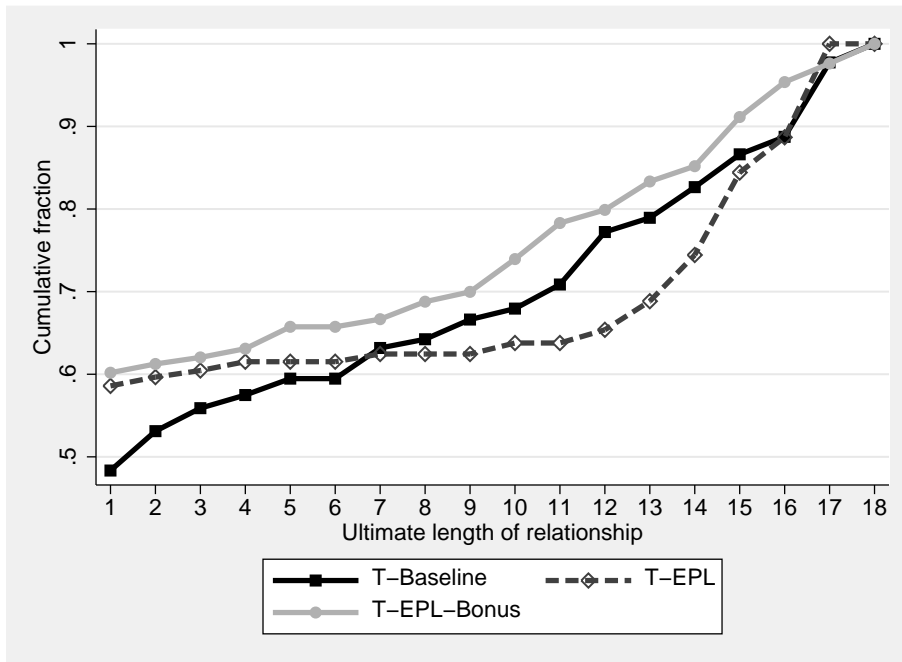
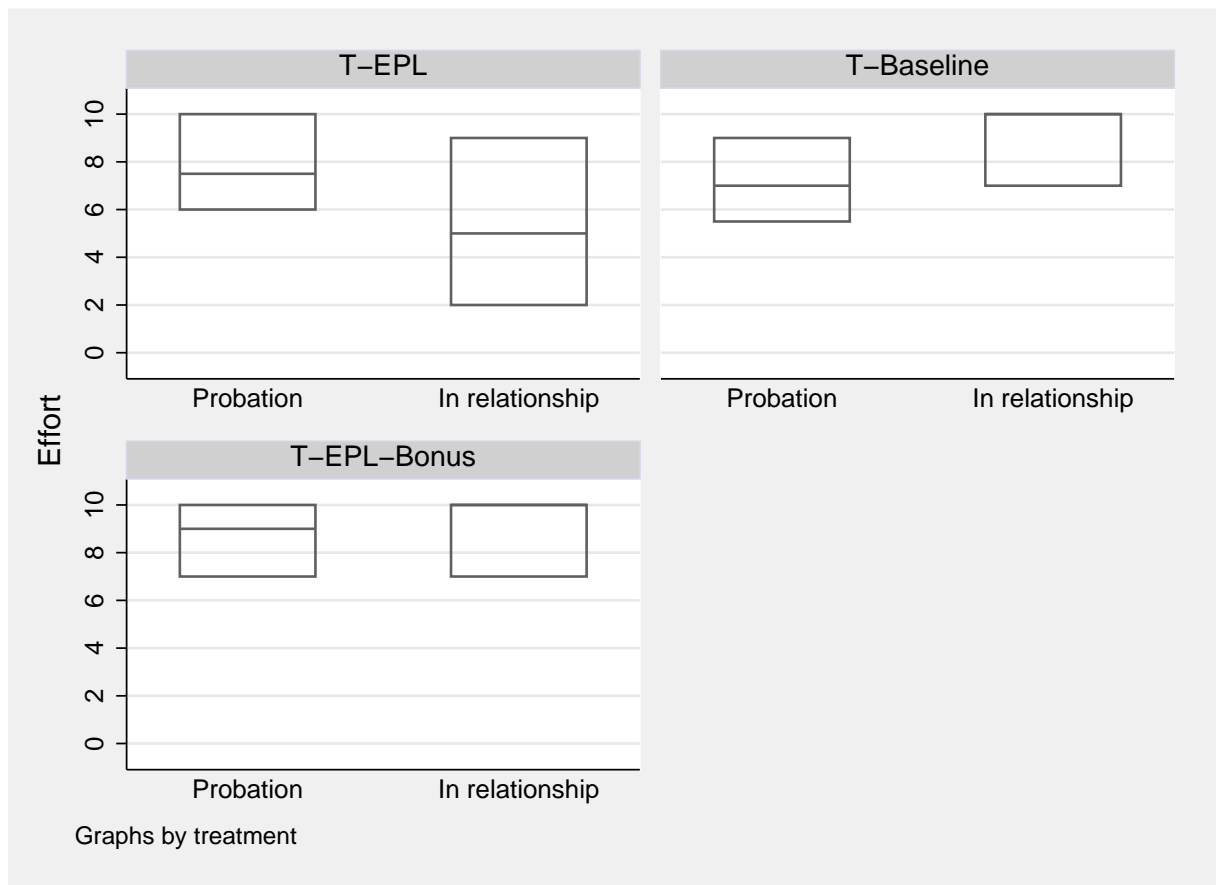


Figure 6: Effort Distributions in Probation and Later Relationship Periods



Notes: The height of a box shows the range from the 25th to the 75th percentiles of the distribution. The center line in each box is the median.

Figure 7: The Impact of EPL on the Use of Wage and Bonus Pay in Long-Term Relationships

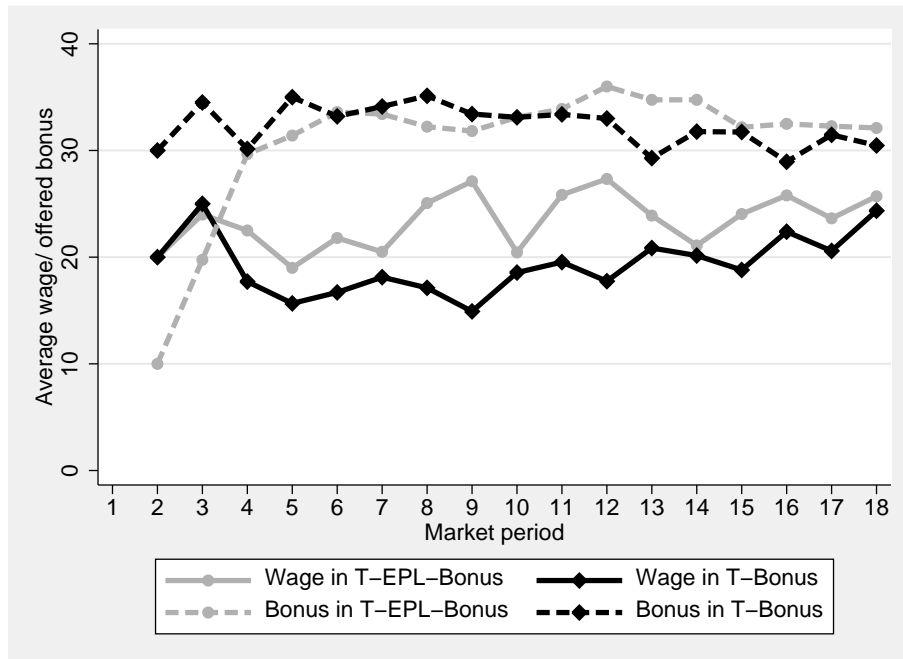


Figure 8: Bonus Pay and Market Efficiency in the Absence of EPL

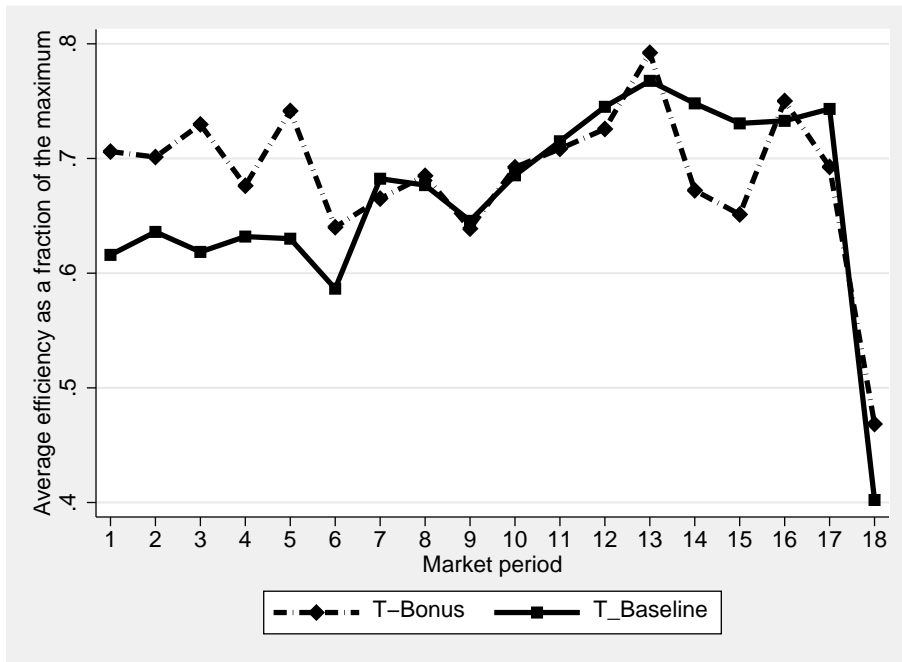


Figure 9: Cumulative Distribution of Relationship Lengths

