When Protection Becomes Exploitation: The Impact of Firing Costs on Naïve Employees*

Florian Englmaier† Matthias Fahn‡ Ulrich Glogowsky§
Marco A. Schwarz¶

November 11, 2025

Abstract

This paper identifies a new channel through which employment protection laws can harm workers: they enable employers to exploit naïve present-biased employees. The theoretical mechanism through which firing costs enable exploitation is built upon a widespread career structure. Employers frequently offer career trajectories with an unattractive entry phase (e.g., with low pay or high effort), an even less attractive mid-career stage, and high rewards promised in a final career phase. Naïve employees accept such contracts, expecting to complete the high-reward track. They, however, eventually drop out or sort into alternative paths mid-career when costs loom large and rewards remain distant. As a result, employers avoid paying the promised rewards while capturing surplus from early-career phases. Firing costs give employers more leeway to exploit workers with such contracts: they reduce the employees' expected risk of dismissal and make long-term rewards appear more credible. Employers can then exploit employees even more—for example, by lowering early-career wages—knowing that employees will still accept such steeper contracts despite not following through. Our model aligns with career patterns in fields such as healthcare, academia, or accounting. It also yields testable predictions on wage structures, attrition, and how firing costs and labor market conditions shape contract design.

Keywords: Employment protection laws, present bias, dynamic contracting

JEL Codes: D21, D90, J33, K31, M52

---

*We are grateful to Daniel Gottlieb, Fabian Herweg, Paul Heidhues, Nicholas Li, Takeshi Murooka, Tomasz Sulka, and seminar participants at HHU Düsseldorf, JKU Linz, LMU Munich, as well as the DICE/ZEW Winter School, for helpful comments. Financial support by the DFG through CRC TRR 190 (Project Number 280092119), the Austrian Central Bank (Anniversary Fund Number 18796), and the Austrian Ministry of Education, Science, and Research through LIFT_C is gratefully acknowledged.

†LMU Munich, Organizations Research Group (ORG), CEPR, CESifo, and IZA, florian.englmaier@econ.lmu.de

‡University of Hong Kong, CESifo, and IZA, mfahn@hku.hk

§Johannes Kepler University Linz and CESifo, ulrich.glogowsky@jku.at

¶Düsseldorf Institute for Competition Economics (DICE), Heinrich Heine University Düsseldorf, ICES, and CESifo, marco.schwarz@hhu.de.

1 Introduction

Employment protection laws (EPLs) are prevalent worldwide. A common feature of these laws is that they impose firing costs on employers. Hereby, policymakers hope to secure the stability of the job of employees (Betcherman, 2013), prevent unfair treatment (OECD, 2013), or foster the development of firm-specific human capital through sustained employment relationships (Pierre and Scarpetta, 2004; Belot et al., 2007; Acharya et al., 2013). Although EPLs appear effective in achieving these objectives (Betcherman, 2013), their overall benefit to employees remains unclear.[1] A particularly underexplored aspect is whether firing costs allow employers to reorganize their labor contracts in ways that adversely affect employees' career trajectories and wages.

This paper demonstrates that firing costs can, indeed, have harmful effects on employees: they allow employers to exploit early-career employees who are not fully rational. Consequently, well-intentioned policies such as EPLs may unintentionally benefit employers at the expense of the very employees they are meant to protect. To reach this conclusion, we depart from conventional models that assume individuals are fully rational and time consistent. As widely documented in the literature (DellaVigna, 2009; Fang and Wang, 2015; Chan, 2017; Cheung et al., 2021), many people are present-biased—putting extra value on current over future consumption—and are naïve about it, i.e., they expect not to be present-biased in the future. Using a principal-agent model that incorporates these insights, we demonstrate that negative consequences of EPLs for employees can emerge.

The theoretical mechanism through which firing costs enable exploitation builds on a widespread career structure: Employers often offer career paths that include (a) an unattractive initial phase with low pay, high effort, or low autonomy (*entry phase*), (b) an even more burdensome mid-career phase (*qualification phase*), and (c) a promised high-reward phase at the end of one's career (*reward phase*). This final phase may involve a high-paying or prestigious position, such as a leadership role. Structures of this kind are common in professions such as healthcare, academia, and accounting.[2] If employees are present-biased and (partially) naïve about it, they may pursue such careers with the expectation of eventually attaining the final reward. However, they often abandon the path midway, when demands are highest and the promised rewards still seem distant—a pattern observed in the settings

[1] There is some past research highlighting potential indirect adverse effects of these policies for employees. Higher firing costs, for example, distort employers' incentives to create new jobs. As a result, employment protection laws can erode overall employment or affect the dynamics of labor markets (Bertola and Rogerson, 1997; Hopenhayn and Rogerson, 1993; Mortensen and Pissarides, 1994).

[2] In healthcare, specialists begin with years of training (entry), proceed through a demanding residency with low pay and low autonomy (qualification), and only later access the financial and professional rewards of senior roles (reward). Similarly, academics begin with a PhD (entry), progress through an insecure and competitive postdoctoral or junior faculty stage (qualification), and may eventually secure a tenured professorship (reward). Accounting offers another example. Here, junior staff commit early (entry), face years of high-intensity client work under pressure to advance (qualification), and a few reach partnership (reward).

mentioned above.[3] This behavior explains why employers may find it optimal to design such paths: Anticipating that only a few employees will persist until the end, they can extract surplus during the early stages while avoiding the cost of long-term rewards. Higher firing costs enhance this mechanism. They reduce the likelihood of dismissal, thereby making the promises of future rewards more credible. As a result, employers can capture even more surplus before the final reward stage by structuring contracts to be more steeply backloaded.[4]

By highlighting this effect, our paper contributes to the vast existing literature on how firms can exploit present-biased consumers (see Kőszegi, 2014 or Heidhues and Kőszegi, 2018 for literature reviews). While previous studies have discussed the role laws and market characteristics play in mitigating or enforcing this exploitation (Handel, 2013; Ericson, 2014; Sulka, 2024), there is a notable lack of understanding about the influence of labor market institutions on firms' ability to exploit present-biased employees. This oversight represents a critical gap in the literature, especially considering (a) the growing evidence that present bias matters in the workplace (Kaur et al., 2015, Mas and Pallais, 2017) and (b) employers more and more leverage people analytics methods and big data to learn about their employees' characteristics and biases. Given their relevance, it is, therefore, crucial to dissect how employers may capitalize on their employees' psychological tendencies and how policies influence this behavior.

Details of baseline model Moving to a more formal description of our setting, we base our analysis on the following principal-agent model. A risk-neutral principal (she) and a risk-neutral agent (he) interact over three periods: the entry phase (period 1), the qualification phase (period 2), and the reward phase (period 3). The principal discounts future profits exponentially; the agent is present-biased and discounts his future utility in a quasi-hyperbolic way (Laibson, 1997). At the beginning of the first period, the principal offers a long-term contract to the agent. While the long-term contract determines both parties' obligations in the event of continued employment, either the principal or the agent can terminate the relationship at the beginning of the second and third periods. A termination

[3] Subsection 4.3 discusses evidence suggesting that divergence from the intended path is widespread in these professions. Empirically, workers may diverge from their intended career paths in multiple ways. Some jobs offer multiple career trajectories from the outset. Employees initially aim for a high-reward path but later sort into a flatter alternative (by choice or implicitly). Other jobs dictate single, structured career trajectories, but employees ultimately abandon the original path and move into different (flatter) careers. Importantly, sorting is often embedded in the contract structure or occurs gradually, without a formal choice. For example, in academia, postdoctoral researchers aim for a professorship. However, if they fail to meet publication expectations before their contract ends, they are often rehired into different roles with flatter trajectories.

[4] Our model deliberately abstracts from the positive roles that firing costs can play in practice (mentioned in the first paragraph of the introduction). We set these aspects aside to isolate a specific mechanism: how firms can use firing costs strategically to exploit present-biased and naïve agents. This simplification is not meant to suggest that firing costs are generally harmful, but rather to highlight one unintended consequence that arises in the presence of present bias.

by the principal requires her to pay a fixed firing cost K, determined by the severity of employment protection laws.[5] By contrast, the agent is always free to leave at no cost. The principal's employment offer contains a wage in exchange for the costly effort exerted by the agent, with effort being verifiable.

Our first contribution is to determine the profit-maximizing contracts for an agent with and without a present bias. If the agent is not present-biased or if he is present-biased but sophisticated (i.e., fully aware of his bias), short-term incentives are optimal (i.e., payments for effort are made in the same period as it is exerted). Intuitively, because effort is verifiable, such short-term incentives secure the first-best effort and leave the agent with his outside option.

By contrast, when employing a naïve agent who is not aware of his future present bias, the principal designs a long-term contract that specifies (a) a wage payment (and first-best effort) in period 1 and (b) a menu of career paths among which the agent can choose in period 2. This menu consists of a "*virtual path*" the agent naïvely intends to choose and a "*real path*" he inadvertently ends up selecting. While the real path contains wage payments that cover the agent's respective effort costs, the principal designs the virtual path so that period 2 serves as the "qualification period," in which the agent's utility is low. In period 3, the virtual path promises the agent a high utility level.

An important question is why it is optimal for the principal to offer this menu of career paths. From the perspective of period 1, which involves an extra weight on period-1 utility but the same weights on utilities in periods 2 and 3, the agent would optimally select the virtual career path in the subsequent period. However, when period 2 comes, the agent puts a higher weight on period-2 than on period-3 utility. Therefore, the relative costs of the qualification period loom larger than they did from the perspective of period 1. He is consequently willing to sacrifice the high period-3 rent in exchange for a moderately higher current period-2 payment, which the real path provides.

Because the naïve present-biased agent does not anticipate his eventual choice of the real career path, the rent promised in the virtual path makes him willing to accept a lower compensation in the first period, leaving him with a utility below his outside option. This dynamic makes it optimal for the principal to offer a career path that is steeply back-loaded: low utility in the early stages, with high rewards only in later periods. The principal, therefore, transforms an inherently static contracting setting—effort can be verified and compensated in the same period as it is exerted—into a dynamic contract. Because the agent is naïve present-biased, he cannot overcome the barriers established by the principal in the form of the qualification period. Consequently, he picks the flat compensation scheme provided by the real career path in period 2, despite having accepted a low initial wage in hopes of a

[5] A common interpretation of firing costs is understanding them as a "tax on job destruction." This tax typically reflects real costs associated with separations. Because it is paid outside the firm-worker pair, the firm cannot include it in the wage bargaining process (Bertola and Rogerson, 1997).

future payoff that never materializes. Such exploitative contracts can be understood as a bet à la Crémer and McLean (1988). The principal anticipates that the agent will act in a present-biased way, whereas the agent believes he will remain time-consistent. The menu of career paths then acts as a wager: the principal bets on the real path, the agent on the virtual one, and it is precisely this divergence that the contract is designed to monetize.

Building on this baseline model, our second contribution is to show that higher *firing cost K* allow the principal to exploit early-career agents by lowering their contractual wages more extensively. Specifically, firing costs affect the structure of both the real and virtual career paths. To understand why, note that the extent to which the principal can exploit the agent during the beginning of his career depends on how appealing the third period appears in the virtual path. The principal's credibility in making such promises is constrained by her ability to fire the agent. Thus, higher firing costs enhance the firm's credibility by making layoffs more expensive, thereby making the virtual path seem more attractive to the agent.

This increased credibility allows the principal to extract greater surplus at the agent's expense. Notably, while higher firing costs lead to lower starting wages in period 1, wages in periods 2 and 3 remain unchanged. In this way, firing costs may not only help (for reasons outside of our model) employees but also reduce their early-career earnings. As a result, the negative impact of firing costs is likely to be most severe for younger employees in new job matches, while older or more established workers are less affected.

Extensions We extend our baseline model along several dimensions and show that the mechanism remains robust. First, the exploitation mechanism survives even when the agent has bargaining power. While surplus sharing reduces the principal's ability to exploit the agent, it does not eliminate the incentive to defer compensation. Second, similar to the effect of bargaining power, labor market competition constrains the firm's ability to extract surplus through back-loaded contracts. However, the key mechanism remains intact as well. Third, we relax the assumption of full naïveté. The results remain robust: The offered contract does not depend on the agent's degree of naïveté, unless he is fully sophisticated. We also incorporate frequency naïveté as described by Eliaz and Spiegler (2006). In this setting, the agent remains present-biased but is uncertain about his future behavior: with probability q, he believes he is time-consistent and will follow through with the virtual path (reflecting full naïveté), while with probability $1-q$, he recognizes his present bias and anticipates abandoning the path later on (reflecting sophistication). Our analysis shows that as the agent's confidence in his ability to stick to the virtual path increases, the potential for exploitation by the principal grows, and the effect of firing costs becomes more significant. Fourth, we analyze settings with asymmetric information and again find that exploitation may remain optimal. Specifically, when the principal cannot observe whether an agent is naïve or time-consistent, she offers the exploitation contract if the expected share of naïve agents exceeds a cutoff. As a result, the principal either engages in exploitation or not, depending on the expected type distribution. Fifth, we consider overconfident workers who underestimate future effort costs. In the adjusted model, the first period of the virtual path might not only include higher effort but also a high rent. However, the role of firing costs remains the same. Such costs increase the firm's credibility in promising future rewards and, thereby, enable greater exploitation. In conclusion, across all extensions, the mechanism continues to generate dynamic exploitation. The principal shifts compensation to later periods and extracts surplus early on. Moreover, higher firing costs lead to more exploitation as they enhance the credibility of promised rewards.

Testable predictions as guidance for future work To guide future empirical work, we also formulate five sets of testable predictions that follow from our model. They link naïve present bias and firing costs to the structure of contracts and the dynamics of careers. Specifically, the predictions describe (a) how firms back-load wages when workers are naïve present-biased, (b) how naïve present-biased workers behave over time, (c) how firing costs shape contract design, and (d) how bargaining power or (e) labor-market competition moderate contracts. The core of all these predictions is that firms exploit present-biased workers by shifting burdens to early stages and deferring rewards. We also briefly discuss how follow-up studies can test our predictions using variation across workers, firms, or labor market institutions.

Evidence Several suggestive pieces of evidence align with our key predictions. For example, in fields such as healthcare (Goldacre et al., 2010; Lambert and Goldacre, 2011), academia (Ehrenberg et al., 1998; Sauermann and Roach, 2012; Roach and Sauermann, 2017; Denton et al., 2022), and accounting (Deloitte Insights, 2016; Going Concern, 2023; CFO, 2023; National Pipeline Advisory Group, 2023), workers often start with ambitious plans to pursue high-upside tracks and expect to follow through. In reality, however, many end up changing course when the challenges of mid-career stages become salient. This gap between initial intentions and actual outcomes is a central prediction of our model, and, in theory, arises from the interplay between back-loaded incentives and agents' misprediction of their own future behavior. Moreover, firing costs in these professions are especially high at senior levels of the hierarchy and, in some cases, even exceed those set by general employment protection laws. We argue that these features help explain the prevalence of steep, back-loaded career paths that enable the exploitation of workers in earlier stages. Also consistent with our mechanism, Carrell and Zinman (2014) present evidence suggesting that individuals who are more likely to be present-biased are also more likely to exit long-term reward tracks prematurely.[6] Other papers confirm our predictions on the role of

[6]They exploit the random assignment of U.S. Air Force personnel to bases in different U.S. states. Some states allow access to payday loans, while others restrict them. They find that, when payday loans are avail-

firing costs. For example, Leonardi and Pica (2013) show that an increase in firing costs in Italy lowered wages for new hires, especially for young workers. The wages of the incumbents remained unchanged. Similarly, Cervini-Pla et al. (2014) find that reducing firing costs increased entry wages. Lastly, in line with our extension that accounts for differences in bargaining power, Leonardi and Pica (2013) find that the negative effect of firing costs on wages is particularly strong for workers with low bargaining power.[7] These findings are difficult to reconcile with standard models but naturally follow from the behavioral and institutional dynamics we highlight.

Literature Our paper relates to three broad fields of literature. First, we contribute to the literature highlighting potential drawbacks of employment-protection policies. Stringent hiring and firing laws can, for example, limit firms' ability to adapt quickly to changes in demand and technology (Kuzmina, 2023) or change the mix of skills workers invest in (Estevez-Abe et al., 2001; Wasmer, 1999). The policies may also undermine labor mobility from declining sectors to new dynamic sectors and, thus, affect (a) the efficient allocation of labor, (b) productivity, or (c) even economic growth (Hopenhayn and Rogerson, 1993; Belot et al., 2007). We add to this literature by demonstrating that employment-protection policies, when interacting with behavioral factors such as present bias, can lead to further unintended consequences that traditional analyses overlook. Specifically, our findings reveal a nuanced dynamic in which these policies, despite their protective intent, can inadvertently support the exploitation of behavioral tendencies and result in suboptimal outcomes for workers. This underscores the need for a more holistic approach in policy formulation that considers psychological insights to ensure the well-being of employees in the labor market.

Second, we contribute to the literature on the implications of present bias for policy design. The earlier papers mainly frame this bias as a form of mis-optimizations that leads to "behavioral mistakes" (such as exercising too little, smoking too much, or under-saving for retirement). Building on this idea, most of the policy-related papers highlight how governments can correct such mistakes with policies such as creating optimal defaults in health insurance (Handel, 2013; Ericson, 2014, 2020), sending reminders (Ericson, 2017), bringing together the time of a decision and its effect (Murooka and Schwarz, 2018, 2019; Johnen, 2019), or setting up mandatory pensions (Sulka, 2024). Recently, several empirical studies have emerged, indicating that the present bias also affects labor supply decisions (Kaur et al., 2015; Mas and Pallais, 2017). The underlying idea is that employment relationships frequently reward up-front effort with future benefits. Present-biased employees may then

able, job performance deteriorates and early attrition increases. In line with present bias, the effects are stronger for younger and financially unsophisticated airman.

[7] A higher bargaining power of the agent reduces the principal's ability to take advantage of the agent's present bias. Thus, the negative effect of a higher firing cost on the agent's compensation gradually diminishes as the agent's bargaining power increases.

inflate their perception of the immediate effort costs and, consequently, exert less effort than their "long-run self" would prefer. Along these lines, the present bias influences not only the design of health and savings policies but also that of labor-market policies. Lockwood (2020), for example, shows that the present bias reduces the optimal income tax rate, especially if the elasticity of the taxable income is high. In our view, however, an in-depth analysis of employment protection policies in the presence of present bias is still missing. We close this gap by analyzing the problem through the lens of a simple principal-agent model.

Third, our paper relates to the literature on optimal exploitation contracts when workers are present-biased. This literature is based on the behavioral IO literature, which has demonstrated that firms can extract rents from consumers who are unaware of future biases and are induced to pay high fees when changing their original plans (DellaVigna and Malmendier, 2004; Eliaz and Spiegler, 2006; Heidhues and Kőszegi, 2010). Gottlieb and Zhang (2021) show that the inefficiency losses of such contracts diminish as the time horizon grows. In an employment setting, Gilpatric (2008), Li et al. (2012), and Yılmaz (2013) study the implications of an employee's present bias if there is moral hazard. We analyze how a naïve employee's naïveté affects contract dynamics, a dimension not explored by these contributions.

The papers closest to ours are Englmaier et al. (2023) and Yu (2021). In contrast to Englmaier et al. (2023), our model allows the principal to terminate the relationship, however, at some cost, and she can condition payments on effort rather than on outcomes. Additionally, our focus is on how present bias influences labor market policies. Yu (2021) explores optimal retirement policies for present-biased agents within a Mirrlees taxation framework. The government offers a menu containing a real allocation with moderate savings and an imaginary allocation with back-loaded rewards. Naïve agents believe they will choose the imaginary option, so they work harder upfront to qualify for it. When the future arrives, however, their present bias leads them to abandon the imaginary option and take the real one. The government, thus, incentivizes effort with promises it never has to fulfill. Finally, Fahn and Seibel (2022) also explore the role of commitment in employment relationships. They show that, if a firm is not able to commit to long-term contracts, naïve agents overestimate the extent of future wage reductions due to non-monetary benefits of employment, leading them to accept smaller reductions in the present. This finding suggests that present-biased employees can also benefit from being naïve. While Fahn and Seibel (2022) focus on a setting where today's effort increases tomorrow's benefits, our paper shows that even with a static production technology, a dynamic compensation system can emerge as it allows firms to increase profits and to this end exploit naïve employees.

2 Model Setup

Technology A risk-neutral principal ("she") can hire a risk-neutral agent ("he") for three periods, $t \in \{1, 2, 3\}$. If employed in period t, the agent receives a wage $w_t \in \mathbb{R}$ and chooses the verifiable effort level $e_t \geq 0$. The costs of effort $c(e_t)$ are strictly increasing, differentiable, and convex (with $c(0) = c'(0) = 0$). Denoting the marginal value of the agent's effort by $\theta > 0$, we assume that the effort level e_t generates a deterministic output $e_t\theta$ that the principal consumes.

Given these assumptions, the agent's payoff in period t when employed by the principal is

$$w_t - c(e_t).$$

The principal obtains

$$e_t\theta - w_t.$$

If the agent does not work for the principal in period t, he receives his outside option $\overline{u} \in \mathbb{R}^+$; the principal's outside option is normalized to zero.

The level of effort that maximizes the total surplus if the agent works for the principal, the first-best effort denoted by e^{FB}, is defined by

$$\theta - c'(e^{FB}) = 0.$$

Time preferences While the principal discounts the future exponentially with a constant factor $\delta \in (0, 1]$, the agent applies quasi-hyperbolic discounting to future payoffs (Phelps and Pollak, 1968; Laibson, 1997): From the perspective of period $t = 1$, the agent discounts future payoffs with $\beta\,\delta$ (period $t = 2$) or $\beta\,\delta^2$ (period $t = 3$), with $\beta \in (0, 1]$; the discounting between payoffs in periods 2 and 3 is exponential, at rate δ. From the perspective of period $t = 2$, the agent discounts period-3 payoffs with $\beta\,\delta$. Hence, the agent is present-biased, and his preferences are dynamically inconsistent.

The agent may misunderstand his own future time preferences, mistakenly believing that he will be time-consistent in the future. In this situation, he is considered *naïve*. This stands in contrast to a *sophisticated* agent, who fully recognizes his tendency for present bias in the future. In Section 6.2, we also discuss cases of partial naïveté. First, we analyze the scenario where the agent expects his future present bias to be determined by the factor $\hat{\beta}$ (where $\beta < \hat{\beta} \leq 1$), as described by O'Donoghue and Rabin (2001). Second, we consider frequency naïveté (Eliaz and Spiegler, 2006), in which the agent believes that he will be time-consistent in the future with probability q, and present-biased with probability $1 - q$. In the following analysis, we focus on the consequences of the agent's present bias and, thus, set

$$\delta = 1.$$

We impose this assumption solely for simplicity; it does not qualitatively affect our results.

Perceptions We assume common knowledge about the principal's time preferences. On the contrary, the agent's time preferences are not common knowledge. Although the principal knows the agent's time preferences and his value β, the agent believes the principal shares his own (potentially incorrect) self-perception. A naïve agent is, therefore, convinced that the principal also perceives him to be time-consistent in the future. This assumption borrows from the behavioral IO literature, which posits that firms, through their experience, understand the agents' systematically changing preferences better than the agents themselves (Eliaz and Spiegler, 2006).

Contracts and commitment The principal can commit to long-term contracts but has the option of firing the agent at the beginning of periods 2 and 3 at a firing cost of $K > 0$. The firing decision is irreversible; subsequently, the principal and the agent consume their outside option utilities in the subsequent periods. The value of K captures the extent of employment protection in the economy, with higher values indicating a more stringent employment protection. Note that the assumption that K is identical in both periods does not affect our results. The reason is that firing costs are only relevant in period 3. We demonstrate this in Section 5.1, where we show how the firing cost K affects the payoffs of both the principal and the agent. Furthermore, for now, we abstract from severance payments (i.e., payments that the agent receives after termination), but discuss them in Section 8.

For the remainder of this paper, we assume

$$e^{FB}\theta - c(e^{FB}) - \bar{u} > -K,$$

indicating that firing the agent is inefficient if he exerts e^{FB}.

The agent cannot commit to long-term contracts and is free to leave at the beginning of every period. Moreover, his effort is verifiable; thus, forcing contracts that specify the required effort level the agent has to exert are possible. Our results would remain unchanged if the agent, instead, did not receive the wage w_t when deviating from the contractually specified effort.

Now, in $t = 1$, the principal makes a take-it-or-leave-it contract offer to the agent. This offer contains the wage and effort for period 1 and a *menu of career paths*, denoted by C. The agent can select one element from C, labeled $i \in \{1, 2, ..., I\}$, at the beginning of period 2. Each element in C specifies wages and efforts for the next two periods, thus

$C = \left\{ \left(w_2^i, e_2^i, w_3^i, e_3^i \right)_{i=1}^{I} \right\}$. Without loss of generality, we can restrict I to 1 or 2, depending on the extent of naïveté of the agent. If the agent is sophisticated or time-consistent, he correctly anticipates his future behavior, in which case the principal sets $I = 1$. By contrast, if the agent is naïve, the principal optimally sets $I = 2$ so that the menu consists of two paths: one that the agent believes he will choose in period 2 (*virtual path*) and one that he actually selects (*real path*). We refer to the virtual path with a superscript "v" and to the real path with the superscript "r." Thus, with a slight abuse of notation, the menu becomes $C = \left\{ \left(w_2^r, e_2^r, w_3^r, e_3^r \right), \left(w_2^v, e_2^v, w_3^v, e_3^v \right) \right\}$.

Therefore, unlike the literature on adverse selection—where principals offer a menu of contracts to screen agents with private information—in our setting, the principal has full knowledge of the agent's characteristics and can perfectly predict the agent's future choices. The principal's objective is to take advantage of the agent's misperceptions by offering one contract the agent expects to choose in the future and another one that the agent will ultimately prefer, but which yields higher profits for the principal. Note that Section 6.3 discusses asymmetric information about the agent's time preferences and perceptions.

Payoffs Next, we describe the real and perceived payoff streams along the equilibrium path where the agent is (and anticipates to be) employed in every period t. His realized utility streams equal

$$\begin{aligned}
U_1^r &= w_1^r - c(e_1^r) + \beta \left(w_2^r - c(e_2^r) + w_3^r - c(e_3^r) \right) \\
U_2^r &= w_2^r - c(e_2^r) + \beta \left(w_3^r - c(e_3^r) \right) \\
U_3^r &= w_3^r - c(e_3^r).
\end{aligned}$$

These U_t^rs correspond to the utilities that a sophisticated or time-consistent agent receives (in the latter case with $\beta = 1$).

Naïve agents expect to select the virtual path in period 2; thus, their perceived utility streams *from the perspective of period* 1 are

$$\begin{aligned}
U_1^v &= w_1^r - c(e_1^r) + \beta \left(w_2^v - c(e_2^v) + w_3^v - c(e_3^v) \right) \\
U_2^v &= w_2^v - c(e_2^v) + w_3^v - c(e_3^v) \\
U_3^v &= w_3^v - c(e_3^v).
\end{aligned}$$

The principal's payoffs are

$$\begin{aligned}
\Pi_1^r &= e_1^r \theta - w_1^r + e_2^r \theta - w_2^r + e_3^r \theta - w_3^r \\
\Pi_2^r &= e_2^r \theta - w_2^r + e_3^r \theta - w_3^r \\
\Pi_3^r &= e_3^r \theta - w_3^r,
\end{aligned}$$

while the naïve agent, from the perspective of period 1, perceives them to be

$$
\begin{aligned}
\Pi_1^v &= e_1^r\theta - w_1^r + e_2^v\theta - w_2^v + e_3^v\theta - w_3^v \\
\Pi_2^v &= e_2^v\theta - w_2^v + e_3^v\theta - w_3^v \\
\Pi_3^v &= e_3^v\theta - w_3^v.
\end{aligned}
$$

Strategies and equilibrium Following O'Donoghue and Rabin (1999), we describe the strategies of the players using the term *perception-perfect strategy*. Such a strategy specifies the actions of the player based on dynamically consistent beliefs about their future behavior. While a time-consistent or sophisticated agent correctly anticipates his future actions, a naïve agent has incorrect beliefs about his future time preferences.

We denote a principal's strategy by σ_P. In period $t = 1$, this strategy determines the long-term contract C. In periods $t = 2, 3$, σ_P specifies whether the principal adheres to the contract or fires the agent at a cost K. Similarly, we denote the agent's strategy by σ_A. His strategy determines in each period whether the agent works for the principal (and exerts the contracted effort level e_t) or opts for his outside option. In period 2, σ_A also specifies his choice from C.

We apply the concept of *perception-perfect equilibrium*. This equilibrium maximizes each player's payoff, given their perception of their own and the other player's future behavior. Because the principal can make a take-it-or-leave offer at the start of period 1, she offers the menu C that maximizes Π_1^r. In all later periods, her decision revolves around firing the agent or not, doing so only if it is optimal. The naïve agent maximizes U_1^v in period 1 and U_t^r in periods 2, 3 and expects the principal to maximize Π_t^v rather than Π_t^r.

Discussion of Assumptions We assume that effort is verifiable, which does not affect our main results as long as payments are not restricted to being non-negative. The reason is that, if transfers are allowed to have negative values, it is possible to induce first-best effort even in the presence of moral hazard (for example, when effort e_t determines the probability of achieving an output with value θ), without providing the agent with any rent. In such a scenario, the first-period wage would be reduced not only by the expected virtual rent but also by the rent linked to exerting effort in the first period (and similarly for all subsequent periods). Consequently, the incentives for providing effort can be disentangled from exploiting the agent's present bias. Our results would only differ if the agent were protected by limited liability. In that case, the inability to employ negative transfers to reallocate rents would influence the structure of the optimal contract menu (see Englmaier et al., 2023 for an analysis of dynamic moral hazard with limited liability and present-biased agents).

3 Optimal Contract: Time-consistent and Sophisticated Agents

We first derive two benchmarks: profit-maximizing contracts for (a) non-present-biased agents and (b) sophisticated agents.

Time-consistent agent Consider an agent without a present bias ($\beta = 1$). Because the agent's effort is verifiable, the contract

$$e_t = e^{FB}, w_t = c(e^{FB}) + \overline{u}$$

in each period t maximizes both the surplus and the principal's profits. The agent always accepts this contract. In addition, the principal extracts the entire surplus, eliminating any incentive to fire the agent.

Sophisticated present-biased agent A sophisticated present-biased agent correctly anticipates his future choices. Thus, the principal lets C consist of only one element, and the same contract as for a time-consistent agent maximizes surplus and profits (i.e., $e_t = e^{FB}$, $w_t = c(e^{FB}) + \overline{u}$ in every t). The payoffs under such a contract are

$$\begin{aligned}\Pi_1 =& 3\left(e^{FB}\theta - c(e^{FB}) - \overline{u}\right) \\ U_1 =& \bar{u}\left(1 + 2\beta\right).\end{aligned}$$

This contract ensures that the agent accepts the contract in every period, induces him to exert the surplus-maximizing effort level, and allows the principal to extract the entire surplus. Note that adjusting this contract to account for the agent's effectively lower discount factor by front-loading payments to period 1 (in exchange for lower payments in later periods) is not beneficial for the principal. In such a case, the agent—who cannot commit—would quit working for the principal after the first period.

Thus, if the agent is sophisticated, his present bias does not affect the profit-maximizing contract. This result follows from (a) the verifiability of effort and (b) the static production technology that allows the effort and compensation to be realized in the same period. However, with a naïve agent, the principal finds it optimal to create a dynamic compensation structure endogenously, as we show in the next section.

4 Optimal Contract: Naïve Agents

This section analyzes the principal's optimization problem when facing a naïve present-biased agent.

4.1 Optimization Problem

The principal can always offer a naïve agent the same contract as a sophisticated agent. Consequently, $3\left(e^{FB}\theta - c(e^{FB}) - \overline{u}\right)$ sets a lower bound for the principal's profits, who therefore never finds it optimal to fire the agent. In the following, we demonstrate that the principal can further increase her profits. To that end, she can design a dynamic incentive scheme containing a menu of career paths to exploit the naïve agent's misperception of his future behavior. The menu C includes both the virtual path (which seems optimal to the agent from the perspective of period 1) and the real path (which the agent ultimately selects). Next, we derive a series of constraints that this menu C must satisfy.

Individual rationality constraints for the agent The first condition ensures that the agent finds it optimal to accept C in period 1, instead of rejecting C and consuming $\overline{u}$ in all periods. He does so under the expectation of choosing the virtual path in period 2, thus

$$\text{(IRA1)} \qquad w_1^r - c(e_1^r) + \beta\left(w_2^v - c(e_2^v) + w_3^v - c(e_3^v)\right) \geq \bar{u} + 2\beta\bar{u}.$$

Furthermore, in periods 2 and 3, the agent's real and perceived utilities must exceed his outside option:

$$\text{(rIRA2)} \qquad U_2^r \geq \bar{u} + \beta\bar{u}$$

$$\text{(rIRA3)} \qquad U_3^r \geq \bar{u}$$

$$\text{(vIRA2)} \qquad U_2^v \geq 2\bar{u}$$

$$\text{(vIRA3)} \qquad U_3^v \geq \bar{u}.$$

Note that a constraint $U_1^r \geq \bar{u} + 2\beta\bar{u}$ is *not* necessary for an equilibrium because the agent does not expect to choose the real path. In fact, under the profit-maximizing contract, this condition turns out to be violated.

Individual rationality constraints for the principal As previously mentioned, because the principal's profits are always larger than with a sophisticated agent, she will never fire the agent. However, the agent's first-period self must believe the principal will not fire him in the periods $t = 2, 3$ if he has chosen the virtual path:

$$\Pi_2^v \geq -K, \tag{vIRP2}$$

$$\Pi_3^v \geq -K. \tag{vIRP3}$$

If either of these constraints is not satisfied, the agent expects to be laid off in a future period. This feature contrasts with studies such as Eliaz and Spiegler (2006, 2008) and Heidhues and Kőszegi (2010), which assume firms have unlimited commitment power. We deviate from this approach to account for the institutional environment of labor markets, which likely restricts commitments.

Selection constraints As a final condition, the agent must expect to choose the virtual path in the second period but actually select the real path.

$$\begin{aligned} & w_2^r - c(e_2^r) + \beta \left(w_3^r - c(e_3^r) \right) \\ \geq & w_2^v - c(e_2^v) + \beta \left(w_3^v - c(e_3^v) \right), \end{aligned} \tag{rC}$$

$$\begin{aligned} & w_2^v - c(e_2^v) + w_3^v - c(e_3^v) \\ \geq & w_2^r - c(e_2^r) + w_3^r - c(e_3^r). \end{aligned} \tag{vC}$$

Objective The principal's objective is to offer a long-term contract C that maximizes her first-period profits Π_1^r, subject to the constraints just derived.

4.2 Profit-Maximizing Contract

A profit-maximizing contract consists of two key elements. First, the principal shifts as much of the agent's compensation as possible to period 3 of the virtual career path. Second, the principal structures the virtual path in period 2 to be less appealing than the naïve agent expects, since the agent does not anticipate the discounting that occurs between periods 2 and 3.

Thought experiment To demonstrate why such a contract structure allows the principal to exploit the agent, let us introduce a thought experiment. Imagine that the principal offers the naïvely present-biased agent the optimal contract for the sophisticated agent. This contract provides the outside option in every period. Starting from this contract, suppose that we reduce the agent's period-1 payoff by $\Delta_1 > 0$ and increase his period-3 payoff by Δ_1/β. Moreover, we reduce his period-2 payoff by Δ_2 and shift this amount to the third period. From the first period's point of view, decreasing w_1 by Δ_1 and w_2 by Δ_2, and increasing w_3 by $\Delta_1/\beta + \Delta_2$ keeps the agent indifferent to the original situation. That is because

$$\underbrace{-\Delta_1}_{t=1} + \beta \underbrace{\left(\underbrace{-\Delta_2}_{t=2} + \underbrace{\Delta_1/\beta + \Delta_2}_{t=3} \right)}_{=\Delta_1} = 0.$$

However, from the perspective of period 2, the agent's payoff from this operation is $-\Delta_2 + \beta\,(\Delta_1/\beta + \Delta_2) = \Delta_1 - (1-\beta)\Delta_2 < \Delta_1$. Thus, if the principal instead offers an increased payment of $\Delta_1 - (1-\beta)\Delta_2$ paid *in period 2* (in combination with no additional payment in period 3), the agent will accept it. This transaction increases the principal's total profits by $(1-\beta)\Delta_2$ compared to the optimal contract for a sophisticated agent.

This discussion demonstrates that the principal should create a menu of career paths that includes (a) a virtual path that the agent expects to select in the second period and (b) a real path that the agent actually chooses. While the principal shifts the payments of the virtual path to the third period, the real path offers higher second-period and lower third-period payments. By designing this menu, the principal exploits the agent's ignorance of discounting the period-3 payoffs from the perspective of period 2.

The following Proposition (1) details how this contract structure determines the components of a profit-maximizing contract.

Proposition 1. *There exists a profit-maximizing contract with the following features:*

- *All effort levels are e^{FB}.*
- *Wages are*

$$\begin{aligned}
w_1^r &= c(e^{FB}) + \bar{u} - \beta\,(1-\beta)\big(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\big)\\
w_2^r &= w_3^r = c(e^{FB}) + \bar{u}\\
w_2^v &= c(e^{FB}) + \bar{u} - \beta\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right)\\
w_3^v &= K + e^{FB}\theta.
\end{aligned}$$

- *Payoffs are*

$$\begin{aligned}
U_1^r &= (1+2\beta)\,\bar{u} - \beta\,(1-\beta)\big(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\big)\\
U_1^v &= (1+2\beta)\,\bar{u}\\
\Pi_1^r &= 3\big(e^{FB}\theta - c(e^{FB}) - \bar{u}\big) + \beta\,(1-\beta)\big(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\big).
\end{aligned}$$

Proposition 1 demonstrates that the real path involves the same second- and third-period components as the contract for a time-consistent or sophisticated agent. However, the first-period wage is lower: The wage component w_1 encompasses $c(e^{FB}) + \bar{u}$, which corresponds to the agent's "fair" compensation; the term $\beta\,(1-\beta)\big(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\big)$ is subtracted from the fair compensation and indicates the extent of his exploitation. This term reflects the total expected and discounted rent that the agent expects from choosing the virtual path in the future (i.e., from making a career), and it serves as a reward for *today's* effort. It also implies that a decrease in his outside option $\bar{u}$ would harm the agent not only by lowering the utility threshold required for him to accept employment but also by increasing the future relationship surplus. This, in turn, determines the extent of his exploitation.

Another insight of the proposition is that the principal's goal is to maximize the agent's payment in the third period of the virtual path. The reason is that this decision allows her to reduce w_1 by more. However, the third-period wage must be sufficiently low to ensure that it does not seem optimal (from the agent's perspective) for the principal to fire the agent to prevent the "promised" payments. Therefore, w_3^v includes the total output and the firing cost, making the principal indifferent between retaining and firing the agent. The principal crafts the virtual path's second period sufficiently unattractive that the agent actually selects the real path.

Three additional aspects are noteworthy. First, under an optimal contract, all effort levels align with the first-best level. This feature maximizes the effective surplus and enables the principal to set the highest w_3^v to maximally exploit the agent. Only in period 2 of the virtual path, the first-best effort is not uniquely optimal. In this case, the difference $w_2^v - c(e_2^v)$ matters, making the "qualification period" unattractive due to either low wages or high effort. We conclude that the role of effort is negligible in our main model. However, the effort level becomes more relevant in Section 6.4, where we consider an agent who is overconfident about his future effort costs.

Second, as discussed as part of the above thought experiment, the agent's exploitation depends solely on the size of the reduction in the second period (Δ_2). The reason why the optimal contract specified in Proposition 1 then involves a *first-period* wage reduction is that the agent is always free to leave. Thus, under the real path, he must at least receive his outside option in period 2. The wage reduction in period 1, therefore, grants the agent a future rent, which is later reduced due to his time inconsistency.

Third, although the principal earns higher profits when dealing with a naïve agent than with a sophisticated or time-consistent one, these profits do not increase monotonically with the agent's degree of present bias. Instead, there is a non-linear relationship: profits rise with $\beta(1-\beta)$, reflecting how effectively the principal can exploit the naïve agent. This non-linearity arises because exploiting the agent requires deferring his compensation to the future. Although a lower β indicates a stronger present bias—which makes the agent more exploitable—it also means that the agent places less value on future rewards, limiting how much he can actually be exploited. As a result, the principal finds it challenging to exploit not only agents who are nearly time-consistent but also those with extremely high present bias. The reason is that such agents are largely indifferent to delayed compensation.

4.3 Mapping the Model to Real-World Career Paths

The optimal contract offers a *virtual* path with a low-utility qualification phase and a large delayed reward, alongside a *real* path with flatter payments that many agents ultimately select. We now connect these elements to real-world career paths. We first discuss how the model's mechanism applies to a broader range of settings where workers face mid-career burdens. Next, we present two complementary interpretations that link the model to practice. Finally, we illustrate these mappings with examples from healthcare, academia, and accounting.

Broader applicability and generalizations Our model is intentionally stylized, yet its mechanism is broadly applicable and theoretically robust. First, although the baseline model operationalizes the burden of the qualification phase as low wages, the mechanism is more general. What matters is not the specific form of the burden, but that the mid-career stage delivers lower utility than the agent's outside option at that time.[8] Therefore, similar dynamics arise when the mid-career utility is relatively low due to non-financial factors such as limited autonomy, poor work-life balance, or high mid-career effort requirements.[9] Second, although the baseline model predicts that a naïve agent never remains on the steep, high-reward track, the framework can also generate cases where some employees do so once we allow for heterogeneity in time consistency (see extension in Section 6.3). Specifically, if the fraction of time-consistent workers is not too large and the principal cannot distinguish them from naïve ones, it is optimal to offer only the menu designed for naïve agents. The time-consistent workers then complete the virtual path, while the naïve ones deviate mid-career. Overall, the core mechanism captures how present-biased agents initially plan to pursue a challenging career path but later deviate from it when the mid-career utility cost of that path—relative to its future rewards—turns out to be higher than anticipated. Building on this broader perspective, we next outline two interpretations of our model that clarify how the virtual and real paths can be understood in real-world contexts.

Interpretation 1: Dual-track careers with sorting A first interpretation of our model is that multiple career trajectories exist within a single organization. New hires face a high-upside track (the virtual path) with a demanding qualification phase and large, conditional rewards. Alongside this runs a flatter, more stable track with lower long-run upside and less sustained performance pressure. In some firms, employees explicitly choose between

[8]For the following interpretations, it is also important to take into account that our model abstracts from human capital accumulation, which would likely improve outside options over time.

[9]One example with differently defined mid-career costs are settings where firms face frictions that discourage wage reductions, such as downward nominal wage rigidity or psychological considerations (see Bewley, 1999). In these cases, firms may require an increase in effort rather than implementing wage reductions during qualification. The agent's mid-career choice then hinges not on accepting lower wages but on sustaining a high effort level.

tracks (formal dual-career frameworks). In others, sorting occurs implicitly via promotion gates or performance thresholds (without a formal declaration). In these implicit cases, sorting reflects workers' naïve expectations about their future effort and persistence, which the firm anticipates and embeds in the contract.

Interpretation 2: Single-track careers with dropouts In the second interpretation, the agents enter a single, structured career track with distinct stages: an early-career entry phase (period 1), a mid-career qualification stage (period 2), and a potential high-reward phase (period 3). The qualification stage and high-reward phase constitute the virtual contract. By contrast, the real contract is effectively chosen when an employee drops out mid-career and moves to a less demanding role elsewhere. This interpretation aligns with our model because periods 2 and 3 of the real contract are equivalent to the optimal spot contract (i.e., the terms profit-maximizing firms would offer when they do not cater to naïve, present-biased employees). Moreover, with multiple firms present and agents being able to leave their job at no cost, it is natural to treat the spot contract as the agent's effective outside option.[10]

Common mechanism Despite their differences, both interpretations highlight the same economic logic. The principal designs a contract with steep back-loaded utility, knowing that most agents will not ultimately receive the promised reward but still accept a relatively low wage at an early career stage. Whether divergence arises from sorting across multiple career tracks or from deviation from a single track, the result is the same: early-career effort is induced by a reward that rarely has to be paid.

Example 1: Healthcare An example of a career path that aligns with the structure of our model is specialist training in the healthcare sector. Aspiring medical specialists typically enter a single trajectory that resembles the model's virtual path: it is demanding, back-loaded in rewards, and includes a burdensome qualification phase. Medical training begins with several years of schooling, which requires substantial time and financial investment and offers little earnings (period 1). Next comes residency and fellowship, which form the burdensome qualification phase (period 2). During this period, aspiring specialists face long hours, high responsibility, and limited autonomy. They must manage complex clinical cases, lead medical teams, and balance clinical duties with research, teaching, and administration. Despite these demands, compensation remains low, especially compared

[10] In practice, firms may rationally specialize in offering only one of the two paths (e.g, due to specialization in scope or labor) or may provide only a limited number of real-contract positions. In academia, for example, the real path often takes the form of moving to industry or transitioning into a teaching-track position (of which there are typically few). Moreover, note that many professions combine elements of both interpretations. They attract optimistic entrants with steep promised rewards, but feature both career-track sorting and dropout before those rewards are realized.

to peers who choose general practice. As a result, this period is likely the least attractive in terms of overall utility. Only in the final period—upon reaching an attending role—do specialists receive the substantial financial rewards and professional status associated with their career (period 3). Given this structure, our model predicts that (present-biased) early-career doctors start with optimistic plans to become specialists but often abandon this path when mid-career burdens become salient. Empirical evidence supports this interpretation. Goldacre et al. (2010), for example, use a longitudinal study to show that UK medical students overestimate their propensity to become a specialist. They find that while many early career physicians express interest in hospital-based specialties (e.g., surgery, cardiology), a significant proportion switch to general practice during the qualification phase. Concerns about work-life balance, training flexibility, or disillusionment with competitive specialties frequently explain this decision. Lambert and Goldacre (2011) similarly show that the share of graduates preferring general practice over specialties increases significantly between the first and third year after qualification.

Example 2: Academia A similar example in line with the model appears in academic careers. The first stage, Ph.D. training, requires years of hard work and delayed financial stability (period 1). This period is followed by an even more challenging postdoctoral phase or other temporary academic positions, which are highly competitive and often insecure (period 2). At this stage, expectations for publishing, teaching, and grant writing intensify, yet job prospects remain uncertain and salaries are lower than those offered in industry jobs. The rewards of an academic career—tenure, autonomy, and better pay—come only for the few who reach a permanent faculty position (period 3). This career trajectory closely mirrors the structure of our model, and empirical evidence supports the mechanism it implies. For example, Sauermann and Roach (2012) and Roach and Sauermann (2017) show evidence supporting the notion that Ph.D. students in science and engineering start with optimistic plans to become faculty. More than half of them begin their programs aiming for academic careers, but they revise their expectations as they advance their qualification.[11] Crucially, this shift was not driven by a worsening job market, but by a reassessment of their own willingness to endure the academic path (Roach and Sauermann, 2017).[12] Ultimately, many drop out (Denton et al., 2022). Longitudinal studies find that only around 10–20% of STEM Ph.D. graduates ever secure a tenure-track faculty position (e.g. 10% in biology and 20% in physics and chemistry). In short, aspiring academics start with optimistic plans but adjust them before the promised reward is reached. Attrition is not formally orchestrated through a multi-track system. Rather, it emerges from individuals abandoning the original

[11] For example, in chemistry Ph.D. programs the share of students finding a faculty career “extremely attractive” drops from 45% to 32% during the course of chemistry Ph.D. programs.

[12] Roach and Sauermann (2017) show that, although many PhD students initially expect to pursue academic careers, they later realize they dislike the burdensome aspects of academic work (e.g., grant writing, publishing pressure, and research specialization).

plan mid-course.[13]

Example 3: Accounting A third example is the career system in accounting. Upon entering large audit or advisory firms (e.g., the Big Four), new hires face a high-upside trajectory aimed at eventual equity partnership (virtual path). Alongside, they find alternative paths in industry or lower-level roles (real path). The first career step consists of early-career positions as staff accountants or associates (period 1). Here, employees study for the CPA exam, work long hours during peak seasons, and receive modest pay relative to effort. Those who persist move into the qualification phase (period 2), typically as seniors or managers. In their new role, they supervise teams and manage client relationships. At this point, the workload frequently skyrockets, but the promised partnership rewards are still distant. Finally, a small minority advances to an equity partner (period 3), a role that offers substantial financial rewards, status, and firm ownership. Empirically, although many early-career accountants enter intending to become partners, surveys and industry reports show that most revise this plan midway (Going Concern, 2023). Turnover peaks around 3-5 years, precisely during the challenging qualification phase (National Pipeline Advisory Group, 2023). Many professionals are then drawn to corporate roles that offer greater flexibility and better work-life balance (Deloitte Insights, 2016). This pattern is consistent with present-biased behavior: Individuals initially expect to endure the burdens of the qualification stage in pursuit of long-term rewards, but later abandon the path as those burdens materialize.

5 The Role of Firing Costs

We have established that the principal can exploit naïve agents. This section explores how firing costs K influence these exploitation possibilities and the structure of the optimal contract.

5.1 Comparative Statics of K

Drawing from the earlier discussion and Proposition 1, we introduce Proposition 2.

Proposition 2. *The first-period wage and real utility decrease in the firing cost K, second- and*

[13] Another prediction of our model is that when the likelihood of achieving the final reward is smaller, we expect (a) a lower initial compensation and (b) larger later rewards. In fact, departments with very low tenure probabilities tend to pay lower starting salaries and higher wages to tenured faculty (Ehrenberg et al., 1998).

third-period real wages are independent of K:

$$\frac{dw_1^r}{dK} = \frac{dU_1^r}{dK} = -\beta(1-\beta)$$
$$\frac{dw_2^r}{dK} = \frac{dw_2^r}{dK} = 0.$$

A higher K enables the principal to promise greater future payments while, at the same time, lowering the first-period wage. The logic is straightforward: Rising firing costs enhance the principal's commitment to the working relationship, thereby increasing the (perceived) future relationship surplus, which the agent expects to be paid in the third period. However, the first-period wage the agent accepts decreases in the (perceived) surplus. As a result, higher firing costs lower first-period wages. The observation that a higher future relationship surplus leads to lower initial wages implies that, empirically, young workers entering new jobs are likely to face larger wage cuts. This is because their longer expected careers increase the perceived future surplus from the employment relationship, giving employers more extensive exploitation opportunities.

One important point to highlight is that we assume the agent can, from his own perspective, accurately assess how the principal views the situation. In particular, the agent can reject contracts in which the principal's profit along the virtual path falls below $-K$. While this assumption seems strong, we suggest that it should be seen more as a heuristic than as a strict requirement. Even if the agent cannot fully solve the principal's optimization problem, it is enough for him to have a general sense of how his compensation is linked to his productivity for the principal. With this understanding, he can deduce that only certain combinations of effort and promised rewards are credible, depending on the surrounding institutional context. He also recognizes that if the gap between his compensation and his productivity becomes too large, the risk of being dismissed increases. Consequently, the agent's willingness to accept a lower wage early in his career is greater when firing costs are higher.

5.2 Mapping the Model to Real-World Career Systems

We now return to the examples discussed in Section 4.3 to exemplify how the comparative statics of firing costs link the theory to real settings (healthcare, academia, accounting). In the model, K represents the cost an employer faces when terminating a worker. Such costs can arise not only from formal employment protection policies but also from institutional arrangements that make separation costly (such as tenure systems, long-term contracts, or reputational concerns). Viewed in this broader sense, we argue that firing costs in these professions are particularly high at senior hierarchical levels and, in some cases, exceed those implied by general employment protection laws. This pattern may help explain the prevalence of steep, back-loaded career paths in such settings.

Example 1: Healthcare In U.S. health care, clinicians often face stronger *de facto* employment protection than workers under general at-will rules. Under the Health Care Quality Improvement Act, hospitals must follow formal peer-review procedures (with notice and hearing) before disciplining or removing physicians. Moreover, adverse privilege actions lasting more than 30 days must be reported to the National Practitioner Data Bank, creating significant legal and reputational stakes beyond ordinary employment settings (Gharagozloo et al., 2022).

Example 2: Academia Similarly, universities rely on tenure, which in many countries offers much stronger job security than statutory protections. Because violating tenure damages an institution's reputation, universities strive to uphold it even in financial stress. This promise of long-term security is likely a key reason why many PhD students and postdocs accept compensation well below their market value.

Example 3: Accounting A similar logic applies to large accounting firms. For example, removing an equity partner in a Big Four firm requires a formal expulsion process defined in the partnership agreement. This process is typically for cause or by a supermajority vote, and always involves a buyout of their equity. Removing a partner is, hence, much more legally constrained than employment termination, reflecting ownership rights and reputational concerns.

Taken together, these examples illustrate that occupations with high firing costs at upper hierarchical levels often exhibit practices that appear exploitative earlier on and produce the career pattern predicted by Proposition 2: lower entry wages, steeper wage profiles, and credible long-term rewards. Our argument is that elevated firing costs bolster the credibility of deferred rewards, which in turn enables the exploitation of naïve, present-biased employees.

6 Model Variations

Building on our initial model, this section explores several variations. We extend the model by exploring the role of labor market competition and bargaining power, examining the consequences of partial naïveté, and allowing for asymmetric information about the agent's future time preferences. Finally, we demonstrate that firing costs K play a similar role with an agent who—instead of underestimating his future present bias—underestimates his future effort costs.

6.1 Labor Market Competition and Bargaining

We have previously assumed that the principal has full bargaining power and, thus, can determine the terms of the employment relationship. This section discusses a scenario where the agent also has some bargaining power. We operationalize bargaining as follows: Instead of explicitly modeling the bargaining process, we assume that the players arrive at a Nash bargaining outcome in period 1. Here, the principal retains the share α of the total relationship surplus, and the agent gets the share $1-\alpha$. More specifically, the agent accepts any offer that leaves him with $1-\alpha$ of his "present-biased view" of the total relationship surplus. Importantly, in periods $t = 2,3$, the relationship surplus includes the principal's firing costs K. Note that the original contract can also specify that, *in later periods*, a party gets more (or less) than their initial share of the surplus. There, it is important that the agent can still leave without costs. Thus, the contract must at least pay the outside option in any future period.[14]

This setup dictates that the agent's first-period utility U_1^r, which accounts for the fact that he anticipates choosing the virtual path in period 2, must satisfy the following condition:

$$\begin{aligned} U_1^r \geq & \overline{u} + (1-\alpha)(e_1\theta - c(e_1) - \overline{u}) \\ & + \beta\left[\overline{u} + (1-\alpha)\left(e_2^v\theta - c(e_2^v) - \overline{u} + K\right)\right. \\ & \left. + \left(\overline{u} + (1-\alpha)\left(e_3^v\theta - c(e_3^v) - \overline{u} + K\right)\right)\right]. \end{aligned}$$

The rest of the analysis resembles that in Section 4; in particular, all other constraints are identical. Consequently, the principal still offers a menu in period 2 that shifts a major part of the compensation to the third period of the virtual path. It also remains optimal (a) to promise the agent the entire third-period surplus, (b) to reduce w_1^r accordingly, and (c) to set all effort levels to the first best. Therefore, Proposition 3 emerges.

Proposition 3. *Assume that in period 1, the agent can secure a share* $(1-\alpha)$ *of the total relationship surplus from his perspective. Then, we have*

$$\begin{aligned} w_1^r = & \overline{u} + c(e^{FB}) - \beta(1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \\ & + (1-\alpha)\left(e^{FB}\theta - c(e^{FB}) - \bar{u}\right) \\ & + 2\beta(1-\alpha)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right), \end{aligned}$$

and

$$\frac{\partial^2 w_1^r}{\partial \alpha \partial K} = -2\beta.$$

[14]This feature is different from Miller and Watson (2013) and Fahn (2017). In these articles, the inability of parties to commit not to renegotiate any agreement undermines the efficiency of long-term employment relationships.

The first line of w_1^r in Proposition 3 contains the wage when the agent lacks bargaining power. By contrast, the second line represents the agent's share of the first-period surplus (which does not include K). Lastly, the third line reflects his share of the second- and third-period surplus. The proposition indicates that the adverse effects of higher firing costs on the wages of young employees are more pronounced for agents with lower bargaining power. This result arises because that smaller values of α increase the utility the agent is bound to receive anyway in future periods. Then, the principal is more constrained to *further* increase the agent's virtual rent in period 3 (compared to the main case, where she has full bargaining power). Consequently, she is also less able to decrease w_1^r.

6.2 Partial Naïveté

While our main model assumes that the agent is fully naïve, we now discuss two versions of partial naïveté that have been applied in the literature. On the one hand, the agent perceives his future present bias parameter to be $\hat{\beta} \in (\beta, 1]$ (O'Donoghue and Rabin, 2001). On the other hand, we assume "frequency naïveté" as defined by Eliaz and Spiegler (2006). In this scenario, the agent believes there is a probability q that he will remain time-consistent in the future, and a probability $1-q$ that he will be present-biased; the principal is aware of the value of q. While our results are unchanged in the first case, they depend on the value of q in the second. In particular, as q increases, the effect of K becomes more significant, as demonstrated by the following proposition.

Proposition 4. *Assume that the agent is partially naïve. If he perceives his future present bias parameter to be $\hat{\beta} \in (\beta, 1]$, the results are equivalent to him being fully naïve. If he expects to be time consistent in the future periods with probability q and present biased with probability $1-q$, where $q \in [0, 1]$, the outcomes are as follows:*

$$
\begin{aligned}
w_1^r &= c(e^{FB}) + \bar{u} - q\beta(1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \\
U_1^r &= (1+2\beta)\bar{u} - q\beta(1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \\
U_1^v &= (1+2\beta)\bar{u} \\
\Pi_1^r &= 3\left(e^{FB}\theta - c(e^{FB}) - \bar{u}\right) + q\beta(1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right),
\end{aligned}
$$

with

$$
\begin{aligned}
\frac{d^2 w_1^r}{dqdK} &< 0 \\
\frac{d^2 \tilde{U}_1^r}{dqdK} &< 0 \\
\frac{d^2 \Pi_1^r}{dqdK} &> 0.
\end{aligned}
$$

The intuition behind the first result is that $\hat{\beta}$ affects only two constraints: (1) the agent's individual rationality constraint in the second period when considering the virtual career path from the perspective of the first period, and (2) the (vC) constraint, which ensures that the agent's first-period self finds it optimal to select the virtual path in the second period. In all other constraints, only the true value of β is relevant. However, these two constraints were not binding in the original problem and continue to be satisfied for the wage and effort allocations derived in Section 4.2, even when $\hat{\beta}$ varies. Thus, for a given β, the contract is independent of the agent's degree of naïveté, unless he is fully sophisticated (i.e., when $\hat{\beta} = \beta$). It is worth noting that this type of discontinuity is a common characteristic in the literature (Heidhues and Kőszegi, 2010).

The situation changes in the case of frequency naïveté. Here, the agent expects to select the virtual contract with probability q. With probability $1-q$, he anticipates choosing the real contract, which specifies wages $w_2^r = w_3^r = c(e^{FB}) + \bar{u}$. Importantly, the structure of the virtual contract does not depend on q. As a consequence, the agent's degree of naïveté mainly affects the weight assigned to the virtual contract and, therefore, the importance of the firing costs K.

6.3 Asymmetric Information

Now, we explore how the presence of time-consistent agents affects our mechanism. Suppose that, with probability $1-\gamma$, an agent is time-consistent, while with probability γ, the agent is naïve present-biased. The principal cannot observe the agent's type. We will show that there is a threshold (independent of K) that determines the principal's strategy: either to offer a menu tailored for naïve agents, accepting that time-consistent ("TC") agents might select the virtual contract, or to avoid exploiting naïve agents entirely. Then, she offers only the optimal contract for time-consistent agents. Furthermore, we will point out that introducing sophisticated present-biased agents does not alter this trade-off.

Proposition 5. *Assume an agent is time-consistent with probability $1-\gamma$, and naïve present-biased with probability γ, and the principal cannot observe the agent's type. Then, a profit-maximizing equilibrium exists with the following properties. If*

$$\gamma \geq 1-\beta,$$

the principal exclusively offers the optimal exploitation contract designed for the naïve agent, as detailed in Proposition 1. In this case, the time-consistent agent ultimately chooses the virtual *contract in period 2. The resulting profits and utilities are*

$$\begin{aligned}\Pi =& 3\left(e^{FB}\theta - c(e^{FB}) - \bar{u}\right) \\ & + (\beta - (1-\gamma))(1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \\ U_1^{rN} =& (1+2\beta)\bar{u} - \beta(1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \\ U_1^{TC} =& 3\bar{u} + (1-\beta)^2\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right),\end{aligned}$$

where the superscript "N" indicates the naïve agent, and "TC" the time-consistent agent.

If $\gamma < 1-\beta$, the principal only offers the profit-maximizing contract tailored for the time-consistent agent. Then,

$$\begin{aligned}\Pi =& 3\left(e^{FB}\theta - c(e^{FB}) - \bar{u}\right) \\ U_1^{rN} =& (1+2\beta)\bar{u} \\ U_1^{TC} =& 3\bar{u}.\end{aligned}$$

The intuition is as follows. To exploit the naïve agent, the principal must offer a virtual contract that shifts the promised rents to period 3. However, these future rents are even more appealing to the TC agent, who does not discount future payoffs. Moreover, since his perceived and actual future time preferences align, he will also select the virtual contract in period 2. Therefore, the TC agent receives an additional rent under *any* exploitation contract. There is no less costly way for the principal to provide this rent to the TC agent than simply offering the optimal exploitation contract intended for the naïve agent and accepting that the TC agent will go for the virtual contract. Because the rent transferred from the naïve agent to the TC agent under an exploitation contract affects the principal's profit linearly, the principal's optimal strategy is to either offer the strongest possible exploitation contract (with the largest feasible rent shifted to period 3 in the virtual contract) or to forgo exploitation contracts entirely. This result is similar to that of Yu (2021), who finds that the presence of TC agents limits the extent to which the mechanism designer can exploit the naïveté of present-biased agents in an optimal taxation framework, because TC agents would actually choose virtual ("imaginary") allocations.

Finally, the (additional) presence of sophisticated agents does not pose a problem if, alongside the menu of contracts, the static optimal contract is also offered—assuming the tie-breaking condition that the naïve agent chooses the menu if indifferent. Under this setup, the sophisticated agent will select the static optimal contract. This contract provides him with a strictly larger utility than the menu, whereas naïve and TC agents opt for the menu.

6.4 Overconfident Agent

The agent in our model mispredicts his future behavior because he overestimates his discount factor in future periods. This misperception can also be interpreted as a form of overconfidence regarding his future characteristics. In this section, we show that firing costs K play a similar role in a situation where an agent underestimates future effort costs. We demonstrate that the principal still exploits the agent by offering him a virtual contract that he expects to select in future periods, while he actually chooses the real contract. However, the structure of the virtual contract might differ. Now, the virtual contract does not entail an unattractive "qualification" period that the agent must first pass before receiving a rent thereafter. Instead, the first period of the virtual contract for an overconfident agent contains both high effort and the rent he is offered in exchange for low payments in the real contract. Therefore, we would predict different (perceived) career paths for overconfident agents compared to those with a present bias: While the latter expects a rent in the career phase *after* passing an unattractive qualification period, the former perceives to be rewarded immediately.

To formally derive these results, we assume that the agent's effort costs still amount to $c(e)$, but that he perceives them to be $\eta c(e)$ in the second period, with $\eta < 1$. Besides, he now discounts the future exponentially. Moreover, he correctly perceives that his third-period effort costs are $c(e)$. We impose these assumptions to maintain comparability with our main model. While (at least) three periods are needed to construct an exploitation contract for a naïve present-biased agent, two periods are sufficient for an agent who is overconfident about his effort costs. If the agent also were overconfident regarding third-period effort costs, it would be possible for the principal to exploit the agent repeatedly. Moreover, we would need to specify whether and how the agent learns and updates his beliefs over time (see Fahn and Klein, 2023, for an analysis of the dynamics of exploitation contracts if players update their beliefs using Bayes' rule).

Most importantly, in the second period, the principal still offers a menu of contracts that contains a real and a virtual contract, as in our main model. Generally, the set of constraints is the same as with a naïve present-biased agent and is described in the proof to Proposition 6.

Proposition 6. *Assume the agent is overconfident regarding his future effort costs, as described above. Then, there exists a profit-maximizing menu of contracts that has the following properties.*

1. *The virtual contract contains high effort in the second period, $e_2^v > e^{FB}$, with $de_2^v/dK > 0$. In all other (real and virtual) periods, effort is e^{FB}.*

(a) *Wages are*

$$
\begin{aligned}
w_1^r &= \bar{u} + c(e^{FB}) - (1-\eta)\, c(e_2^v) \\
w_2^r = w_3^r = w_3^v &= \bar{u} + c(e^{FB}) \\
w_2^v &= \bar{u} + c(e_2^v),
\end{aligned}
$$

with

$$
\frac{dw_1^r}{dK} < 0 \text{ and } \frac{dw_2^v}{dK} > 0.
$$

(b) *Realized payoffs are*

$$
\begin{aligned}
U_1^r &= 3\bar{u} - (1-\eta)\, c(e_2^v) \\
\Pi_1^r &= 3\left(e^{FB}\theta - c(e^{FB}) - \bar{u}\right) + (1-\eta)\, c(e_2^v),
\end{aligned}
$$

with

$$
\frac{dU_1^r}{dK} < 0 \text{ and } \frac{d\Pi_1^r}{dK} > 0.
$$

There is one main difference between the model with a present-biased agent and the model with an overconfident agent. In the latter, the agent's compensation for a reduced e_1^r is concentrated in period 2 in the virtual contract; he is left with no rent in period 3. Instead, a naïve present-biased agent is offered this rent in period 3 in the virtual contract. In addition to that, the implications of the two models are the same. Therefore, what matters for our results is the agent's naïveté about his future characteristics, not their exact nature. Most importantly, whenever the agent is naïve, (a) the principal can exploit the agent and (b) the extent of exploitation increases in firing costs, which determines her credibility when promising future rents.

7 Predictions and Evidence

The model gives rise to several clear predictions about career dynamics. Subsequently, we outline key patterns that should emerge if the model captures features of real-world careers and discuss evidence that is in line with some of them.

Back-loaded contracts The model implies that when agents are naïve about their present bias, principals can exploit this misperception by deferring compensation to later periods:

Prediction 1a: Firms will design back-loaded compensation profiles (with lower entry wages) when workers are naïve present-biased.[15]

> ***Prediction 1b:*** The degree of back-loading (with lower entry wages) increases (a) with the intensity of present bias at the individual level (as long as $\beta > 0.5$) and (b) in groups of workers where present bias is more prevalent.

Causally testing these predictions requires (a) identifying variation in present bias and (b) datasets that contain information on both compensation trajectories and workers' present bias. Ideally, researchers would exploit exogenous variation in present bias—arising from natural experiments, shocks, or randomized debiasing interventions—to identify its causal effect on contract design. However, in the absence of such variation, one can still gather suggestive evidence by comparing wage-tenure profiles across sectors, occupations, or firms that differ in the prevalence of present-biased preferences. Beyond identification, a central challenge is measurement. A possible (though imperfect) approach is to proxy present bias using observable traits linked to time-inconsistent behavior. For example, present bias seems to be more prevalent among individuals who are younger, have a lower level of education or cognitive reflection, have lower savings rates, use payday loans more frequently, or carry persistent credit card debt (e.g., Frederick, 2005; Meier and Sprenger, 2010; Sutter et al., 2013; Carvalho et al., 2016). Firms may offer these groups more back-loaded wage profiles and lower entry wages.[16] Alternatively, one can measure present bias directly at the individual level using survey or experimental elicitation methods.[17]

Dropout behavior Another core behavioral implication of the model is that present-biased agents initially pledge to back-loaded career paths, but they abandon them later. They enter contracts expecting to follow through to the delayed reward, but revise their plans once mid-career burdens become salient. In addition, in environments with more present-biased individuals, firms design steeper contracts and attrition rates rise accordingly (Section 6.3). This mechanism gives rise to three distinct but related empirical predictions:

> ***Prediction 2a:*** Steeply back-loaded compensation schemes will lead to high attrition rates in the qualification phase (before workers receive the promised reward).

[15] Deferred compensation is a well-established incentive scheme (Lazear, 1979; Lazear and Rosen, 1981). Our contribution is to show that such contracts can also emerge as a response to time-inconsistent preferences.

[16] Several studies are in line with this idea: For example, using firm-level personnel data, Pfeifer (2013) shows that, within firms, younger hires receive lower starting wages and experience higher wage growth compared to older hires. Similarly, van Huizen and Alessie (2015) document that on-the-job search (for promotions) and work effort increase with patience. Firms may exploit this tendency by offering back-loaded contracts that naïve workers are unlikely to realize.

[17] Such measures are available in survey datasets such as the German Socio-Economic Panel (SOEP, 2006 special module) or the Understanding Society Innovation Panel (UKHLS-IP, Wave 6). Pinger (2017) shows that the survey measure of present bias is a strong predictor of real-life behavior.

> ***Prediction 2b:*** The attrition rates will increase (a) with the intensity of present bias at the individual level (as long as $\beta > 0.5$) and (b) in groups of workers where present bias is more prevalent.
>
> ***Prediction 2c:*** Naïve present-biased workers will exhibit a systematic gap between intended and actual behavior: they enter with plans to pursue long-term rewards but later abandon these plans when burdens become salient.

These predictions, again, lend themselves to both descriptive and causal testing strategies at individual and aggregate levels. Researchers could, in principle, assess *Prediction 2a* by examining career trajectories in professions with delayed payoffs, such as the above-mentioned healthcare, academia, or accounting. One possibility involves estimating survival rates and identifying dropout or switching patterns, particularly during the demanding mid-career phase. For example, in up-or-out systems, one could document the share of an initial cohort that (a) achieves promotion, (b) exits the firm, or (c) switches to less demanding tracks, especially in the qualification phase.[18] An alternative approach is to use quasi-experimental variation in compensation structure (e.g., changes to promotion systems, partnership rules, or pay schedules) to test whether steeper back-loading increases attrition. Instead, researchers could test *Prediction 2b* by comparing attrition rates across sectors, firms, or cohorts that vary in the prevalence of present-biased behavior.[19] Lastly, to assess *Prediction 2c*, researchers may use longitudinal data that track workers' early-career expectations (e.g., plans to become a partner, secure tenure, remain in a high-intensity track, or stay in the firm) and compare them to their realized career outcomes. For example, one could test whether early-career workers systematically overestimate the likelihood of persisting along the intended career path (such as remaining with their employer through qualification or attaining promotion). Suggestive evidence from academia or healthcare points to such intention-outcome mismatches (see Section 4.3).

Firing costs The model predicts that higher firing costs improve the firm's ability to credibly promise future rewards because layoffs become more costly. This stronger commitment power allows firms to offer more steeply back-loaded contracts and reduce entry wages. Moreover, the model predicts heterogeneous effects: mostly new employment relationships should respond to changes in firing costs. Incumbent workers who are already beyond the exploitative entry phase should remain unaffected.

[18] There is already some evidence in line with this idea. Existing studies, for example, find that only around 10–20% of Ph.D. students secure tenure, and the likelihood of an associate making partner at a large law firm ranges between 15% and 20% (e.g., Ehrenberg et al., 1998; BCG, 2023).

[19] Carrell and Zinman (2014) present evidence suggesting that individuals who are more likely to be present-biased are also more likely to exit long-term reward tracks prematurely. They exploit the random assignment of U.S. Air Force personnel to bases in different U.S. states. Some states allow access to payday loans, while others restrict them. They find that, when payday loans are available, job performance deteriorates and early attrition increases. In line with present bias, the effects are stronger for younger and financially unsophisticated airmen.

Prediction 3a: When firing costs increase, the entry wages will fall and the wage profiles become more back-loaded. When firing costs fall, the entry wages will rise and the profiles will flatten.

Prediction 3b: These effects will be concentrated among newly hired workers. The wages of incumbent workers should remain largely unchanged.

These predictions are best tested in quasi-experimental settings and, indeed, such studies testing *Prediction 3a and 3b* already exist. For example, Leonardi and Pica (2013) analyze the impacts of a 1990 labor market reform in Italy. The reform raised firing costs for smaller firms (up to 15 employees) but not for larger ones (more than 15 employees). Analyzing administrative data with a regression discontinuity difference-in-difference design, the authors document that increased firing costs slightly lower average wages. The effects are highly heterogeneous. New hires during the reform period suffer a drop in the entry wage, while incumbent workers are left unaffected. In line with the notion that younger people face a stronger present bias, the reduction is more substantial for young workers below 30.[20] Reversely, Cervini-Pla et al. (2014) study the effects of a 1997 Spanish labor market reform that reduced dismissal costs and non-wage labor costs for certain worker groups. Using a difference-in-differences approach with administrative data, their main finding suggests that decreased firing costs have a positive effect on the wages of newly hired workers.

Bargaining power One of our model extensions shows that if the agent has more bargaining power, the back-loading of compensation becomes less extreme. We find:

Prediction 4a: Firms will offer more steeply back-loaded contracts (with lower entry wages) to workers or groups who have weaker bargaining power.

Prediction 4b: The effect of firing costs on contract back-loading will be weaker for workers or groups with stronger bargaining power (e.g., due to unions or legal renegotiation rights).

Ideally, researchers would use exogenous variation in workers' bargaining power to examine its impact on entry-level wages and pay growth. Potential sources of variation in bargaining include changes in collective representation, labor law, or union coverage. For example, Budde et al. (2024) use worker council elections and retirement-induced turnover in

[20] Past papers try to explain such results with "conventional" models of labor market frictions and decentralized bargaining. Leonardi and Pica (2013), for example, argue that incumbents gain bargaining power when firing becomes more expensive, leading to higher wages. Instead, new hires "pre-pay" for the added job security and accept lower wages. However, Cervini-Pla et al. (2014) and Leonardi and Pica (2013) observe mainly wage reductions for "job switchers," while the "incumbents" wages remain mostly unaffected. This observation contradicts incumbent-focused logic. Instead, it aligns with our model, where a higher K does not impact existing relationships. Leonardi and Pica (2013) attribute the lack of incumbent wage gains to firms' refusal to renegotiate. However, if workers anticipate this, they would not initially accept a wage cut.

Germany to show that increases in blue-collar representation affect employment outcomes without reducing wages. These results suggest that stronger worker-side representation constrains firms' ability to shift compensation into later periods. Future studies can apply similar or related quasi-experimental strategies to examine how bargaining power (on its own or in interaction with firing costs) shapes entry wages and the steepness of wage profiles. Descriptive strategies can complement this approach. One way is to compare wage trajectories across groups that differ in bargaining strength, particularly before and after changes in firing costs. For example, Leonardi and Pica (2013) find that the detrimental effect of higher firing costs on wages is greater for workers with low bargaining power. Examples include young blue-collar workers or workers with earnings just above the sectoral contractual minimum compensation.

Labor market competition Similarly to the effect of bargaining power, labor market competition also constrains the firm's ability to extract surplus through back-loaded contracts. When workers have credible outside options, firms must offer better initial terms or risk losing them. As a result, we expect less exploitative contract structures in competitive labor markets:

> ***Prediction 5a:*** In more competitive labor markets, entry wages will be higher and wage profiles less steeply back-loaded.
>
> ***Prediction 5b:*** The gap between promised (virtual) and realized (actual) rewards will be smaller when workers have more attractive outside options.

Researchers can test these predictions by comparing industries or occupations that differ in monopsony power or turnover costs. For example, in competitive tech sectors, firms face pressure to offer better initial terms. At the same time, in more monopsonistic environments (e.g., rural manufacturing, specialized academia, or where non-compete agreements are more common), deferred compensation structures should be more pronounced. Empirically, one could exploit shocks to labor market competition (such as the entry of new firms, changes in non-compete enforceability, or expansions in union rights). The goal would be to examine whether these shifts lead to higher entry wages, flatter wage profiles, or fewer delayed rewards.

8 Discussion and Conclusion

The key message of this paper is that employment protection laws can have unintended consequences: they allow firms to exploit naïve present-biased employees. The mechanism arises from the structure of dynamic exploitation contracts, involving a virtual career path (which the agent expects to choose in the future) and a real path (which he ends up selecting). Higher firing costs increase a firm's commitment when promising a future (virtual) compensation, allowing for a larger wage reduction early on.

A natural question is whether this logic extends to alternative forms of firing costs. The literature distinguishes between firing costs that are transfers to workers (e.g., severance pay) and those that are non-recoverable costs to firms (such as firing taxes). Lazear (1990) argues that severance payments may be passed on to workers through lower wages or performance bonds, thereby leaving total labor costs unchanged. In contrast, most empirical work treats firing costs as a tax on job destruction (Bertola and Rogerson, 1997; Betcherman, 2013). This literature then notes that wage rigidities and financial constraints often prevent firms from offsetting these costs (Martin and Scarpetta, 2012). Leonardi and Pica (2013), for example, show that firing costs can depress wages for new hires, even when they fail to improve bargaining outcomes for workers. Our model helps explain this empirical pattern: such wage suppression can arise even if firing costs take the form of a tax, provided that firms use dynamic exploitation contracts to extract surplus from naïve present-biased employees.

This raises the question of how our mechanism would respond if firing costs were structured as severance payments rather than taxes. To clarify this, first, suppose the agent also receives a severance payment if he chooses to leave voluntarily in periods $t = 2$ or $t = 3$. In this case, severance effectively increases his outside option.

Second, suppose instead that the principal paid severance payments only after firing the agent (and the courts were able to verify that). Then, severance payments would leave our results unchanged. In this case, only the cost component—not the amount captured by the agent—affects the virtual path. In period $t = 3$, the principal optimally promises a high rent, which is determined solely by her termination costs. In period 2, the qualification period, effort requirements can be adjusted to make it sufficiently unattractive, regardless of how much the agent would be paid if he was laid off. Nevertheless, severance payments affect the real path,[21] which is costly for the principal but not anticipated by the agent. Therefore, the agent is only willing to accept an early-career wage reduction for the costly component to the principal, not for the higher pay he could extract in later stages of his career. Taken together, these considerations suggest that our core mechanism may remain robust to institutional variation in the form of firing costs.

[21] In periods $t = 2, 3$, the agent's real compensation is determined by his effective outside option

References

Acharya, V. V., Baghai, R. P., and Subramanian, K. V. (2013). Wrongful discharge laws and innovation. *Review of Financial Studies*, 27(1):301–346. 1

BCG (2023). What percentage of associates at big law firms make partner? https://www.bcgsearch.com/videos/11976/What-percentage-of-associates-at-big-law-firms-make-partner. Accessed: 2024-06-03. 18

Belot, M., Boone, J., and von Ours, J. (2007). Welfare-improving employment protection. *Economica*, 74(295):381–396. 1, 1

Bertola, G. and Rogerson, R. (1997). Institutions and labor reallocation. *European Economic Review*, 41(6):1147–1171. 1, 5, 8

Betcherman, G. (2013). Labor market institutions: A review of the literature. *World Bank Policy Research Working Paper*, (6276). 1, 8

Bewley, T. F. (1999). *Why Wages Don't Fall during a Recession*. Harvard University Press. 9

Budde, J., Dohmen, T., Jäger, S., and Trenkle, S. (2024). Worker representatives. Working Paper 32740, NBER. 7

Carrell, S. E. and Zinman, J. (2014). In harm's way? Payday loan access and military personnel performance. *Review of Financial Studies*, 27(9):2805–2840. 1, 19

Carvalho, L. S., Meier, S., and Wang, S. W. (2016). Poverty and economic decision-making: Evidence from changes in financial resources at payday. *American Economic Review*, 106(2):260–284. 7

Cervini-Pla, M., Ramos, X., and Silva, J. I. (2014). Wage effects of non-wage labour costs. *European Economic Review*, 72:113–137. 1, 7, 20

CFO (2023). 71% of Big 4 auditors worry about mental health. https://www.cfo.com/news/71-of-big-4-auditors-worry-about-mental-health/712063. Accessed: 2024-06-03. 1

Chan, M. K. (2017). Welfare dependence and self-control: An empirical analysis. *The Review of Economic Studies*, 84(4):1379–1423. 1

Cheung, S. L., Tymula, A., and Wang, X. (2021). Quasi-hyperbolic present bias: A meta-analysis. *Working Paper*, (2021-15). 1

Crémer, J. and McLean, R. P. (1988). Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247. 1

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2):315–372. 1

DellaVigna, S. and Malmendier, U. (2004). Contract design and self-control: Theory and evidence. *Quarterly Journal of Economics*, 119(2):353–402. 1

Deloitte Insights (2016). Global human capital trends 2016: The new organization—different by design. Deloitte University Press. 1, 4.3

Denton, M., Borrego, M., and Knight, D. B. (2022). US postdoctoral careers in life sciences, physical sciences and engineering: Government, industry, and academia. *PLoS One*, 17(2):e0263185. 1, 4.3

Ehrenberg, R. G., Pieper, P. J., and Willis, R. A. (1998). Do economics departments with lower tenure probabilities pay higher faculty salaries? *Review of Economics and Statistics*, 80(4):503–512. 1, 13, 18

Eliaz, K. and Spiegler, R. (2006). Contracting with diversely naive agents. *The Review of Economic Studies*, 73(3):689–714. 1, 1, 2, 2, 4.1, 6.2

Eliaz, K. and Spiegler, R. (2008). Consumer optimism and price discrimination. *Theoretical Economics*, 3(4):459–497. 4.1

Englmaier, F., Fahn, M., and Schwarz, M. A. (2023). Long-term employment relations when agents are present biased. *Working Paper*. 1, 2

Ericson, K. M. (2014). Consumer inertia and firm pricing in the Medicare Part D prescription drug insurance exchange. *American Economic Journal: Economic Policy*, 6(1):38–64. 1, 1

Ericson, K. M. (2017). On the interaction of memory and procrastination: Implications for reminders, deadlines, and empirical estimation. *Journal of the European Economic Association*, 15(3):692–719. 1

Ericson, K. M. (2020). When consumers do not make an active decision: Dynamic default rules and their equilibrium effects. *Games and Economic Behavior*, 124:369–385. 1

Estevez-Abe, M., Iversen, T., and Soskice, D. (2001). *Social protection and the formation of skills: A reinterpretation of the welfare state*, pages 145–183. Oxford University PressOxford. 1

Fahn, M. (2017). Minimum wages and relational contracts. *Journal of Law, Economics, and Organization*, 33(2):301–331. 14

Fahn, M. and Klein, N. (2023). Non-common priors, incentives, and promotions: The role of learning. *Working Paper*. 6.4

Fahn, M. and Seibel, R. (2022). Present bias in the labor market – when it pays to be naive. *Games and Economic Behavior*, 135:144–167. 1

Fang, H. and Wang, Y. (2015). Estimating dynamic discrete choice models with hyperbolic discounting, with an application to mammography decisions. *International Economic Review*, 56(2):565–596. 1

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4):25–42. 7

Gharagozloo, F., Poston, R., and Gruessner, R. (2022). The health care quality improvement act of 1986: What every surgeon needs to know. *World Journal of Surgery and Surgical Research*. 5.2

Gilpatric, S. M. (2008). Present-biased preferences, self-awareness and shirking. *Journal of Economic Behavior & Organization*, 67(3):735–754. 1

Going Concern (2023). How bad are the odds of making partner at a Big 4 Firm? https://www.goingconcern.com/how-bad-are-the-odds-of-making-partner-at-a-big-4-firm. Accessed: 2024-06-03. 1, 4.3

Goldacre, M. J., Laxton, L., and Lambert, T. W. (2010). Medical graduates' early career choices of specialty and their eventual specialty destinations: UK prospective cohort studies. *BMJ*, 341(jul06 1):c3199–c3199. 1, 4.3

Gottlieb, D. and Zhang, X. (2021). Long-term contracting with time-inconsistent agents. *Econometrica*, 89(2):793–824. 1

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7):2643–82. 1, 1

Heidhues, P. and Kőszegi, B. (2010). Exploiting naïvete about self-control in the credit market. *American Economic Review*, 100(5):2279–2303. 1, 4.1, 6.2

Heidhues, P. and Kőszegi, B. (2018). *Behavioral industrial organization*, pages 517–612. Elsevier. 1

Hopenhayn, H. and Rogerson, R. (1993). Job turnover and policy evaluation: A general equilibrium analysis. *Journal of Political Economy*, 101(5):915–938. 1, 1

Johnen, J. (2019). Automatic-renewal contracts with heterogeneous consumer inertia. *Journal of Economics & Management Strategy*, 28(4):765–786. 1

Kaur, S., Kremer, M., and Mullainathan, S. (2015). Self-control at work. *Journal of Political Economy*, 123(6):1227–1277. 1, 1

Kőszegi, B. (2014). Behavioral contract theory. *Journal of Economic Literature*, 52(4):1075–1118. 1

Kuzmina, O. (2023). Employment flexibility and capital structure: Evidence from a natural experiment. *Management Science*, 69(9):4992–5017. 1

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112(2):443–477. 1, 2

Lambert, T. and Goldacre, M. (2011). Trends in doctors' early career choices for general practice in the UK: Longitudinal questionnaire surveys. *British Journal of General Practice*, 61(588):e397–e403. 1, 4.3

Lazear, E. P. (1979). Why is there mandatory retirement? *Journal of Political Economy*, 87(6):1261–1284. 15

Lazear, E. P. (1990). Job security provisions and employment. *The Quarterly Journal of Economics*, 105(3):699. 8

Lazear, E. P. and Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5):841–864. 15

Leonardi, M. and Pica, G. (2013). Who pays for it? The heterogeneous wage effects of employment protection legislation. *The Economic Journal*, 123(573):1236–1278. 1, 7, 20, 7, 8

Li, S., Yan, J., and Yu, J. (2012). When Holmström and Milgrom meet Laibson: Moral hazard with short-run impatience and imperfect self-awareness. *Working Paper*. 1

Lockwood, B. B. (2020). Optimal income taxation with present bias. *American Economic Journal: Economic Policy*, 12(4):298–327. 1

Martin, J. P. and Scarpetta, S. (2012). Setting it right: Employment protection, labour reallocation and productivity. *De Economist*, 160(2):89–116. 8

Mas, A. and Pallais, A. (2017). Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–3759. 1, 1

Meier, S. and Sprenger, C. (2010). Present-biased preferences and credit card borrowing. *American Economic Journal: Applied Economics*, 2(1):193–210. 7

Miller, D. A. and Watson, J. (2013). A theory of disagreement in repeated games with bargaining. *Econometrica*, 81(6):2303–2350. 14

Mortensen, D. T. and Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *The Review of Economic Studies*, 61(3):397–415. 1

Murooka, T. and Schwarz, M. A. (2018). The timing of choice-enhancing policies. *Journal of Public Economics*, 157:27–40. 1

Murooka, T. and Schwarz, M. A. (2019). Consumer exploitation and notice periods. *Economics Letters*, 174:89–92. 1

National Pipeline Advisory Group (2023). Accounting staff turnover is highest after 3-5 years, shows NPAG poll. https://www.accountingpipeline.org/accounting-staff-turnover-is-highest-after-3-5-years-shows-npag-poll. Accessed: 2024-06-03. 1, 4.3

O'Donoghue, T. and Rabin, M. (1999). Incentives for procrastinators. *The Quarterly Journal of Economics*, 114(3):769–816. 2

O'Donoghue, T. and Rabin, M. (2001). Choice and procrastination. *Quarterly Journal of Economics*, 116(1):121–160. 2, 6.2

OECD (2013). Protecting jobs, enhancing flexibility: A new look at employment protection legislation. In *OECD Employment Outlook 2013*. 1

Pfeifer, C. (2013). Adjustments of wage-tenure profiles with respect to entry age. *Journal of Business Economics*, 83:171–191. 16

Phelps, E. S. and Pollak, R. A. (1968). On second-best national saving and game-equilibrium growth. *Review of Economic Studies*, 35(2):185–199. 2

Pierre, G. and Scarpetta, S. (2004). How labor market policies can combine workers' protection and job creation: A partial review of some key issues and policy options. *Working Paper*. 1

Pinger, P. R. (2017). Predicting experimental choice behavior and life outcomes from a survey measure of present bias. *Economics Bulletin*, 37(3):2162–2172. 17

Roach, M. and Sauermann, H. (2017). The declining interest in an academic career. *PloS one*, 12(9):e0184130. 1, 4.3, 12

Sauermann, H. and Roach, M. (2012). Science PhD career preferences: Levels, changes, and advisor encouragement. *PLoS One*, 7(5):e36307. 1, 4.3

Sulka, T. (2024). Exploitative contracting in a life cycle savings model. *Review of Economics and Statistics*, pages 1–47. 1, 1

Sutter, M., Kocher, M. G., Ruetzler, D., and Trautmann, S. T. (2013). Experimental decisions predict adolescents' field behavior. *American Economic Review*, 103(1):510–531. 7

van Huizen, T. and Alessie, R. (2015). Time preferences and career investments. *Labour Economics*, 35:77–92. 16

Wasmer, E. (1999). Competition for jobs in a growing economy and the emergence of dualism. *The Economic Journal*, 109(457):349–371. 1

Yılmaz, M. (2013). Repeated moral hazard with a time-inconsistent agent. *Journal of Economic Behavior & Organization*, 95:70–89. 1

Yu, P. C. (2021). Optimal retirement policies with present-biased agents. *Journal of the European Economic Association*, 19(4):2085–2130. 1, 6.3

A Appendix: Omitted Proofs

Proof to Proposition 1 The objective is to maximize

$$\Pi_1 = e_1^r\theta - w_1^r + e_2^r\theta - w_2^r + e_3^r\theta - w_3^r,$$

subject to

$$\begin{aligned}
&\text{(IRA1)} && w_1^r - c(e_1^r) + \beta\left(w_2^v - c(e_2^v) + w_3^v - c(e_3^v)\right) \geq \bar{u} + 2\beta\bar{u} \\
&\text{(rIRA2)} && w_2^r - c(e_2^r) + \beta\left(w_3^r - c(e_3^r)\right) \geq \bar{u} + \beta\bar{u} \\
&\text{(rIRA3)} && w_3^r - c(e_3^r) \geq \bar{u} \\
&\text{(vIRA2)} && w_2^v - c(e_2^v) + w_3^v - c(e_3^v) \geq 2\bar{u} \\
&\text{(vIRA3)} && w_3^v - c(e_3^v) \geq \bar{u} \\
&\text{(vIRP2)} && w_2^v - e_2^v\theta + w_3^v - e_3^v\theta \leq K \\
&\text{(vIRP3)} && w_3^v - e_3^v\theta \leq K \\
&\text{(rC)} && w_2^r - c(e_2^r) + \beta\left(w_3^r - c(e_3^r)\right) \geq w_2^v - c(e_2^v) + \beta\left(w_3^v - c(e_3^v)\right) \\
&\text{(vC)} && w_2^v - c(e_2^v) + w_3^v - c(e_3^v) \geq w_2^r - c(e_2^r) + w_3^r - c(e_3^r)
\end{aligned}$$

For the following, we omit the constraints (vIRA2), (vIRA3), (vIRP2) and (vC), and check ex post whether they hold for the derived contract.

First, (IRA1) binds. If it did not bind, we could reduce w_1^r without violating any constraint. This yields

$$w_1^r = c(e_1^r) + \bar{u} + \beta\left(c(e_2^v) + \bar{u} - w_2^v + c(e_3^v) + \bar{u} - w_3^v\right),$$

and the "new" optimization problem that maximizes

$$\begin{aligned}
\Pi_1^r =& e_1^r\theta - c(e_1^r) - \bar{u} - \beta\left(c(e_2^v) + \bar{u} - w_2^v + c(e_3^v) + \bar{u} - w_3^v\right) \\
&+ e_2^r\theta - w_2^r + e_3^r\theta - w_3^r,
\end{aligned}$$

subject to

$$\begin{aligned}
&\text{(rIRA2)} && w_2^r - c(e_2^r) + \beta\left(w_3^r - c(e_3^r)\right) \geq \bar{u} + \beta\bar{u} \\
&\text{(rIRA3)} && w_3^r - c(e_3^r) \geq \bar{u} \\
&\text{(vIRP3)} && w_3^v \leq K + e_3^v\theta \\
&\text{(rC)} && w_2^r - c(e_2^r) + \beta\left(w_3^r - c(e_3^r)\right) \geq w_2^v - c(e_2^v) + \beta\left(w_3^v - c(e_3^v)\right).
\end{aligned}$$

Second, we show that the constraints (rIRA2), (rIRA3), and (rC) bind. To the contrary, assume that (rC) is slack. Then, the principal can increase w_2^v, which increases profits but does not violate any constraint. Furthermore, if (rIRA2) is slack, the principal can reduce both, w_2^r and w_2^v, by a small ε. Thereby, (rC) remains satisfied, whereas profits increase by $(1-\beta)\,\varepsilon$. Finally if (rIRA3) is slack, the principal can reduce reduce w_3^r and w_3^v by a small ε. Thereby, (rC) remains satisfied, whereas profits increase by $(1-\beta)\,\varepsilon$.

Using these results yields

$$\begin{aligned} w_2^r &= c(e_2^r) + \bar{u} \\ w_3^r &= c(e_3^r) + \bar{u} \\ w_2^v &= c(e_2^v) + \bar{u} - \beta\left(w_3^v - \bar{u} - c(e_3^v)\right), \end{aligned}$$

and the optimization problem maximizes

$$\begin{aligned} \Pi_1^r =& e_1^r\theta - c(e_1^r) - \bar{u} + e_2^r\theta - \bar{u} - c(e_2^r) + e_3^r\theta - c(e_3^r) - \bar{u} \\ & - \beta\left[\beta\left(e_3^v\theta - c(e_3^v) - \bar{u} + K\right) + c(e_3^v) + \bar{u} - w_3^v\right] \end{aligned}$$

subject to

$$\text{(vIRP3)} \qquad w_3^v \leq K + e_3^v\theta.$$

Since Π_1^r increases in w_3^v, (vIRP3) binds as well, and profits are

$$\begin{aligned} \Pi_1^r =& e_1^r\theta - c(e_1^r) - \bar{u} + e_2^r\theta - \bar{u} - c(e_2^r) + e_3^r\theta - c(e_3^r) - \bar{u} \\ & + \beta\,(1-\beta)\left(e_3^v\theta - c(e_3^v) - \bar{u} + K\right). \end{aligned}$$

It immediately follows that $e_1^r = e_2^r = e_3^r = e_3^v = e^{FB}$. Moreover, since e_2^v only enters $w_2^v = c(e_2^v) + \bar{u} - \beta\left(e_3^v\theta - c(e_3^v) - \bar{u} + K\right)$, it is without loss to set $e_2^v = e^{FB}$ as well.

Taking these results into account yields

$$\begin{aligned} w_1^r &= c(e^{FB}) + \bar{u} - \beta\,(1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \\ w_2^r &= c(e^{FB}) + \bar{u} \\ w_3^r &= c(e^{FB}) + \bar{u} \\ w_2^v &= c(e^{FB}) + \bar{u} - \beta\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \\ w_3^v &= K + e^{FB}\theta \end{aligned}$$

Finally, we have to confirm that these outcomes satisfy the omitted constraints we omit the constraints (vIRA2), (vIRA3), (vIRP2) and (vC). These conditions become

$$
\begin{aligned}
&\text{(vIRA2)} && (1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \geq 0 \\
&\text{(vIRA3)} && e^{FB}\theta - c(e^{FB}) - \bar{u} + K \geq 0 \\
&\text{(vIRP2)} && \left(e^{FB}\theta - c(e^{FB}) - \bar{u}\right) + \beta\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \geq 0 \\
&\text{(vC)} && (1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \geq 0,
\end{aligned}
$$

and clearly hold.

Finally, plugging effort and wages into the payoff functions yields

$$
\begin{aligned}
U_1^r &= (1+2\beta)\bar{u} - \beta(1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \\
\Pi_1^r &= 3\left(e^{FB}\theta - c(e^{FB}) - \bar{u}\right) + \beta(1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right).
\end{aligned}
$$

■

Proof to Proposition 4 First, assume the agent perceives his future present bias parameter to be $\hat{\beta} \in (\beta, 1]$. Then, in the optimization problem formulated in the proof to Proposition 1, $\hat{\beta}$ only affects two constraints: the (vIRA2) constraint and the (vC) constraint. These constraints become

$$
\begin{aligned}
&\text{(vIRA2)} && w_2^v - c(e_2^v) + \hat{\beta}\left(w_3^v - c(e_3^v)\right) \geq \bar{u} + \hat{\beta}\bar{u} \\
&\text{(vC)} && w_2^v - c(e_2^v) + \hat{\beta}\left(w_3^v - c(e_3^v)\right) \geq w_2^r - c(e_2^r) + \hat{\beta}\left(w_3^r - c(e_3^r)\right).
\end{aligned}
$$

Substituting the optimal values obtained in Proposition 1 into these constraints yields

$$
\begin{aligned}
&\text{(vIRA2)} && \left(\hat{\beta} - \beta\right)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \geq 0 \\
&\text{(vC)} && \left(\hat{\beta} - \beta\right)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \geq 0,
\end{aligned}
$$

which holds for any $\hat{\beta} \in (\beta, 1]$.

Now, assume frequency naïveté, i.e., that the agent believes he will remain time-consistent in the future with probability q and present–biased with probability $1-q$.

Then, the objective is to maximize

$$
\Pi_1 = e_1^r\theta - w_1^r + e_2^r\theta - w_2^r + e_3^r\theta - w_3^r,
$$

subject to

$$
\begin{aligned}
&\text{(IRA1)} && w_1^r - c(e_1^r) + q\beta\left(w_2^v - c(e_2^v) + w_3^v - c(e_3^v)\right) \\
& && + (1-q)\beta\left(w_2^r - c(e_2^r) + w_3^r - c(e_3^r)\right) \geq \bar{u} + 2\beta\bar{u} \\
&\text{(rIRA2)} && w_2^r - c(e_2^r) + \beta\left(w_3^r - c(e_3^r)\right) \geq \bar{u} + \beta\bar{u} \\
&\text{(rIRA3)} && w_3^r - c(e_3^r) \geq \bar{u} \\
&\text{(vIRA2)} && w_2^v - c(e_2^v) + w_3^v - c(e_3^v) \geq 2\bar{u} \\
&\text{(vIRA3)} && w_3^v - c(e_3^v) \geq \bar{u} \\
&\text{(vIRP2)} && w_2^v - e_2^v\theta + w_3^v - e_3^v\theta \leq K \\
&\text{(vIRP3)} && w_3^v - e_3^v\theta \leq K \\
&\text{(rC)} && w_2^r - c(e_2^r) + \beta\left(w_3^r - c(e_3^r)\right) \geq w_2^v - c(e_2^v) + \beta\left(w_3^v - c(e_3^v)\right) \\
&\text{(vC)} && w_2^v - c(e_2^v) + w_3^v - c(e_3^v) \geq w_2^r - c(e_2^r) + w_3^r - c(e_3^r)
\end{aligned}
$$

As in the proof to Proposition 1, we omit the constraints (vIRA2), (vIRA3), (vIRP2) and (vC), and then verify afterwards whether they are satisfied by the resulting contract. Additionally, by applying similar reasoning as in the proof of Proposition 1, we can demonstrate that the constraints (rIRA), (rIRA2), (rIRA3), (rC), and (vIRP3) are binding. This allows us to replace the payments in the principal's profits, which then become

$$
\begin{aligned}
\Pi_1 = & e_1^r\theta - c(e_1^r) - \bar{u} + e_2^r\theta - c(e_2^r) - \bar{u} + e_3^r\theta - c(e_3^r) - \bar{u} \\
& + q\beta(1-\beta)\left(e_3^v\theta - c(e_3^v) - \bar{u} + K\right)
\end{aligned}
$$

It immediately follows that $e_1^r = e_2^r = e_3^r = e_3^v = e^{FB}$. Moreover, since e_2^v only enters $w_2^v = c(e_2^v) + \bar{u} - \beta\left(e_3^v\theta - c(e_3^v) - \bar{u} + K\right)$, it is without loss to set $e_2^v = e^{FB}$ as well.
Taking these results into account yields

$$
\begin{aligned}
& w_1^r = c(e^{FB}) + \bar{u} - q\beta(1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \\
& w_2^r = c(e^{FB}) + \bar{u} \\
& w_3^r = c(e^{FB}) + \bar{u} \\
& w_2^v = \bar{u} + c(e^{FB}) - \beta\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right) \\
& w_3^v = K + e^{FB}\theta
\end{aligned}
$$

It follows that these outcomes satisfy the omitted constraints we omit the constraints (vIRA2), (vIRA3), (vIRP2) and (vC).

Finally, plugging effort and wages into the payoff functions yields

$$U_1^r = (1+2\beta)\bar{u} - q\beta(1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right)$$
$$\Pi_1^r = 3\left(e^{FB}\theta - c(e^{FB}) - \bar{u}\right) + q\beta(1-\beta)\left(e^{FB}\theta - c(e^{FB}) - \bar{u} + K\right),$$

from which the stated comparative statics follow immediately.

■

Proof to Proposition 5 We show that the principal cannot earn higher profits than by implementing either the optimal static contract or the optimal exploitation contract identified in Subsection 4.2. In other words, it is not possible to profitably screen agent types through contract selection in period 1; instead, screening occurs in period 2, when the TC agent chooses the virtual contract and the naïve agent chooses the real contract.

First, observe that *any* exploitation contract works by shifting promised rents into period 3, thereby taking advantage of the naïve agent's misperception regarding the relative weights of future-period utilities. Moreover, since the naïve agent discounts future payoffs by $\beta < 1$, any reduction in period-1 compensation must be offset by larger increases in later periods to keep the agent indifferent.

Second, consider the case where a TC agent is present. Then, if we hold the naïve agent's (perceived) period-1 utility at his outside-option level, any exploitation contract provides the TC agent with strictly more utility than his outside option. Therefore, when faced with a choice between a profit-maximizing static contract (which is a profit-maximizing contract among the set of non-exploitation contracts) and any exploitation contract that provides the naïve agent with at least his outside option, the TC agent will *always select the latter*—and actually go for the virtual contract in period 2.

Third, we profitable screening types is not feasible because the only difference between the two agent types lies in their discount factor; their effort costs and outputs are otherwise identical. As a result, it is impossible to profitably separate the types through effort choices. The principal must, at minimum, pay the TC agent the additional rent required to induce acceptance of the exploitation contract designed for the naïve agent. To show that no payment structure exists that profitably screens the two types suppose the principal increases her profit by Δ when contracting with the naïve agent. Relative to the static benchmark (i.e., wages $= c(e^{FB}) + \bar{u}$), it follows from Proposition 1 that the wage changes under the exploitation contract are as follows:

$$\Delta w_1^r = -\Delta, \qquad \Delta w_2^v = -\frac{\Delta}{1-\beta}, \qquad \Delta w_3^v = \frac{\Delta}{\beta(1-\beta)}.$$

These changes increase the TC agent's utility by $(1-\beta)/\beta \cdot \Delta$ and decrease the principal's profits by the same amount when dealing with the TC agent. It remains to show that the specific timing of how the principal provides this extra rent to the TC agent does not impact overall profits. Offering Δ in period 2 could achieve screening through contract selection, but it would not raise the principal's profits. Providing Δ in period 1 would prompt the naïve agent to also choose the contract meant for the TC agent, thereby reducing the principal's profits. Thus, delivering Δ in the third period of the virtual contract is weakly optimal. Accordingly, the principal's best achievable outcomes are either the optimal static contract or the optimal exploitation contract, as characterized in Subsection 4.2.

Finally, we derive the conditions under which each choice is optimal. Recall that if the principal increases her profits with a naïve agent by an amount Δ, her profits when dealing with the TC agent decrease by $(1-\beta)/\beta \cdot \Delta$. The agent is time consistent with probability $1-\gamma$, and naïvely present-biased with probability γ. Therefore, the principal's profits for an arbitrary Δ are

$$\begin{aligned}\Pi =&\gamma\Pi_1^N + (1-\gamma)\Pi_1^{TC}\\ =&3\left(e^{FB}\theta - c(e^{FB}) - \bar{u}\right) + \gamma\Delta - (1-\gamma)\frac{1-\beta}{\beta}\Delta.\end{aligned}$$

This implies that the principal's profits are maximized by choosing Δ as large as possible, if

$$\begin{aligned}&\gamma\Delta - (1-\gamma)\frac{1-\beta}{\beta}\Delta \geq 0\\ \Leftrightarrow&\gamma \geq 1-\beta.\end{aligned}$$

Otherwise, the optimal choice is $\Delta = 0$.

■

Proof to Proposition 6 The objective is to maximize

$$\Pi_1 = e_1^r\theta - w_1^r + e_2^r\theta - w_2^r + e_3^r\theta - w_3^r,$$

subject to

$$
\begin{aligned}
&\text{(IRA1)} && w_1^r - c(e_1^r) + \left(w_2^v - \eta c(e_2^v) + w_3^v - c(e_3^v)\right) \geq 3\bar{u} \\
&\text{(rIRA2)} && w_2^r - c(e_2^r) + w_3^r - c(e_3^r) \geq 2\bar{u} \\
&\text{(rIRA3)} && w_3^r - c(e_3^r) \geq \bar{u} \\
&\text{(vIRA2)} && w_2^v - \eta c(e_2^v) + w_3^v - c(e_3^v) \geq 2\bar{u} \\
&\text{(vIRA3)} && w_3^v - c(e_3^v) \geq \bar{u} \\
&\text{(vIRP2)} && w_2^v - e_2^v\theta + w_3^v - e_3^v\theta \leq K \\
&\text{(vIRP3)} && w_3^v - e_3^v\theta \leq K \\
&\text{(rC)} && w_2^r - c(e_2^r) + w_3^r - c(e_3^r) \geq w_2^v - c(e_2^v) + w_3^v - c(e_3^v) \\
&\text{(vC)} && w_2^v - \eta c(e_2^v) + w_3^v - c(e_3^v) \geq w_2^r - \eta c(e_2^r) + w_3^r - c(e_3^r)
\end{aligned}
$$

For the following, we omit the constraints (vIRA2), (vIRA3), and (vC), and check ex post whether they hold for the derived contract.

First, (rIRA) binds. If it did not bind, the principal could reduce w_1^r without violating any constraint. This yields

$$w_1^r = c(e_1^r) + 3\bar{u} + \eta c(e_2^v) - w_2^v + c(e_3^v) - w_3^v,$$

and the "new" optimization problem that maximizes

$$
\begin{aligned}
\Pi_1^r =& e_1^r\theta + e_2^r\theta + e_3^r\theta - c(e_1^r) - \eta c(e_2^v) - c(e_3^v) - 3\bar{u} \\
&+ w_2^v - w_2^r + w_3^v - w_3^r,
\end{aligned}
$$

subject to

$$
\begin{aligned}
&\text{(rIRA2)} && w_2^r - c(e_2^r) + w_3^r - c(e_3^r) \geq 2\bar{u} \\
&\text{(rIRA3)} && w_3^r - c(e_3^r) \geq \bar{u} \\
&\text{(vIRP2)} && w_2^v - e_2^v\theta + w_3^v - e_3^v\theta \leq K \\
&\text{(vIRP3)} && w_3^v \leq K + e_3^v\theta \\
&\text{(rC)} && w_2^r - c(e_2^r) + w_3^r - c(e_3^r) \geq w_2^v - c(e_2^v) + w_3^v - c(e_3^v).
\end{aligned}
$$

Second, we show that the constraints (rIRA2), (rIRA3), and (rC) bind. To the contrary, assume that (rC) is slack. Then, the principal can increase w_2^v, which increases profits but does not violate any constraint. If (rIRA3) is slack, the principal can reduce both, w_3^r and w_3^v, by a small ε. Thereby, (rC) remains satisfied and profits are unaffected, whereas (vIRP2) and (vIRP3) are relaxed. Furthermore, if (rIRA2) is slack, the principal can reduce both, w_2^r and w_2^v, by a small ε. Thereby, (rC) remains satisfied and profits are unaffected, whereas (vIRP2) is relaxed.

Using these results yields

$$\begin{aligned} w_2^r &= \bar{u} + c(e_2^r) \\ w_3^r &= \bar{u} + c(e_3^r) \\ w_2^v &= 2\bar{u} + c(e_2^v) + c(e_3^v) - w_3^v \end{aligned}$$

and the optimization problem maximizes

$$\begin{aligned} \Pi_1^r =& e_1^r\theta + e_2^r\theta + e_3^r\theta - c(e_1^r) - c(e_2^r) - c(e_3^r) \\ &+ (1-\eta)\,c(e_2^v) - 3\bar{u}, \end{aligned}$$

subject to

(vIRP2) $$2\bar{u} - \left[e_2^v\theta - c(e_2^v) + e_3^v\theta - c(e_3^v)\right] \leq K$$

(vIRP3) $$w_3^v \leq K + e_3^v\theta.$$

Π_1^r increases in e_2^v, thus it should be as high as possible until (vIRP2) binds.Then,

$$\frac{de_2^v}{dK} = \frac{1}{c'(e_2^v) - \theta} > 0.$$

Moreover, $e_3^v = e^{FB}$ is clearly optimal because it minimizes the left hand side of (vIRP2), therefore $e_1^r = e_1^r = e_1^r = e_3^v = e^{FB}$.

Finally, we have to confirm that these outcomes satisfy the omitted constraints (vIRA2), (vIRA3), (vIRP2) and (vC). These conditions become

(vIRA2) $$(1-\eta)\,c(e_2^v) \geq 0$$

(vIRA3) $$w_3^v \geq \bar{u} + c(e^{FB})$$

(vC) $$c(e_2^v) \geq c(e^{FB})$$

and clearly hold, in particular setting $w_3^v = \bar{u} + c(e^{FB})$ satisfies (vIRP3) and (vIRA3).

Then, wages and payoffs are

$$\begin{aligned} w_1^r &= \bar{u} + c(e^{FB}) - (1-\eta)\,c(e_2^v) \\ w_2^r = w_3^r = w_3^v &= \bar{u} + c(e^{FB}) \\ w_2^v &= \bar{u} + c(e_2^v) \end{aligned}$$

$$\begin{aligned} U_1^r &= 3\bar{u} - (1-\eta)\,c(e_2^v) \\ \Pi_1^r &= 3\left(e^{FB}\theta - c(e^{FB}) - \bar{u}\right) + (1-\eta)\,c(e_2^v). \end{aligned}$$

■