Long-Term Employment Relations When Agents are Present Biased*

Florian Englmaier† Matthias Fahn‡ Marco A. Schwarz§

March 11, 2020

Abstract

We analyze how agents' present bias affects optimal contracting in a dynamic employment setting. The principal maximizes profits by offering a menu of contracts to naive agents: a *virtual* contract – which agents plan to choose in the future – and a *real* contract which they end up choosing. This virtual contract motivates the agent and allows the principal to keep the agent below his outside option. We show that such a profit-maximizing menu of contracts has important implications on how labor market institutions such as employment protection affect profits, wages, and effort.

JEL Codes: D03, D21, J31, M52

Keywords: employment relations, dynamic contracting, present bias

*We are grateful to Dan Barron, Paul Heidhues, Christian Holzner, Botond Kőszegi, Jin Li, Dylan Minor, Takeshi Murooka, David Schindler, Klaus M. Schmidt, Peter Schwardmann, Ran Spiegler, and to seminar audiences in Munich, at IZA, at the 2014 SFB/TR-15 conference in Berlin, the COPE 2015 in Vienna, at the ISNIE 2015 at Harvard Law School, at the World Congress of the Econometric Society 2015 in Montreal, at the Annual Congress of the European Economic Association 2015 in Mannheim, at the 2015 CESifo conference on Behavioral Economics in Munich, and at the LEOH 2018 in Hamburg for valuable comments and suggestions. Anne Niedermaier and Paula von Gablenz provided very helpful research assistance. Financial support by Deutsche Forschungsgemeinschaft through SFB/TR–15 and CRC/TRR–190 and the Austrian Science Fund (SFB F63) is gratefully acknowledged.

†LMU Munich, Department of Economics & Organizations Research Group & CEPR & CESifo; florian.englmaier@econ.lmu.de

‡JKU Linz & CESifo; matthias.fahn@jku.at

§University of Innsbruck, Department of Economics; marco.schwarz@uibk.ac.at

1 Introduction

Many individuals are present biased and put extra weight on current over future consumption (see DellaVigna, 2009, for an excellent survey). Firms have developed numerous ways to exploit consumers, in particular if those are naive and thus fail to correctly anticipate the future extent of their present bias; examples include domains such as retirement savings (Laibson, 1997), health club attendance (DellaVigna and Malmendier, 2006), or credit card usage (Shui and Ausubel, 2004). But individuals are employees as well, and also their career choices involve important intertemporal trade-offs. Still, firms' responses to their employees' present bias have so far attracted much less attention.

Nonetheless, there is evidence that present-biased preferences also matter in labor markets. Kaur, Kremer, and Mullainathan (2010, 2015) use a field experiment with data entry workers to demonstrate that many of them have present biased preferences; Frakes and Wasserman (2016) provide evidence for U.S. patent examiners. Moreover, employees suffer substantial wage losses if they smoke (Böckerman, Hyytinen, and Kaprio, 2015), drink heavily (Boeckerman, Hyytinen, and Maczulskij, 2017), or are obese (Brunello and D'Hombres, 2007, Baum and Ford, 2004). These traits have been found to correlate with present-biased preferences and naiveté about it,[1] which provides indirect evidence for the latter's negative effect on labor market outcomes.

In this paper, we explore how a long-term employment contract can be devised to exploit naive present-biased employees. Firms set up a career path that *endogenously* creates a cost/reward structure with immediate costs but delayed benefits. Employees who are naive about their present bias are offered a menu of contracts, consisting of a *virtual* contract and a *real* contract. The virtual contract contains a major part of the employee's promised compensation, but also an unattractive "qualification period" that has to be taken before enjoying these benefits. Thus, employees who are naive about their present bias plan to select the virtual contract in the future, yet end up choosing the real contract.

Our benchmark setting involves a risk neutral principal and a risk neutral agent who interact over an infinite time horizon. Whereas the principal discounts future profits exponentially, the agent discounts his future utility in a quasi-hyperbolic way. The principal offers a long-term contract to the agent and is able to commit to its terms as long as

[1] See Kang and Ikeda (2016a) for smoking; Kang and Ikeda (2016b) for smoking and obesity, and Ikeda, Kang, and Ohtake (2010) for overweight.

her profits do not fall short of her outside option. If the agent works for the principal, he chooses costly effort. Effort generates a deterministic output that is consumed by the principal, after which the agent receives this period's compensation. Since the output of a period's effort is generated in the very same period, and since effort and output are verifiable, the agent's present bias per se does not represent an obstacle for efficiency. A spot contract could implement efficient outcomes, allowing the principal to extract the full surplus. Indeed, such a spot contract is offered to agents without a present bias and those who are present-biased but fully aware of their bias.

A naive agent who is not aware of his future present bias, however, is offered a menu of contracts. It consists of a virtual contract which the agent intends to choose in the future and a real contract which he ends up choosing. The first period of the virtual contract is designed as a "qualification period" in which the agent's utility is low. Thereafter, the virtual contract grants the agent a high utility level. From today's perspective, which involves an extra weight on current utility, it is optimal to select the virtual contract in the subsequent period. Then, the benefits of the later periods of the virtual contract outweigh the costs of the qualification period. This makes the agent willing to accept a lower compensation today, i.e., in the real contract. However, when tomorrow comes, the lower utility of the qualification period is not worth the future benefits anymore and the agent postpones taking up the virtual contract. Consequently, the agent always ends up choosing the real contract, leaving him with a utility below his outside option.

Thus, we derive a mechanism that can deliver an explanation for the lower compensation of naive present-biased employees. It is grounded in their inability to overcome short-run barriers established by their employer before letting them get onto the career trajectory. As a consequence, the firm can provide incentives at substantially lower costs. Of course, most employment relationships do not have detailed contracts that explicitly describe future compensation. However, there often is an implicit understanding that, if an employee exerts a lot of effort or takes up certain career development options, he will be promoted or rewarded in another way.[2] After deriving the general design of a profit-maximizing long-term contract for a naive present-biased agent, we explore a number of implications that follow from our main mechanism.

[2] An area in which those agreements are a more explicit is multilevel marketing (as used by e.g., Herba Life), where there are clear rules what level of sales and what number of sub-contractors will result in what level compensation.)

Analyzing a version of our model with a finite horizon, we show in Section 6.1 that (young) agents who are at the beginning of their working lives will be exploited more. The reason is that they can be more easily lured by career prospects: They work weakly more (strictly so if there is a binding limited liability constraint; see below) and are paid less. This broadly resembles actual labor market patterns where we find age earning profiles with relatively better paid but less productive older workers.[3]

Generally, how much the principal can exploit the agent depends on how auspicious the virtual contract appears. The agent is promised the total surplus of the virtual contract, and this (discounted) value is subtracted from the agent's real payments. Thus, anything that makes this surplus look more attractive – such as a higher standard discount factor or lower outside options – harms a naive present-biased agent. Because one might regard the principal's outside option to also include the costs of firing the agent, this can give rise to the additional interpretation that less stringent employment protection increases a naive present-biased agent's real compensation. Indeed, Cervini-Pl, Ramos, and Silva (2014) find that a Spanish reform reducing firing costs led to higher wages for affected workers. Furthermore, Leonardi and Pica (2013) analyze the consequences of a reform that increased firing costs in Italy and observe a slight average reduction of wages. This effect is mainly driven by entry costs of workers that moved jobs, which is consistent with our view that – due to the principal's commitment to long-term contracts – wage reductions should mostly be observed in new matches. To explain their results, Leonardi and Pica (2013) focus on models of labor market frictions and decentralized bargaining. There, higher firing costs increase incumbent workers' bargaining position and thus should allow them to raise their wages. New workers, on the other hand, "pre-pay" for the increased job security and accept lower wages. However, Leonardi and Pica (2013) only observe wage reductions for "outsiders", whereas "insiders" are left unaffected, which is at odds with their theoretical explanations, but is in line with our results.[4]

[3]See Waldman (2012), for an overview of empirical evidence for upward-sloping wage curves, as well as alternative theoretical explanations. Moreover, there is evidence indicating that older workers are less productive. For example, Haltiwanger, Lane, and Spletzer (1999) show that firms with higher shares of workers above 55 are less productive; similar results are obtained by Skirbekk (2004) and Lallemand and Rycx (2009), who find that firms with high proportions of older workers are less productive, in particular in relation to the wages they receive. Note that some of these studies find that productivity initially increases with age. Such a result could be generated by our model if we also allowed for human capital accumulation. Then, an increased productivity driven by a higher human capital would dominate early on, whereas the contractual forces which reduce effort would matter more at later stages of an employment relationship.

[4]Leonardi and Pica (2013) aim at explaining this discrepancy with the absence of a credible threat

Moreover, the negative effect of higher firing costs on wages in Leonardi and Pica (2013) is significantly more pronounced for workers below 30. Such an outcome would be generated by the aforemetioned version of our model with a finite time horizon, but would be more difficult to explain by the alternative models mentioned by Leonardi and Pica (2013). Finally, Leonardi and Pica (2013) find that the negative effect of firing costs on wages is particularly strong for workers with low bargaining power. This result is consistent with an extension of our model where we take different values of bargaining power into account (Section 6.2). There, we allow the agent to capture a varying share of the relationship surplus (in our benchmark model, the principal enjoys full bargaining power). We show that a higher bargaining power of the agent reduces the principal's ability to take advantage of the agent's present bias. Then, the negative effect of a higher future (virtual) relationship surplus – for example caused by higher firing costs – on the agent's compensation gradually diminishes. Note that this extension can also be applied to pursue the effect of labor market competition. If we assume a more competitive labor market to increase the agent's bargaining power, the extent of competition does not change the qualitative structure of the optimal menu of contracts. However, more competition for workers gives less scope to exploit the agent, and a higher virtual surplus will eventually even benefit him.

Moreover, we analyze further variations of the model that deliver additional insights on how dealing with present-biased employees might affect labor market outcomes.

If the agent is protected by limited liability (Section 6.3), the principal is constrained in reducing the agent's real compensation. Then, once limited liability is binding and the agent receives a zero payment in the real contract, the principal resorts to extracting additional rent by inducing the agent to work harder than the efficient effort level (and harder than an agent without a present bias).[5] Such an inefficiency is more pronounced for (ceteris paribus) more productive employment relationships. Furthermore, having limited liability in a finite time horizon can give rise to an effort profile that gradually goes down over time. We also show that our results are robust to incorporating moral hazard, partial naiveté, learning, a weaker present bias in the monetary domain, and unobservable agent

by workers in case firms refuse to renegotiate wages. However, if workers anticipated their inability to renegotiate higher wages later on, they should not be willing to accept upfront wage cuts.

[5]This result relates to a number of papers that have demonstrated that (naive) present-biased agents might actually perform better than time-consistent agents (Fahn and Hakenes, 2019, and Weinschenk, 2019, show that this can hold in team settings; Fahn and Seibel, 2020, show that naive agents might conduct more on-the-job search). Different from these papers, the higher effort in our setting destroys surplus.

types.

Our results deliver a number of testable implications beyond the ones already stated. In Section 8, we discuss ways to potentially assess their empirical validity.

Finally, note that being naive about one's future time preferences could also be perceived as a form of overoptimism: Principal and agent disagree about a future state (the agent's time preferences), and the compensation structure is designed in a way to exploit this disagreement. Indeed, Gervais, Heaton, and Odean (2011) and de la Rosa (2011) analyze principal-agent settings with moral hazard and find that the principal can incentivize the agent more cheaply if he overestimates the likelihood to succeed.[6] To work out the precise role of the agent's present bias, and to potentially separate the channels by which overoptimism and present bias allow the principal to exploit the agent, we analyze a version of our model in which the agent is an exponential discounter but underestimates his future effort costs (Section 7). Then, the principal still offers a menu of contracts with a virtual and a real contract. However, while it is uniquely optimal to offer an unattractive first virtual period followed by high payoffs *thereafter* to a naive present-biased agent, the first virtual period for an overoptimistic agent might not only include higher effort, but also the agent's complete virtual rent (with the principal keeping the full surplus of all subsequent virtual periods; such an arrangement is strictly optimal if the principal has a larger discount factor than the agent). Thus, we argue that the pattern of an unattractive qualification period being followed by a promising career characterizes the exploitative contract for a present-biased agent, whereas an overoptimistic agent would rather be offered a contract with virtual costs and benefits being condensed into one period.

2 Related Literature

The paper relates to the literature on time-inconsistent preferences, first formalized by Strotz (1955) who allows for an individual's discount rate between two periods to depend on the time of evaluation. Phelps and Pollak (1968) argue that, since time-inconsistent preferences affect savings, the possibility of individuals having a present bias should be in-

[6] Evidence is provided by Hoffman and Burks (2015), who analyze data from a trucking firm with detailed information on drivers' beliefs about their future performance. Many drivers overestimate their future productivity and adjust their beliefs only slowly. The firm benefits from drivers' overoptimism as those with a larger bias are less likely to quit.

corporated into growth models. Laibson (1997) shows that given people are present-biased, choices that seem suboptimal – for example the purchase of illiquid assets – can actually increase an individual's utility by binding future selves and hence providing commitment. He develops the workhorse model to analyze present-biased time preferences, the $\beta - \delta$-model: An individual gives extra weight to utility today over any future period, but discounts all future instances in a standard exponential way. O'Donoghue and Rabin (1999a) compare the welfare of so-called "sophisticated" and "naive" individuals, where the former are aware of their present bias and the latter are not.[7]

Besides numerous theoretical contributions, there is substantial evidence suggesting that people make decisions that are not consistent over time. For example, consider Shui and Ausubel (2005) or DellaVigna and Malmendier (2004, 2006), who document present-biased behavior for credit card usage and health club attendance respectively. Kaur, Kremer, and Mullainathan (2010, 2015) provide evidence from a field experiment with data entry workers that present bias at work is important and Frakes and Wasserman (2016) provide real-world evidence from the U.S. patent office. Augenblick, Niederle, and Sprenger (2015) document that subjects show a considerable present bias in effort (while they only find limited present bias in monetary choices). This suggests that studying the role of present bias in the workplace and in workers' careers is a particularly relevant and promising topic for research.

We also relate to the literature on optimal contracting choices when agents are present biased. O'Donoghue and Rabin (1999b) develop a model where a principal hires an agent with present-biased preferences in order to complete a task, but the agent's costs of doing so vary. If the latter is the agent's private information, it can be optimal to employ a scheme of increasing punishments if the task has not been completed yet. Whereas the interaction in O'Donoghue and Rabin (1999b) is basically one-shot, i.e., the relationship between principal and agent is terminated once the task has been completed, we show how repeated interaction can have a substantial impact on optimal contracts, by allowing the agent to choose among a menu of contracts (with history-dependent elements) in every period.

Similar to O'Donoghue and Rabin (1999b), Gilpatric (2008) analyzes a contracting prob-

[7]Further theoretical contributions like Harris and Laibson (2001) or Krusell and Smith (2003) show that (quasi-)hyperbolic discounting of individuals can have important implications for consumption-savings decisions.

lem between a principal and a (risk-neutral) agent where the latter's effort choice is observable. He shows that it might be optimal to let the agent slack off after signing the contract (where this slacking off is not anticipated by a naive agent), however requiring the agent to "compensate" the principal by accepting a lower compensation. We relate to Gilpatric (2008) in the sense that agents deviate from the planned choices post entering the contract, which can be exploited by the principal. However, since principal and agent only interact once in Gilpatric (2008), his model contains no dynamics. Moreover, it does not allow to make inferences on the effects of the institutional environment (for example captured by the extent of employment protection, bargaining power, or limited-liability) on labor market outcomes.

Both Li, Yan, and Yu (2012) and Yılmaz (2013) analyze a repeated moral hazard setting with a risk-neutral principal and a risk-averse agent, where the latter has β–δ-preferences. Yılmaz (2013) shows that due to present-biased preferences, a lesser degree of consumption smoothing is optimal. Similarly, Li, Yan, and Yu (2012) find that the principal might optimally sell a risky project to a risk-averse agent. However, both restrict the contracting space to consist of only one element, and hence the agent sticks to his planned actions in future periods.

From a modeling perspective the paper closest to ours is Heidhues and Kőszegi (2010). However, they analyze a rather different environment in which firms in a competitive credit market interact with present-biased consumers. They find that naive consumers are attracted by contracts which are characterized by cheap baseline repayment terms and large fines for delays. Naive consumers over–borrow and end up paying fines and suffering welfare losses. The results in our present paper, as in many other papers with present-biased agents, are driven by a fundamentally similar basic intuition. Nevertheless, our papers differs in terms of application, interpretation of the contract features, relevance of institutional details, and policy implications.

Furthermore, Eliaz and Spiegler (2006) analyze optimal screening contracts for agents with different degrees of sophistication concerning their future preferences. They show that the principal benefits monotonically from a higher degree of an agent's naiveté. We extend their basic idea to a labor market relationship, show how a dynamic exploitative contract is optimally designed, and explore the interaction of an agent's bias with various properties of the labor market.

3 Model Setup

Technology & Per-Period Utilities There is one risk neutral principal ("she") who can employ a risk neutral agent ("he"). We consider an infinite time horizon with discrete time periods $t = 0, 1, 2, ...$ If the agent is employed by the principal in any period t, he chooses effort $e_t \geq 0$. This choice is associated with effort costs $c(e_t)$, where $c(e_t)$ is a strictly increasing, differentiable and convex function, with $c(0) = 0, c'(0) = 0$, and $\lim_{e_t \to \infty} c' = \infty$. Effort e_t generates a deterministic output $y(e_t) = e_t\theta$ which is consumed by the principal and where $\theta > 0$ captures the value of the agent's effort. Furthermore, the agent receives a fixed wage payment w_t and a bonus $b_t(e_t)$. Note that we do not impose non-negativity constraints on payments, which hence can (and under some instances will) be negative (implying payments from the agent to the principal). We consider the case of limited liability below, in Section 6.3. The agent's payoff in a period t when employed by the principal is

$$w_t + b_t(e_t) - c(e_t),$$

whereas the principal receives

$$e_t\theta - b_t(e_t) - w_t.$$

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 in this case is denoted by $\overline{\pi} \in \mathbb{R}$.

The effort level maximizing total surplus if the agent works for the principal, denoted by e^{FB}, is defined by

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

For the remainder of the paper, we assume $\theta e^{FB} - c(e^{FB}) - \overline{u} - \overline{\pi} > 0$, i.e., the employment relationship is socially efficient.

Time Preferences The principal discounts the future exponentially with a constant factor $\delta \in (0, 1]$, whereas the agent discounts his future utility in a quasi-hyperbolic way according to Phelps and Pollak (1968) and Laibson (1997): While current payoffs are not discounted, future (period-t) payoffs are discounted with a factor $\beta\,\delta^t$, with $\beta \in (0, 1]$ and δ being identical to the principal's discount factor. Hence, the agent is present-biased and his preferences are dynamically inconsistent. Concerning the agent's belief about his future preferences for instant gratification, we follow O'Donoghue and Rabin (2001) and their description of partial naiveté. While an agent discounts the future with the factor β, he thinks that in any future period, he will discount the future with the factor $\hat{\beta}$, where $\beta \leq \hat{\beta} \leq 1$. In other words, the agent may be aware of his present bias but expects it to be weaker than it actually is. We will mainly focus on the two extreme cases $\hat{\beta} = 1$ and $\hat{\beta} = \beta$. The first case describes a fully naive agent who – in every period – thinks that from tomorrow on, his present bias will disappear and he will discount the future exponentially. The second case describes a sophisticated agent who is fully aware of his (future) present bias.

Perceptions We have to distinguish between intra- and inter-player perceptions. Concerning the first, we assume the agent's beliefs to be dynamically consistent as defined by O'Donoghue and Rabin (2001) (p. 129), i.e., the agent's belief of what he will do in period τ must be the same in all $t < \tau$.

Concerning inter-player perceptions, we assume common knowledge concerning the principal's time preferences. Furthermore, the principal is aware of the agent's time preferences and knows his values β and $\hat{\beta}$. This implies that, for a (partially) naive agent, the principal correctly anticipates any (potential) discrepancy between intended and realized behavior. Finally, we assume that the agent believes the principal to share his own perception of himself. A (partially) naive agent hence is convinced that the principal also perceives the agent's future present bias to be characterized by $\hat{\beta}$. In Appendix C.5 we explicitly allow for unobservable types and derive optimal screening contracts.

Note that our assumptions that the principal is not being present-biased and has better information about the agent's naiveté than the agent himself (without the agent being aware of the principal's superior information) borrows from the literature on behavioral IO. As Eliaz and Spiegler (2006), p. 693, put it, "in many markets, it is reasonable to assume that the firm – with its army of marketing experts – has better knowledge of the agents' systematically changing tastes (e.g. their vulnerability to temptation, propensity to procrastinate, or sensitivity to reference points) than some of the agents themselves. [...] In order to justify our interpretation [...], one should assume that either the agent is unaware of the principal's superior knowledge or believes (due to overconfidence) that this knowledge does not apply in his particular case, even if it does apply to the rest of the population." We believe that similar arguments can be applied to labor markets. Firms have HR departments and accumulate vast experience on how agents assess their own characteristics. Moreover, even if the principal herself was present biased (but still was aware of the agent's naiveté), our results would persist because the principal reduces current (real) payments to the agent by as much as possible.

Contractability, Timing, and Histories To isolate the effect of the agent's present bias on the structure of the employment relationship, we abstract from other potential agency problems. Thus, we assume that the agent's effort as well as wage and bonus payments are verifiable. Moreover, the principal can commit to long-term contracts, where however her commitment is limited by the firm value. Hence, she can always escape her obligations by shutting down.[8] Shutting down is an irreversible decision, inducing principal and agent to consume their outside utilities $\overline{\pi}$ and $\overline{u}$ in every subsequent period. Alternatively, we might assume that the principal can fire the agent at some cost, for example reflecting the degree of employment protection. The costs of shutting down or firing the agent would then be given by $-\frac{\overline{\pi}}{1-\delta}$.

We do not allow the agent to commit to long-term contracts. Note that this assumption turns out to be without loss of generality: Since the principal benefits from the (partially) naive agent's misperceptions of his future choices, she actually does not want to grant the agent the opportunity to commit to a long-term contract. Hence, we can restrict attention to situations where the agent is free to leave in every period.[9]

Now, in the first period of the game, in $t = 0$, the principal makes a take-it-or-leave-it long-term contract offer to the agent, which is denoted by $\boldsymbol{C}$. This offer consists of a menu of contracts for every future period, contingent on any possible history. Each of these contracts contains a fixed wage payment and a bonus for the desired effort level. The menu of contracts offered in period t is denoted by C_t. Without loss of generality, we can restrict C_t to consist of either one or two elements, depending on the agent's naiveté. If the agent is sophisticated or not present-biased, he correctly anticipates his future behavior, and C_t consists of exactly one element. If the agent is (partially) naive, the principal finds it optimal to let C_t consist of exactly two elements: the contract that the agent believes

[8]Restricting the principal's commitment is necessary to ensure existence of an equilibrium, an aspect we further elaborate below.

[9]Again, any costs for the agent to leave his current occupation could be captured by an appropriate choice of $\overline{u}$.

to choose in the future, and the contract the agent is actually going to select.[10] We call the former virtual contract and describe the respective components with a superscript "v". The latter is called real contract and its components are described with the superscript "r". We interpret the agent's selection of a contract as his choice of one of two available career paths.

Thus, $C_t = \{C_t^r, C_t^v\} = \{(w_t^r, b_t^r(e_t)), (w_t^v, b_t^v(e_t))\}$, where $w_t^{r/v} \in \mathbb{R}$, and $b_t^{r/v}(e_t) \in \mathbb{R}$ for each $e_t \in \mathbb{R}^+$. At the beginning of every period t, the principal first decides whether to shut down or whether to offer C_t. In the former case, principal and agent consume their outside utilities. In the latter case, if choosing to work for the principal, the agent selects one element out of C_t. Then, the agent receives $w_t^{r/v}$ and makes his effort choice, triggering the (automatically enforced) bonus payment $b_t^{r/v}(e_t)$. Finally, the principal consumes the output $e_t\theta$, and the game moves on to the next period.

Strategies Following O'Donoghue and Rabin (1999b), we use the phrase *perception-perfect strategy* to describe players' strategies. Such a strategy specifies a player's behavior given dynamically consistent beliefs about future behavior. Whereas a time-consistent or a sophisticated agent correctly anticipates his future behavior, the same is not necessarily true for a (partially) naive agent who has wrong beliefs concerning his future time preferences.

The principal's strategy is denoted by σ_P. In period $t = 0$, it determines the long-term contract $\boldsymbol{C}$. In every period $t \geq 0$, $\boldsymbol{C}$ maps the history of the game into an offered menu of contracts, C_t. Finally, σ_P determines whether the principal follows $\boldsymbol{C}$ or terminates the contract by shutting down.

An agent's strategy is denoted by σ_A and – given the history – determines whether he decides to work for the principal in a period t, as well as his choice from C_t and how much effort e_t to exert.

Note that it is optimal to have stationary real and virtual contracts (besides the first period of the virtual contract), which we formally prove in an earlier version of this paper (see Englmaier, Fahn, and Schwarz (2018)). Therefore, in the remainder of this article, we use the subscript 1 for the first period of the virtual contract. For all subsequent periods, we omit time subscripts. The real contract hence consists of w^r, e^r, b^r, and the virtual contract of w_1^v, e_1^v, b_1^v for its first and w^v, e^v, b^v for all subsequent periods after it has been

[10] Similar features of an "optimal" exploitative contract have been derived previously by Eliaz and Spiegler (2006) and Heidhues and Kőszegi (2010).

selected.

Payoffs In the following, we describe real and perceived payoff streams on the path of an equilibrium in which the agent is employed in every period t. His realized utility stream equals

$$U^r = b^r + w^r - c(e^r) + \beta \frac{\delta}{1-\delta} \left(b^r + w^r - c(e^r)\right).$$

Naive types perceive their utility stream in a period t and from the perspective of period t (indicated by the superscript "rv") to be

$$U^{rv} = b^r + w^r - c(e^r) + \beta\delta \left[(b_1^v + w_1^v - c(e_1^v)) + \frac{\delta}{1-\delta} \left(b^v + w^v - c(e^v)\right) \right].$$

The principal's payoff in any period t equals

$$\Pi^r = \frac{(e^r\theta - b^r - w^r)}{1-\delta},$$

whereas the naive agent perceives it to be

$$\Pi^{rv} = e^r\theta - b^r - w^r + \delta \left[(e_1^v\theta - b_1^v - w_1^v) + \frac{\delta}{1-\delta} \left(e^v\theta - b^v - w^v\right) \right].$$

For later use, we further define perceived future virtual payoffs, which amount to $U^v = \left(b^v + w^v - c(e^v)\right)/(1-\delta)$ for the agent and $\Pi^v = \left(e^v\theta - b^v - w^v\right)/(1-\delta)$ for the principal.

Equilibrium We apply the concept of *perception-perfect equilibrium* which maximizes players' payoffs, given each player's perception of their own future behavior as well as of the other's future behavior.

The principal hence chooses $\boldsymbol{C}$ to maximize Π^r. In all later periods, the principal is left with the decision whether to shut down or not, and proceeds to maximize Π^r.

The (partially) naive agent chooses her actions to maximize U^{rv} in every period. He further expects the principal to also consider Π^{rv} (and later on Π^v) instead of Π^r.

4 Benchmarks: Agent Witout Present Bias and Sophisticated Agent

We first derive two benchmarks, profit-maximizing contracts for agents who are not present-biased, as well as for sophisticated agents, who are fully aware of their bias.

Agent without present bias Consider an agent without a present bias, hence $\beta = \hat{\beta} = 1$. Since the agent's effort is verifiable, a series of spot contracts maximizes the surplus as well as the principal's profits. One possibility to generate such an outcome is to offer the following contract in every period t:

$w = \overline{u}$, $b(e^{FB}) = c(e^{FB})$, and $b(\tilde{e}) = 0$ for $\tilde{e} \neq e^{FB}$.

The agent always accepts such a contract and exerts effort e^{FB}. Since the principal extracts the whole surplus, she has no incentive to ever shut down.

Sophisticated present-biased agent A sophisticated present-biased agent (i.e., $\hat{\beta} = \beta$), correctly anticipates his future choices. Thus, it is sufficient to let $\boldsymbol{C}$ consist of only one element in every period, and a series of spot contracts (as with an agent without present bias) maximizes surplus and profits: $w = \overline{u}$, $b(e^{FB}) = c(e^{FB})$, and $b(\tilde{e}) = 0$ for $\tilde{e} \neq e^{FB}$.

This contract makes the agent accept the contract in every period, induces him to choose the surplus-maximizing effort level, and allows the principal to extract the whole surplus. Note that this contract cannot be improved upon by taking into account the agent's effective lower discount factor and shifting payments to period 0. In such a case, the agent would simply walk away after this period.

Thus, the agent's present bias does not affect the profit-maximizing contract if he is sophisticated. This follows from effort being verifiable and the static production technology, which allows to immediately compensate the agent for his effort.

5 The Principal's Problem with a Naive Agent

5.1 Preliminaries

This section considers the principal's problem when facing a naive present-biased agent, i.e., an agent with $\hat{\beta} = 1$. In Subsection 6.4, we argue that the results derived here remain unaffected for any $\hat{\beta} \in (\beta, 1)$. First, note that the same contract as for a sophisticated agent could be offered to a naive agent. Thus, the existence of a present bias does not automatically trigger inefficiencies. This also implies that the principal never shuts down on the equilibrium path, and the agent always remains with the principal.

However, the principal can design a menu of contracts to exploit the naive agent's misperception of his future behavior and thus make higher profits than with a sophisticated agent. Now, $\boldsymbol{C}$ includes the virtual contract that appears optimal for the agent from the perspective of earlier periods, as well as the real contract that is actually selected by the agent in every period. Note that here it is without loss of generality to assume that, once the agent has chosen the virtual contract, he subsequently cannot go back to the real contract.

The elements of $\boldsymbol{C}$ must satisfy four classes of constraints, where the first three are standard in contracting problems: 1) **Individual rationality constraints** for the agent (IRA), which make him accept one of the offered contracts. 2) **Individual rationality constraints** for the principal (IRP), which keep her from shutting down. 3) **Incentive compatibility constraints** for the agent (IC), which induce him to select equilibrium effort. These constraints must also hold for the virtual contracts, i.e., for future histories that never materialize on the equilibrium path. Finally, 4) **Selection constraints** ensure that the agent keeps choosing the real contract in every period while intending to select the virtual contract in all future periods.

Individual rationality constraints for the agent It must be optimal for the agent to accept the real contract (when expecting to select the virtual contract in all future periods) instead of rejecting all contracts and consuming $\overline{u}$ in the respective period. There, it is optimal to assume that, once the agent deviates, he receives his outside option in all later periods (for example by subsequently being offered the profit-maximizing contract for a sophisticated agent):

$$w^r + b^r - c(e^r) - \overline{u} + \beta\delta \left[(b_1^v + w_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta} (b^v + w^v - c(e^v) - \overline{u}) \right] \geq 0. \tag{rIRA}$$

Furthermore, the agent must expect to accept the contract in all future periods, i.e., in the first period of the virtual contract (v1IRA) as well as in all following periods (vIRA).

$$b_1^v + w_1^v - c(e_1^v) - \overline{u} + \frac{\delta}{1-\delta} (b^v + w^v - c(e^v) - \overline{u}) \geq 0, \tag{v1IRA}$$

$$b^v + w^v - c(e^v) - \overline{u} \geq 0. \tag{vIRA}$$

Note that an individual rationality constraint for U^r (i.e., $w^r + b^r - c(e^r) - \overline{u} \geq 0$) does *not* have to hold because the agent does not expect to select the real contract in future periods. In fact, this constraint will indeed be violated and the agent will receive less than his outside option.

Individual rationality constraints for the principal First, it must not be optimal for the principal to shut down, hence $e^r\theta - b^r - w^r - \overline{\pi} \geq 0$ has to hold. Since the principal cannot do worse than with a sophisticated agent, her real profits will always be positive and this constraint can be omitted. Moreover, the constraint $\Pi^{rv} \geq \overline{\pi}/(1-\delta)$ implies that shutting down must not be optimal for the principal if the agent follows up on his plan to select the real contract today and the virtual contract from tomorrow on. Furthermore, there are (IR) constraints for the first period of the virtual contract (v1IRP), as well as all following periods (vIRP),

$$e_1^v\theta - b_1^v - w_1^v - \overline{\pi} + \frac{\delta}{1-\delta} (e^v\theta - b^v - w^v - \overline{\pi}) \geq 0, \tag{v1IRP}$$

$$e^v\theta - b^v - w^v - \overline{\pi} \geq 0. \tag{vIRP}$$

If either of these constraints was not satisfied, the agent would expect the principal to shut down and not honor her obligations in the virtual contract.

Incentive compatibility constraints It has to be in the agent's interest to choose equilibrium effort e^r. He will only do so if he is sufficiently punished after a deviation, and many arrangements can be used for that purpose because of the verifiability of effort. For example, a "forcing contract" with infinitely negative payments in case the agent does not choose equilibrium effort could be written (and IC constraints omitted). We abstain from such a simplification because it allows us to easily extend our setup to analyze moral hazard. Moreover, one could argue that the agent's punishment is restricted, for example by his outside option. Here, we assume that, after a deviation, the agent only receives his output $\overline{u}$ in this and all subsequent periods, noting that any stronger punishment would not change our results. Thus, the agent will choose an effort level of zero in case he deviates, which yields the following incentive compatibility constraints for real and virtual effort levels.

$$b^r - c(e^r) - \overline{u} + \beta\delta\left[(b_1^v + w_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}(b^v + w^v - c(e^v) - \overline{u})\right] \geq 0, \tag{rIC}$$

$$b_1^v - c(e_1^v) - \overline{u} + \frac{\delta}{1-\delta}(b^v + w^v - c(e^v) - \overline{u}) \geq 0, \tag{v1IC}$$

$$-c(e^v) + b^v + \frac{\delta}{1-\delta}(b^v + w^v - c(e^v) - \overline{u}) \geq 0. \tag{vIC}$$

Selection constraints Finally, the agent must expect to select the virtual contract in the future, while going for the real (instead of the virtual) contract when actually making the choice. This yields the constraints

$$\begin{aligned} & w^r + b^r - c(e^r) + \beta\delta\left[(b_1^v + w_1^v - c(e_1^v)) + \frac{\delta}{1-\delta}(b^v + w^v - c(e^v))\right] \\ \geq & b_1^v + w_1^v - c(e_1^v) + \beta\frac{\delta}{1-\delta}(b^v + w^v - c(e^v)), \end{aligned} \tag{rC}$$

$$\begin{aligned} & b_1^v + w_1^v - c(e_1^v) + \frac{\delta}{1-\delta}(b^v + w^v - c(e^v)) \\ \geq & w^r + b^r - c(e^r) + \delta\left[(b_1^v + w_1^v - c(e_1^v)) + \frac{\delta}{1-\delta}(b^v + w^v - c(e^v))\right]. \end{aligned} \tag{vC}$$

Objective The principal's objective is to offer a menu of contracts $\boldsymbol{C}$ that maximizes her real profits π^r in every period, subject to the constraints just derived.

5.2 Profit-Maximizing Contract

In this section, we fully characterize a profit-maximizing menu of contracts $\boldsymbol{C}$. First, the program can be substantially simplified, which we do in Appendix B. For example, due to the verifiability of the agent's effort, it is without loss of generality (Lemma B1) to set all wages to zero, hence $w^r = w_1^v = w^v = 0$. Thus, the bonus is used to compensate the agent for his effort cost, but also for his outside option and to grant him a potential rent.

In order to extract rents from the agent, the principal shifts as much as possible of the agent's compensation into the virtual contract. However, the principal has to ensure that the virtual contract is never selected by the agent. This is achieved by designing the first period of the virtual contract sufficiently unattractive for the agent who discounts the future with an additional factor β (but attractive enough for an exponential discounter that the agent anticipates to be in future periods). In Appendix B (Lemma B2), we further show that $-c(e^r)+b^r < \overline{u}$, which implies that the naive agent is indeed exploited by the principal.

Lemma 1. *A profit-maximizing menu of contracts has the following features:*

- *The constraints (rIRA), (rC), and (vIRP) hold with equality,*
- $e^r = e^v = e^{FB}$,
- $b^r = c(e^{FB}) + \overline{u} - \beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$.

The first bullet point of Lemma 1 indicates that the principal promises the whole virtual rent to the agent (binding (vIRP))[11], with the exception of the first virtual period which is designed sufficiently unattractive to never be selected by the agent (binding (rC)). Furthermore, even though the agent expects a rent in the future, his perceived rent from today's perspective (including today's payoffs from the real contract) is zero (binding (rIRA)).

First-best effort levels in the real and virtual contract (with the exception of e_1^v) not only maximize the surplus that the principal can reap in any case; $e^v = e^{FB}$ also maximizes the future rent the agent expects, making it possible to reduce real payments by the highest feasible amount.

[11]This result illustrates the need to limit the principal's commitment and allowing her to shut down. If she was able to (credibly) make arbitrary promises (i.e., if she did not face (IR) constraints), there would be no profit-maximizing equilibrium. In this case, the principal would promise infinitely high payments to the agent in the virtual contract, and extract infinitely high payments from the agent in the real contract.

Finally, the components determining the value of b^r capture the link between the agent's real and virtual compensation. The first elements, $c(e^{FB}) + \overline{u}$, are the agent's effort costs and outside option, hence would constitute his "fair" compensation. The last term yields the extent to which the agent is exploited compared to this fair compensation. It amounts to the total expected and discounted rent the agent expects from choosing the virtual contract in the future, i.e., from making a career, and is supposed to serve as the reward for *today's* effort.

The possibility to exploit agents who do not stick to their planned action is reminiscent of extant results in the literature; see, e.g., DellaVigna and Malmendier (2004), Eliaz and Spiegler (2006, 2008), or Heidhues and Kőszegi (2010). In the following, we thus focus on the properties of a profit-maximizing menu of contracts and later show how variations of the original model generate further results on how the presence of present-biased agents might affect labor market outcomes (Section 6).

Corollary 1. *The real bonus b^r is increasing in $\overline{\pi}$ and decreasing in θ.*

The agent's compensation b^r is smaller if the (perceived) relationship surplus is larger. This allows to generate a number of testable predictions which relate determinants of this surplus to workers' wages. As a first example, recall that the principal's outside option $\overline{\pi}$ might also include costs of firing the agent. Such firing costs are supposed to be larger in environments with more stringent employment protection laws. Hence, if regulatory changes increase firing costs and thus reduce $\overline{\pi}$, we expect wages to go down, in particular those of newly formed employment relationships (the principal's long-term commitment restricts her ability to ex-post worsen the conditions for an agent that she already employed). In this vein, Leonardi and Pica (2013) analyze the impact of a 1990 labor market reform in Italy which increased firing costs for firms with 15 employees or less but left costs for larger firms unaffected. Using administrative data from the Italian Social Security Institute for the Italian provinces of Vicenza and Treviso, they indeed find that higher firing costs lead to a slight reduction of average wages. Moreover, this effect is mostly driven by entry wages of workers that have moved jobs. Below, we use variations of our model to derive additional predictions, and relate them to empirical observations. Before doing so, we take a closer look at further implications of our main model.

First, *real* payoffs of the principal and agent are characterized in Proposition 1.

Proposition 1. *Real net per-period payoffs of principal and agent are*

$$\pi^r - \overline{\pi} = \left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\left(1 + \beta(1-\beta)\frac{\delta^2}{1-\delta}\right)$$

and

$$u^r - \overline{u} = -\beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right),$$

respectively.

Proposition 1 implies that agents with levels of β close to 0 or 1 can hardly be exploited by the principal. The real loss for the agent is maximized for intermediate values of β (see Figure 1). Hence, types with intermediate values of β are preferred by the principal (see Figure 2). The reason for this is that agents with a β close to 1 do not depart from time-consistency very much and agents with a β close to 0 do not care very much about the future, so the virtual future surplus cannot have a big impact on today's choices.

****

Figure 1 (agent) and Figure 2 (principal) to appear about here

****

Furthermore, we take a look at the first period of the virtual contract which we interpret as a qualification period to deter the agent from ever selecting C^v. It only depends on the difference $c(e_1^v) - b_1^v$, without the exact values of b_1^v and e_1^v being relevant.[12] Binding (rC), (rIRA), and (vIRP) constraints yield

$$c(e_1^v) - b_1^v = \frac{\delta\beta}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) - \overline{u}.$$

Hence, a policy of subsidizing career-development (which could be captured by a subsidy to effectively increase b_1^v) would not help, since the principal could adjust the required level of e_1^v accordingly in order to keep the agent from selecting the virtual contract.

[12] Note that this also implies that if the agent's present bias would only manifest in one domain, i.e., either in monetary payments or effort, our results would (qualitatively) be the same as we discuss in Appendix C.4.

Figure 3 depicts the path of effort over time and Figure 4 the path of utility over time.

****
Figure 3 and Figure 4 to appear about here
****

A high relationship surplus as well as a high long-run discount factor δ reduce the naive agent's real payoff. Put differently, if the future is not valued very much because making a career does not appear too attractive, then the agent does not want to trade today's payment against future benefits and hence cannot be exploited as much.

6 Variations and Additional Implications

In this section, we look at variations of the original model. Besides serving as robustness tests, these allow to generate additional results on how the existence of (naive) present-biased employees can affect labor market outcomes.

6.1 Finite Time Horizon

While employment relations are in general long term, they still might have a pre-defined (maximum) tenure. In particular, in many markets and countries there exists a mandatory retirement age. Here we document that our results are qualitatively robust to considering a finite time horizon.

Proposition 2. *Assume the game is finite and ends in period $T > 3$. Then, a contract with the following features for all periods $t \leq T$ maximizes the principal's profits:*

- $e_t^r = e_t^v = e^{FB}$.
- $b_t^r = c(e_t^r) + \overline{u} - \beta(1-\beta)\delta^2 \sum_{j=0}^{T-t-2} \delta^j \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$, *with* $\sum_{j=0}^{k} x_j := 0\ \forall k < 0$.
- $u_t^r - \overline{u} = -\beta(1-\beta)\delta^2 \sum_{j=0}^{T-t-2} \delta^j \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$, *with* $\sum_{j=0}^{k} x_j := 0\ \forall k < 0$.
- $\pi_t^r - \overline{\pi} = \left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\left(1 + \beta(1-\beta)\delta^2 \sum_{j=0}^{T-t-2} \delta^j\right)$, *with* $\sum_{j=0}^{k} x_j := 0\ \forall k < 0$.

A longer remaining time horizon implies a larger (virtual) future surplus and consequently lower payments to the agent today, which immediately yields the following corollary.

Corollary 2. *For a finite time horizon T, the real bonus b_t^r is decreasing in $T-t$. That is, the principal can exploit the agent more in the early stages of his career. Moreover, the negative effect of a smaller $\overline{\pi}$ on b_t^r decreases over time.*

Not only is b_t^r larger in later periods of the employment relationship, but also the negative effect of a higher relationship surplus on the agent's compensation is less pronounced. This result can again be related to observations made by Leonardi and Pica (2013). They find that the aforementioned negative effect of a more stringent employment protection on wages is stronger for young workers below 30.

6.2 Labor Market Competition and Bargaining

So far we have assumed that the principal has full bargaining power and can hence determine the terms of the employment relationship. In this section, we explore the implications of the agent having positive bargaining power. We do not explicitly model the bargaining process, but assume that players arrive at a Nash bargaining outcome where the principal keeps a share α and the agent a share $1-\alpha$ of the relationship surplus. More precisely, the agent accepts any offer that leaves him with $1-\alpha$ of his "present-biased view" of the total relationship surplus. Note that, because the agent is free to leave at any time, he cannot commit not to renegotiate an agreement. Thus, a front-loading of the agent's compensation is not feasible – although it might be optimal given the agent's lower effective discount factor – and a condition stating that the agent must receive at least $1-\alpha$ of the total relationship surplus has to hold in every period. The principal, on the other hand, can commit to long-term contracts, hence she is able to backload the agent's compensation.[13]

Then, the structure of the optimal menu of contracts for a naive agent remains unaffected. However, although the agent is still exploited for any strictly positive α, a lower value of α makes him better off in real terms.

As we show in Appendix C.1, the agent's real utility and his utility in the first period of the virtual contract are

[13]This is different from, e.g., Miller and Watson (2013) and Fahn (2017), where the principal's inability to commit not to renegotiate any agreement has substantial negative consequences on the efficiency of a long-term employment relationship.

$$u^r - \overline{u} = \left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\left((1-\alpha) - \frac{\alpha\,(1-\beta)\,\beta\delta^2}{1-\delta}\right)$$

and

$$u_1^v - \overline{u} = \left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\left((1-\alpha) - \frac{\alpha\beta\delta}{1-\delta}\right).$$

Therefore, the agent receives less than his fair share which would amount to $\overline{u} + (1-\alpha)\left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)$ in every period. Furthermore, the agent's real payoff is decreasing in α. Only for $\alpha = 0$, the agent cannot be exploited. Then, $u^r = u_1^v = u^v$, and the distinction between real and virtual contract becomes immaterial.

The above analysis delivers a further interesting takeaway. In our benchmark model (which corresponds to $\alpha = 1$), the agent is better off with a smaller (perceived future) relationship surplus. This comparative static reverses if α is small (more precisely, if $\alpha \leq (1-\delta)\,/\,[1-\delta+(1-\beta)\,\beta\delta^2]$), in which case the agent benefits from a higher relationship surplus. Taking up the interpretation of a lower outside option corresponding to higher layoff costs or more restrictive employment protection laws, the agent is only harmed by those laws if his bargaining power is relatively low. Indeed, Leonardi and Pica (2013) find that the negative effect of higher dismissal costs on wages is larger for workers with low bargaining power such as young blue-collar workers or those with a small premium over the sectoral contractual minimum compensation.

Finally, note that relative bargaining power can also be driven by the competitiveness of the labor market in which case a smaller α would correspond to more competition for agents. Then, our results indicate that more competition for workers helps present-biased individuals, without letting the scope for their exploitation completely disappear.

6.3 Limited Liability

We have not imposed any restrictions on the size of b^r which thus might actually be negative, then indicating payments from the agent to the principal in addition to the agent's effort. In many cases, though, payments are restricted by some lower bound. In this subsection, we assume that the agent is protected by limited liability, i.e., payments can not be negative ($b^r \geq 0$). Proposition 3 characterizes a profit-maximizing contract when payments have to be non-negative.

Proposition 3. *Assume $b, w \geq 0$ in every period t. Then, a profit-maximizing contract $\boldsymbol{C}$ has $e^r \geq e^{FB}$, $e^v = e^{FB}$, and $b^r = max\left\{0,\ c(e^{FB}) + \overline{u} - \beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)\right\}$. Moreover,*

- *if $c(e^{FB}) + \overline{u} - \beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) < 0$, the limited liability constraint for the agent's real bonus binds. Then $e^r > e^{FB}$ and is chosen to satisfy $c(e^r) + \overline{u} - \beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) = 0$.*
- *If $c(e^{FB}) + \overline{u} - \beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) \geq 0$, the limited liability constraints do not bind and hence do not affect $\boldsymbol{C}$.*
- *In either case, the agent's real payoff is $u^r - \overline{u} = -\beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$.*

Interestingly (and different from moral hazard problems with limited liability where the agent generally receives a rent), the agent is not better off than without non-negativity constraints on payments but receives exactly the same level of real utility. This is because the principal optimally responds by letting the agent work harder and setting e^r above e^{FB}. Hence, we might face a situation where the naive present-biased agent works harder than a sophisticated agent or one without a present bias. Moreover, the full burden of the inefficient outcome is borne by the principal: Because inefficiently high effort reduces total surplus, she can extract less from the agent than when payments are not restricted.

Proposition 3 also allows us to derive an interaction between the productivity of the employment relationship and inefficiencies caused by the limited liability constraint.

Corollary 3. *For any given cost function $c(\cdot)$ and parameters $\overline{u}, \overline{\pi}$, there exists a $\underline{\theta}$ such that for all $\theta > \underline{\theta}$, $e^r > e^{FB}$. Moreover, for all θ, $de^r/d\theta \geq 0$.*

A more productive employment relationship generally triggers larger inefficiencies in the presence of limited liability constraints.[14] The latter are not only more likely to bind if θ is larger, but already existing inefficiencies due to inefficiently high effort are also exacerbated.

Finally, note that having a limited liability constraint with a finite time horizon (as analyzed in the previous section) also yields interesting effort dynamics: With limited

[14]Note that, if we allowed for differences between future and current levels of productivity, this effect would be driven by the former.

liability, (real) effort is increasing in the future virtual surplus. With a finite time horizon, this future surplus decreases over time. Therefore, effort gradually decreases in this case, until limited liability cedes to bind.

Corollary 4. *Assume $b, w \geq 0$ in every period t. For a finite time horizon T, the optimal e^r_t is weakly decreasing in t. It is strictly decreasing as long as $e^r_t > e^{FB}$.*

6.4 Further Variations

Here, we provide an overview of how our results are robust to additional variations. A more detailed analysis can be found in subsections of Appendix C.

Moral Hazard The results of our basic model do not change qualitatively with moral hazard, i.e., if the agent's effort is his private information and only output is contractible. The structure of the optimal contract menu for the naive agent is the same as without moral hazard, with the exception that the bonus is only paid if output has been high (see Appendix C.2). In fact, the results of our model do not depend on effort provision but on different streams of utility in different contracts.

Partial Naiveté and Learning In our main model, we assume that the agent fails over and over again: In every period he selects the real contract, although having planned to take up the virtual contract. This assumption seems quite strong. One might expect that after a couple of failed attempts to actually select the virtual contract, the agent might realize that he has a present bias (for fully naive agents) or that his present bias is stronger than he thought it would be (for partially naive agents). In Appendix C.3, we show that even with some learning by the agent, the principal continues to offer the same contracts as long as the agent does not become fully sophisticated. The intuition is as follows. A partially naive agent expects that in the future choosing the virtual contract is strictly better than choosing the real contract. But tomorrow he will be just indifferent between the two contracts, his assessment being driven by his true β but not his expectation, $\hat{\beta}$. Hence, for a given β, the principal will only offer two different sets of contracts to all agents: the virtual and the real contract for naive agents (no matter whether they are fully or only partially naive) and spot contracts for sophisticated and time-consistent agents.[15] Because the latter would not be chosen by an agent who is (still) naive, it could always be added to the menu of contracts – in particular if a naive agent might eventually learn his true β.

Weaker Present Bias in the Monetary Domain In Appendix C.4, we assume that the agent has a stronger present bias in effort than in money. We show that our results are unchanged, except that having a weaker bias in the monetary domain *hurts* a naive present-biased agent. The reason is that a less severe present bias in the monetary domain lets the agent discount his virtual compensation to a lesser extent, making the virtual contract appear even more appealing.

Unobservable Agent Types So far we have assumed that the principal can tailor her offer to the agent's type. However, in real-world contexts the principal might not know the exact extent of the agent's present bias. We show in Appendix C.5 that our results are robust to considering the case of unobservable agent types: Agents are optimally separated, and each receives a menu tailored to his type. Moreover, menus generally are still exploitative in the same vein as before, only the extent of exploitation is reduced for some in order to prevent one type from mimicking another. Intuitively, the principal will focus on exploiting types she is more likely to face, and hence induces a separation by reducing the exploitative rents she collects from less frequent types.

7 Discussion: Present Bias versus Overoptimism

In the above analysis we have assumed that the agent mispredicts his future behavior because he underestimates his discount factor in future periods. In this section, we analyze whether the generated outcomes can be distinguished from a setting in which the agent is *overoptimistic* regarding his productivity in the employment relationship. Assuming that the agent underestimates the size of his future effort costs, we show that the principal still exploits the agent by offering a virtual contract the agent expects to select in future periods, whereas actually going for the real contract.

[15] Such a discontinuity in the form of contracts is a common feature in the literature that look at different degrees of sophistication, whether all but the fully sophisticated types receive exactly the same contract (as in our paper; also see Heidhues and Kőszegi, 2010) or very similar contracts (e.g., DellaVigna and Malmendier, 2004).

However, under overoptimism the structure of the virtual contract can differ compared to the setting with a naive present-biased agent as the virtual contract does not have to entail an unattractive "qualification" period which the agent first has to pass before receiving a rent thereafter. Instead, the first period of the virtual contract does not only have to contain high effort, but can also include the *complete* rent the agent is offered in exchange for low payments in the real contract. In all further periods of the virtual contract, the agent is left with his outside option. Therefore, we would predict different (perceived) career paths for overoptimistic agents compared to those with a present bias.

To formally derive this result, we assume that the agent's effort costs amount to $c(e)$, but that he perceives his effort costs in all future periods to be $kc(e)$, with $k < 1$; hence, the agent is overoptimistic with respect to his future productivity. Moreover, he now discounts the future exponentially. As before, the principal is aware of the agent's perceptions, as well as his true characteristics.

In any period of an employment relationship, the principal still offers a menu of contracts which contains a real contract with components e^r and b^r and a virtual contract with components e_1^v and b_1^v (for the first period of the virtual contract), as well as e^v and b^v (for all later periods of the virtual contract). Several constraints make sure that the agent i) accepts the real contract (compared to going for his outside option and compared to selecting the virtual contract), but ii) expects to choose the virtual contract from next period on. Moreover, the principal's profits must be sufficiently large (in terms of realized outcomes, but also from the agent's perspective) to rule out that firing the agent and consuming her outside option is optimal. Generally the set of constraints is the same as with a present biased but naive agent and delivered in the proof to Proposition 4.

Proposition 4. *Assume the agent is overoptimistic regarding his future effort costs, as described above. Then, there exists a profit-maximizing menu of contracts that has the following properties.*

1. *For the first period of the virtual contract, the agent is offered the complete relationship surplus supposed to be generated in this and all subsequent periods. Hence,* $b_1^v = e_1^v\theta - \overline{\pi} + \frac{\delta}{1-\delta}\left(e^v\theta - kc(e^v) - \overline{u} - \overline{\pi}\right)$.
2. *For all later periods of the virtual contract, the principal keeps the complete relationship surplus supposed to be generated in the employment relationship. Hence,* $b^v = kc(e^v) + \overline{u}$.

3. e^r *is characterized by* $\theta - c' = 0$, e^v *by* $\theta - kc' = 0$.

4. *If* $c(e^v)\,(1-k) \geq (e^v\theta - kc(e^v) - \overline{u} - \overline{\pi})\,/\,(1-\delta)$, *then* $e_1^v = e^v$. *Moreover,*

$$u^r - \overline{u} = -\frac{\delta\,(e^v\theta - kc(e^v) - \overline{u} - \overline{\pi})}{1-\delta}.$$

Otherwise, e_1^v *is characterized by*

$$e_1^v\theta - c(e_1^v) - \overline{u} - \overline{\pi} + \frac{\delta}{1-\delta}\,(e^v\theta - kc(e^v) - \overline{u} - \overline{\pi}) = 0.$$

Moreover,

$$u^r - \overline{u} = -\delta\,(1-k)\,c(e_1^v).$$

In the derived contract, the agent's compensation for low real payments is concentrated in the first period of the virtual contract. In later periods, the agent still works hard, but the generated rents remain with the principal and provide commitment to credibly promise larger payments in the qualification period. On the contrary, a naive present-biased agent is only supposed to work hard in the first period of the virtual contract, with the compensation following in the remainder of his career. Thus, a (naive) present-biased agent has a considerably different expected career path than an agent who is overoptimistic regarding his future effort costs.

However, the derived contract for an overoptimistic agent is not unique in maximizing the principal's profits. Both the principal and the agent are indifferent between (appropriately discounted) payoffs in the first or later periods of the virtual contract. Thus, reducing payments in the first virtual period and increasing it in later periods would also maximize profits and resemble the profit-maximizing contract for a naive present-biased agent. But outcomes would be unique if the principal had a (if only marginally) larger discount factor than the agent. Then, the principal's profits would be uniquely maximized by concentrating the agent's rent in the first period of the virtual contract, as specified in Proposition 4. In this case, an even sharper distinction between outcomes for an overoptimistic and a naive present-biased agent could be drawn.

In any case, having a present-biased agent allows for clearer predictions since it is uniquely optimal to grant the agent the complete rent in all but the first periods of the virtual contract. The reason is that moving wage payments from the first to later periods in the virtual contract reduces the attractiveness of selecting the virtual contract today (and thereby relaxes the rC constraint), but does not affect the relative attractiveness between real and virtual contract from the perspective of earlier periods.

Finally, the real rent of an overoptimistic agent is monotonically decreasing in the extent of his overoptimism (u^r is increasing in k). A naive present-biased agent, on the other hand, can best be exploited for intermediate values of β (see Section 5.2). This means, even if overoptimism and present bias were highly correlated, one could empirically asses the relative importance of the two in profit-maximizing dynamic incentive schemes.

8 Conclusion

We have shown how a principal can design a profit-maximizing long-term contract when employing a present-biased agent. The principal can take advantage of an agent's naiveté by promising career prospects that involve short-term sacrifices. In doing so and by offering a menu of contracts, the principal can push the agent even below his outside option: a virtual contract, which consists of a relatively low compensation in the initial period but promises high future benefits, and a real contract, which keeps the agent below his outside option. The agent expects that he will choose the virtual contract from the next period on and therefore accepts a lower compensation today. However, he always ends up choosing the real contract and hence never gets to enjoy the generous benefits from the virtual contract. A higher surplus in the virtual contract affects the agent's real compensation and potentially also implemented effort (in the case of limited liability). These results give rise to a number of testable implications – beyond the ones already explored above (such as the consequences of a more stringent employment protection).

For example, our results would predict that a raise in the productivity parameter θ increases a naive agent's compensation if α (the principal's bargaining power in Section 6.2) is small, but decreases his compensation otherwise. Generally, empirical evidence suggests that some rent-sharing between firms and workers exists (Van Reenen, 1996, Bell, Bukowski, and Machin, 2018). However, the link between productivity and wages seems to have been vanishing in many Western countries (Benmelech, Bergman, and Kim, 2018, Table V), particularly for lower parts of the wage distribution (Acemoglu and Autor, 2011). One prominent explanation is a lower bargaining power of workers, for example caused by declining unionization (Benmelech, Bergman, and Kim, 2018). If our mechanism also contributed to this development (besides the direct consequences of a reduced worker bargaining power), one could explain a negative (and not just declining but still positive) relationship between a firm's productivity and the wages of (at least some) workers and would expect one in settings in which workers do not have a lot of bargaining power. However, the focus of the aforementioned papers has been to establish a *reduced* link between pay and productivity, and more empirical research is needed to explore the existence of a potential negative interaction. Yet, some results seem to indicate that such a negative link might indeed exist. For example, Bell, Bukowski, and Machin (2018) explore the extent of rent sharing between firms and workers in Britain, and focus on the decline from the 80s and 90s to the 2000s. Although the link is in general still positive, a closer look at their results reveals that some interactions between profits and wages are (significantly) negative, in particular in more recent years when workers arguably had less bargaining power (Bell, Bukowski, and Machin, 2018, tables 4 and 5). Furthermore, Autor et al. (2017a,b) explore causes for the fall of labor's share of GDP in the US and provide micro-evidence on its evolution. They find that a declining labor share goes together with an increased concentration on product markets (in which case we would expect remaining firms to experience a larger θ). Their explanation involves a re-allocation of market shares to "superstar firms" with low labor costs and can account for a major part of the decline. Our model highlights a complementing (micro-)channel (that an increase in productivity is not completely reflected in wage increases). Moreover, also within-firm effects are negative and significant in their paper, but not at the center of the authors' attention. For our purpose, it would be relevant to further pursue these within-firm effects, and in particular whether an increased concentration on product markets also reduced the wages of individual workers. This is of particular interest if one believes that productivity and the firms' bargaining power will increase in the future.

Finally, taking our results with limited liability into account, a larger θ (in combination with low worker bargaining power) might not only lower wages, but could also increase the effort of (in particular, younger) workers. Thus, we would predict that higher productivity/profits can also yield an (unpaid) increase in working hours.

Taken together, we regard it a fruitful exercise to further explore the forces underlying the tremendous transformations of the conditions workers had to experience in the last decades, and what role behavioral biases, in particular present bias, have played in mediating those transformations.

References

Acemoglu, D., and D. Autor (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, ed. by D. Card, and O. Ashenfelter, vol. 4, pp. 1043 – 1171. Elsevier.

Augenblick, N., M. Niederle, and C. Sprenger (2015): "Working Over Time: Dynamic Inconsistency in Real Effort Tasks," *Quarterly Journal of Economics*, 130(3), 1067–1115.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2017a): "Concentrating on the Fall of the Labor Share," *American Economic Review*, 107(5), 180–85.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2017b): "The Fall of the Labor Share and the Rise of Superstar Firms," Working Paper 23396, National Bureau of Economic Research.

Baum, C. L., and W. F. Ford (2004): "The wage effects of obesity: a longitudinal study," *Health Economics*, 13(9), 885–899.

Bell, B., P. Bukowski, and S. Machin (2018): "Rent Sharing and Inclusive Growth," Working Paper 1584, Center for Economic Performance.

Benmelech, E., N. Bergman, and H. Kim (2018): "Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?," Working Paper 24307, National Bureau of Economic Research.

Böckerman, P., A. Hyytinen, and J. Kaprio (2015): "Smoking and long-term labour market outcomes," *Tobacco Control*, 24(4), 348–353.

Boeckerman, P., A. Hyytinen, and T. Maczulskij (2017): "Alcohol Consumption and Long-Term Labor Market Outcomes," *Health Economics*, 26(3), 275–291.

Brunello, G., and B. D'Hombres (2007): "Does body weight affect wages?: Evidence from Europe," *Economics & Human Biology*, 5(1), 1 – 19.

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

de la Rosa, L. E. (2011): "Relational Incentives and Moral Hazard in Teams," *Games and Economic Behavior*, 73(2), 429–451.

DellaVigna, S. (2009): "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, 47(2), 315–372.

DellaVigna, S., and U. Malmendier (2004): "Contract Design and Self-Control: Theory and Evidence," *Quarterly Journal of Economics*, 119(2), 353–402.

——— (2006): "Paying Not to Go to the Gym," *American Economic Review*, 96(3), 694–719.

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

Englmaier, F., M. Fahn, and M. Schwarz (2018): "Long-Term Employment Relations When Agents are Present Biased," Working Paper DP13227, CEPR.

Fahn, M. (2017): "Minimum Wages and Relational Contracts," *Journal of Law, Economics, and Organization*, 33(2), 301–331.

Fahn, M., and H. Hakenes (2019): "Teamwork as a Self-Disciplining Device," *American Economic Journal: Microeconomics*, 11(4), 1–32.

Fahn, M., and R. Seibel (2020): "Inconsistent Time Preferences and On-the-job Search When it Pays to be Naive," Working Paper 7604, CESifo.

Fong, Y.-f., and J. Li (2017): "Relational Contracts, Limited Liability, and Employment Dynamics," *Journal of Economic Theory*, 169, 270–293.

Frakes, M. D., and M. F. Wasserman (2016): "Procrastination in the Workplace: Evidence from the U.S. Patent Office," Working Paper 22987, NBER.

Galperti, S. (2015): "Commitment, Flexibility, and Optimal Screening of Time Inconsistency," *Econometrica*, 83(4), 1425–1465.

Gervais, S., J. Heaton, and T. Odean (2011): "Overconfidence, Compensation Contracts, and Capital Budgeting," *Journal of Finance*, 66(5), 1735–1777.

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

Haltiwanger, J. C., J. I. Lane, and J. R. Spletzer (1999): "Productivity Differences across Employers: The Roles of Employer Size, Age, and Human Capital," *The American Economic Review*, 89(2), 94–98.

Harris, C., and D. Laibson (2001): "Dynamic Choices of Hyperbolic Consumers," *Econometrica*, 69(4), 935–957.

Heidhues, P., and B. Kőszegi (2010): "Exploiting Naïvete about Self-Control in the Credit Market," *American Economic Review*, 100(5), 2279–2303.

Hoffman, M., and S. Burks (2015): "Training Contracts, Worker Overconfidence, and the Return to Firm-Sponsored General Training," *Working Paper*.

Ikeda, S., M.-I. Kang, and F. Ohtake (2010): "Hyperbolic discounting, the sign effect, and the body mass index," *Journal of Health Economics*, 29(2), 268 – 284.

Kang, M.-I., and S. Ikeda (2016a): "Time Discounting and Smoking Behavior: Evidence from a Panel Survey," in *Behavioral Economics of Preferences, Choices, and Happiness*, ed. by S. Ikeda, H. K. Kato, F. Ohtake, and Y. Tsutsui, pp. 197–226. Springer Japan, Tokyo.

——— (2016b): "Time discounting, present biases, and health-related behaviors: Evidence from Japan," *Economics & Human Biology*, 21, 122 – 136.

Kaur, S., M. Kremer, and S. Mullainathan (2010): "Self -Control and the Development of Work Arrangements," *American Economic Review Papers and Proceedings*, 100(2), 624–628.

——— (2015): "Self-Control at Work," *Journal of Political Economy*, 123(6), 1227–1277.

Krusell, P., and A. A. Smith, Jr. (2003): "Consumption–Savings Decisions with Quasi–Geometric Discounting," *Econometrica*, 71(1), 365–375.

Laibson, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112(2), 443–477.

Lallemand, T., and F. Rycx (2009): "Are Young and Old Workers Harmful for Firm Productivity?," IZA Discussion Papers 3938, Institute of Labor Economics (IZA).

Leonardi, M., and G. Pica (2013): "Who Pays for it? The Heterogeneous Wage Effects of Employment Protection Legislation," *The Economic Journal*, 123(573), 1236–1278.

Li, S., J. Yan, and B. Xiao (2014): “Contract Design and Self-Control with Asymmetric Information,” *Economic Inquiry*, 52(2), 618–624.

Li, S., J. Yan, and J. Yu (2012): “When Holmström and Milgrom meet Laibson: Moral Hazard with Short-run Impatience and Imperfect Self–Awareness,” *Working Paper.*

Miller, D. A., and J. Watson (2013): “A Theory of Disagreement in Repeated Games with Bargaining,” *Econometrica*, 81(6), 2303–2350.

O’Donoghue, T., and M. Rabin (1999a): “Doing it Now or Later,” *American Economic Review*, 89(1), 103–124.

O’Donoghue, T., and M. Rabin (1999b): “Incentives for Procrastinators,” *The Quarterly Journal of Economics*, 114(3), 769–816.

O’Donoghue, T., and M. Rabin (2001): “Choice and Procrastination,” *Quarterly Journal of Economics*, 116(1), 121–160.

Phelps, E. S., and R. A. Pollak (1968): “On Second-best National Saving and Game-Equilibrium Growth,” *Review of Economic Studies*, 35(2), 185–199.

Rogerson, W. P. (1985): “Repeated Moral Hazard,” *Econometrica*, 53(1), 69–76.

Shui, H., and L. M. Ausubel (2005): “Time inconsistency in the credit card market,” *Working Paper.*

Skirbekk, V. (2004): “Age and Individual Productivity: A Literature Survey,” *Vienna Yearbook of Population Research*, 2, 133–153.

Spear, S., and S. Srivastava (1987): “On Repeated Moral Hazard with Discounting,” *Review of Economic Studies*, 54(4), 599–617.

Strotz, R. H. (1955): “Myopia and Inconsistency in Dynamic Utility Maximization,” *Review of Economic Studies*, 23(3), 165–180.

Van Reenen, J. (1996): “The Creation and Capture of Rents: Wages and Innovation in a Panel of UK Companies,” *The Quarterly Journal of Economics*, 111(1), 195–226.

Waldman, M. (2012): *Theory and Evidence in Internal Labor Markets*pp. 520–572. Princeton University Press.

Weinschenk, P. (2019): “On the Benefits of Time-Inconsistent Preferences,” Working paper, SSRN.

Yan, J. (2011): "Contracting with Heterogeneous Quasi-hyperbolic Agents: Investment Goods Pricing under Asymmetric Information," *Working Paper.*

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

A Figures

Figure 1: Utility as Function of β

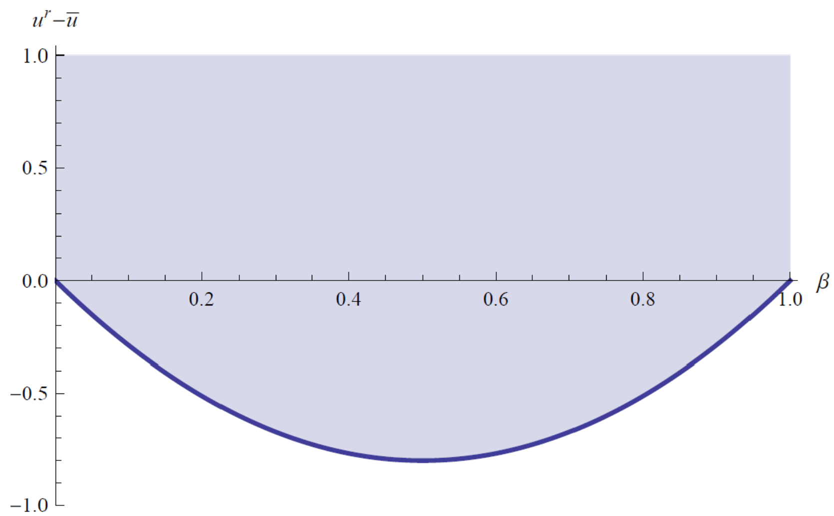

Figure 2: Profit as Function of β

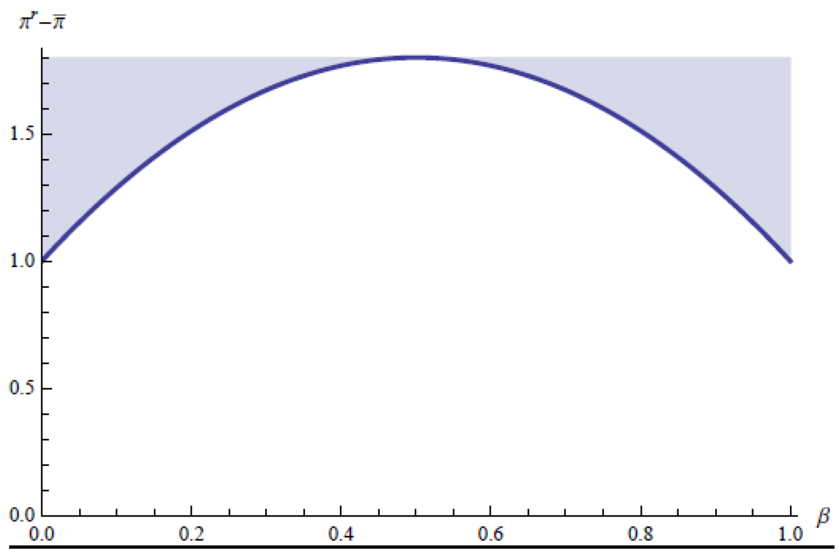

Figure 3: Effort Path in Virtual Contract

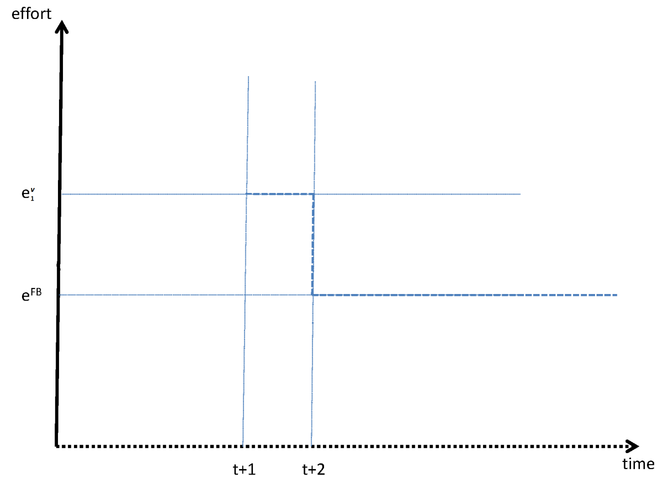

Figure 4: Utility Path in Virtual Contract

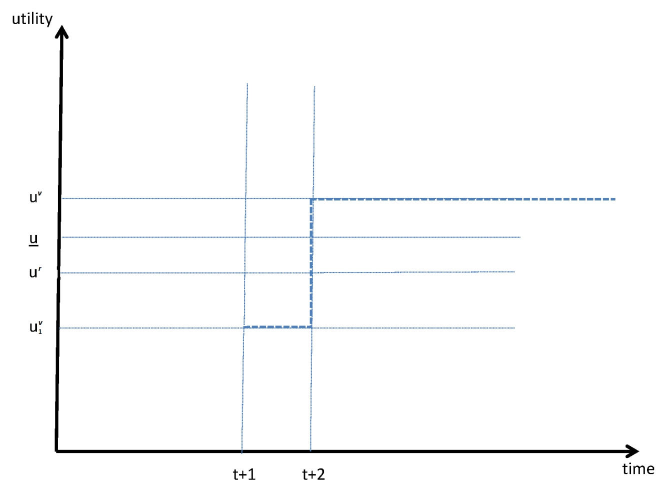

B Preliminaries: Simplifying the Problem

First, we derive a number of preliminary results that help us to prove our main results.

Lemma B1. *It is without loss of generality to set* $w^r = w_1^v = w^v = 0$.

Proof. Assume $w^v > 0$. Reducing w^v by ε and increasing b^v by $\varepsilon > 0$ does not tighten any constraint, but relaxes (vIC). Assume $w^r > 0$. Reducing w^r by $\varepsilon > 0$ and increasing b^r by ε does not tighten any constraint but relaxes (rIC). An equivalent argument can be applied to show that $w_1^v = 0$. □

Hence, the constraints (rIRA) and (rIC), (v1IRA) and (v1IC), and (vIRA) and (vIC), respectively, are identical, allowing us to omit (rIC), (v1IC) and (vIC).

Therefore, the remaining constraints are

$$b^r - c(e^r) - \overline{u} + \beta\delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u})\right] \geq 0 \qquad \text{(rIRA)}$$

$$(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u}) \geq 0 \qquad \text{(v1IRA)}$$

$$(b^v - c(e^v) - \overline{u}) \geq 0 \qquad \text{(vIRA)}$$

$$e^r\theta - b^r + \delta\left[(e_1^v\theta - b_1^v) + \frac{\delta}{1-\delta}(e^v\theta - b^v)\right] \geq \frac{\overline{\pi}}{1-\delta} \qquad \text{(vrIRP)}$$

$$(e_1^v\theta - b_1^v) + \frac{\delta}{1-\delta}(e^v\theta - b^v) \geq \frac{\overline{\pi}}{1-\delta} \qquad \text{(v1IRP)}$$

$$e^v\theta - b^v \geq \overline{\pi} \qquad \text{(vIRP)}$$

$$-c(e^r) + b^r - \overline{u} + \beta\delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u})\right]$$
$$\geq (b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta\beta}{1-\delta}(b^v - c(e^v) - \overline{u}) \qquad \text{(rC)}$$

$$(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{(1-\delta)}(b^v - c(e^v) - \overline{u})$$
$$\geq (b^r - c(e^r) - \overline{u}) + \delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{(1-\delta)}(b^v - c(e^v) - \overline{u})\right] \qquad \text{(vC)}$$

Note that we added $\overline{u}\left(1 + \frac{\beta\delta}{1-\delta}\right)$ on both sides of (rC) and $\frac{\overline{u}}{1-\delta}$ on both sides of (vC).

In a next step, we prove the principal can push the agent below his outside option:

Lemma B2. *If the agent is naive and has $\beta \in (0,1)$, then in the profit-maximizing menu of contracts, $-c(e^r) + b^r < \overline{u}$.*

Proof. First, assume $-c(e^r)+b^r > \overline{u}$. Then, (vC) and (v1IRA) imply that $(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{(1-\delta)}(b^v - c(e^v) - \overline{u})$ must be strictly positive as well. Change $\boldsymbol{C}$ in the following way: Set $b^r - c(e^r) - \overline{u} = b_1^v - c(e_1^v) - \overline{u} = (b^v - c(e^v) - \overline{u}) = 0$, which leaves all constraints satisfied and increases the principal's profits. Note that these considerations already allow us to omit (vrIRP), (v1IRA) and (vC).

Now, assume that $-c(e^r)+b^r = \overline{u}$. Change $\boldsymbol{C}$ in the following way: First, set $(b_1^v - c(e_1^v) - \overline{u}) = (b^v - c(e^v) - \overline{u}) = 0$ for given effort levels, which satisfies all constraints. Then, reduce b_1^v by ε and increase b^v by $\varepsilon\frac{1-\delta}{\delta\beta}$. This increases $b_1^v + \frac{\delta}{1-\delta}b^v$ and hence relaxes (rIRA), (v1IRA), (vIRA) and (rC) (and does not violate limited liability as well as (v1IRP) and (vIRP) constraints for ε sufficiently small), and therefore allows the principal to reduce b^r. □

$-c(e^r) + b^r < \overline{u}$ immediately implies that, given (rIRA), the (v1IRA) constraint automatically holds and can be omitted. The same is true for (vC). Furthermore, the (vrIRP) constraint can be omitted: $e^r\theta - b^r - \overline{\pi}$ will be strictly positive in a profit-maximizing equilibrium, and (v1IRP) yields $(e_1^v\theta - b_1^v) + \frac{\delta}{1-\delta}(e^v\theta - b^v) \geq \frac{\overline{\pi}}{1-\delta}$.

$$b^r - c(e^r) - \overline{u} + \beta\delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u})\right] \geq 0 \qquad \text{(rIRA)}$$

$$(b^v - c(e^v) - \overline{u}) \geq 0 \qquad \text{(vIRA)}$$

$$(e_1^v\theta - b_1^v) + \frac{\delta}{1-\delta}(e^v\theta - b^v) \geq \frac{\overline{\pi}}{1-\delta} \qquad \text{(v1IRP)}$$

$$e^v\theta - b^v \geq \overline{\pi} \qquad \text{(vIRP)}$$

$$-c(e^r)+b^r-\overline{u}+\beta\delta\left[(b_1^v-c(e_1^v)-\overline{u})+\frac{\delta}{1-\delta}\left(b^v-c(e^v)-\overline{u}\right)\right]$$
$$\geq (b_1^v-c(e_1^v)-\overline{u})+\frac{\delta\beta}{1-\delta}\left(b^v-c(e^v)-\overline{u}\right) \qquad \text{(rC)}$$

Lemma B3. *(vIRA) and (v1IRP) constraints are slack and can hence be omitted in a profit-maximizing equilibrium.*

Proof. First, assume that (vIRA) binds. Increasing b^v by ε and reducing $e_1^v\theta-b_1^v$ by $\frac{\delta}{1-\delta}\varepsilon$ keeps all constraints unaffected with the exception of (vIRA) and (rC) which are relaxed.

Concerning (v1IRP), first note that if it binds, the same has to be true for (vIRP). Otherwise, we could reduce b_1^v by ε (if (vIRP) is slack but (v1IRP) binds, $b_1^v>0$ for sure) and increase b^v by $\varepsilon\frac{1-\delta}{\delta}$. This would not affect (v1IRP) and (rIR), but relax (rC).

Now, assume that (v1IRP) and (vIRP) bind. This implies that $b_1^v>0$, that the agent gets the whole virtual surplus, and constraints are:

$$b^r-c(e^r)-\overline{u}+\beta\delta\left[(e_1^v\theta-c(e_1^v)-\overline{\pi}-\overline{u})+\frac{\delta}{1-\delta}\left(e^v\theta-c(e^v)-\overline{\pi}-\overline{u}\right)\right]\geq 0 \qquad \text{(rIR)}$$

$$-c(e^r)+b^r-\overline{u}+\beta\delta\left[(e_1^v\theta-c(e_1^v)-\overline{\pi}-\overline{u})+\frac{\delta}{1-\delta}\left(e^v\theta-c(e^v)-\overline{\pi}-\overline{u}\right)\right]$$
$$\geq (e_1^v\theta-c(e_1^v)-\overline{\pi}-\overline{u})+\frac{\delta\beta}{1-\delta}\left(e^v\theta-c(e^v)-\overline{\pi}-\overline{u}\right) \qquad \text{(rC)}$$

There, the right hand side of (rC) is positive, hence that (rIR) is slack. Therefore, a slight reduction of b_1^v, accompanied with a reduction of b^r or an increase of e^r (to keep (rC) unaffected) would keep (rIR) satisfied and increase the principal's profits. □

C Proofs to Lemmas and Propositions from the Main Part

Proof to Lemma 1 and Proposition 1

Proof. Given the remaining constraints after simplifying the problem in Section B, i.e., constraints (rIR), (rC), and (vIRP), the principal's maximization problem gives rise to the following Lagrange function:

$$L = \frac{e^r\theta - b^r}{1-\delta} + \lambda_{rIR}\left[b^r - c(e^r) - \overline{u} + \beta\delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u})\right]\right]$$
$$+ \lambda_{rC}\left[-c(e^r) + b^r - \overline{u} - (b_1^v - c(e_1^v) - \overline{u})(1-\beta\delta) - \delta\beta(b^v - c(e^v) - \overline{u})\right]$$
$$+ \lambda_{vIRP}(e^v\theta - b^v - \overline{\pi}),$$

with first-order conditions

$$\frac{\partial L}{\partial b^r} = \frac{-1}{1-\delta} + \lambda_{rIR} + \lambda_{rC} = 0$$
$$\frac{\partial L}{\partial e^r} = \frac{\theta}{1-\delta} - c(e^r)'(\lambda_{rIR} + \lambda_{rC}) = 0$$
$$\frac{\partial L}{\partial (b_1^v - c(e_1^v))} = \lambda_{rIR}\beta\delta - \lambda_{rC}(1-\beta\delta) = 0$$
$$\frac{\partial L}{\partial b^v} = \lambda_{rIR}\beta\delta\frac{\delta}{1-\delta} - \lambda_{rC}\delta\beta - \lambda_{vIRP} = 0$$
$$\frac{\partial L}{\partial e^v} = -\lambda_{rIR}\beta\delta\frac{\delta}{1-\delta}c(e^v)' + \lambda_{rC}\delta\beta c(e^v)' + \lambda_{vIRP}\theta = 0.$$

Hence, $\lambda_{rIR} = \frac{1}{1-\delta} - \lambda_{rC}$ and $\theta - c(e^r)' = 0$, i.e., $e^r = e^{FB}$. Rearranging these conditions further yields $\lambda_{rC} = \lambda_{rIR}\frac{\beta\delta}{(1-\beta\delta)}$, $\lambda_{vIRP} = \lambda_{rIR}\frac{\beta\delta^2(1-\beta)}{(1-\delta)(1-\beta\delta)}$, and $\frac{\beta\delta^2(1-\beta)}{(1-\delta)(1-\beta\delta)}(\theta - c(e^v)') = 0$, i.e., $e^v = e^{FB}$.

Hence, (rC), (rIR) and (vIRP) can only bind simultaneously, which also implies that all of them must bind. To the contrary, assume that (rC) and (rIR) do not bind. Then, b^r can be further reduced until one of them binds, further increasing the principal's profits. Using these results gives the values for b^r, u^r and π^r. □

Proof to Proposition 2

Proof. We solve this by backward induction: In the last period, there are no future periods left, so the agent is just compensated for first-best effort: $e_T^r = e^{FB}$, $b_T^r = c(e^{FB})$. In the second to last period, the agent cannot be fooled regarding the last period as he will end up choosing the contract giving him the highest utility. So again, the agent is just compensated for first-best effort: $e_{T-1}^r = e^{FB}$, $b_{T-1}^r = c(e^{FB})$. In $T-2$, the agent can be fooled by offering him a virtual contract in which he earns the whole production in the last period if he has earned less (or has worked harder) in the second to last period: $e_T^v = e_{T-1}^v = e^{FB}$, $b_{T-1}^v = c(e^{FB}) + \overline{u} - \beta\delta\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$, $b_T^v = e^{FB}\theta - \overline{\pi}$, where b_{T-1}^v makes the $T-1$ agent just indifferent between the virtual and the real contract. If the virtual contract above is promised to him, he is willing to work below his outside option in $T-2$: $e^r_{T-2} = e^{FB}$, $b^r_{T-2} = c(e^{FB}) + \overline{u} - \beta(1-\beta)\delta^2 \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$. Analogously, in earlier periods the promised virtual contract will always promise the whole surplus to the agent from the period after the entry period on, so the entry period bonus that makes the agent indifferent is given by $b^v_t = c(e^{FB}) + \overline{u} - \beta\delta \sum_{j=0}^{T-t-1} \delta^j \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$. This in turn allows the principal to exploit the agent by setting $b^r_t = c(e^r_t) + \overline{u} - \beta(1-\beta)\delta^2 \sum_{j=0}^{T-t-2} \delta^j \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$. The expression for π^r_t follows by plugging in b^r_t and $c(e^r_t)$ □

Proof to Corollary 2

Proof. This follows immediately from Proposition 2. □

Proof to Proposition 3

Proof. Note that none of the steps we performed to simplify the original problem is affected by the presence of a limited liability constraint. Hence, relevant constraints are the same, and the Lagrange function becomes:

$$\begin{aligned} L =& \frac{e^r\theta - b^r}{1-\delta} + \lambda_{rIR}\left[b^r - c(e^r) - \overline{u} + \beta\delta\left[(b^v_1 - c(e^v_1) - \overline{u}) + \frac{\delta}{1-\delta}\left(b^v - c(e^v) - \overline{u}\right)\right]\right] \\ &+ \lambda_{rC}\left[-c(e^r) + b^r - \overline{u} - \left(b^v_1 - c(e^v_1) - \overline{u}\right)(1-\beta\delta) - \delta\beta\left(b^v - c(e^v) - \overline{u}\right)\right] + \lambda_{b^r} b^r \\ &+ \lambda_{vIRP}\left(e^v\theta - b^v - \overline{\pi}\right), \end{aligned}$$

with first-order conditions:

$$\begin{aligned} \frac{\partial L}{\partial b^r} &= \frac{-1}{1-\delta} + \lambda_{rIR} + \lambda_{rC} + \lambda_{b^r} = 0 \\ \frac{\partial L}{\partial e^r} &= \frac{\theta}{1-\delta} - c(e^r)'\left(\lambda_{rIR} + \lambda_{rC}\right) = 0 \\ \frac{\partial L}{\partial\left(b^v_1 - c(e^v_1)\right)} &= \lambda_{rIR}\beta\delta - \lambda_{rC}\left(1-\beta\delta\right) = 0 \\ \frac{\partial L}{\partial b^v} &= \lambda_{rIR}\beta\delta\frac{\delta}{1-\delta} - \lambda_{rC}\delta\beta - \lambda_{vIRP} = 0 \\ \frac{\partial L}{\partial e^v} &= -\lambda_{rIR}\beta\delta\frac{\delta}{1-\delta}c(e^v)' + \lambda_{rC}\delta\beta c(e^v)' + \lambda_{vIRP}\theta = 0 \end{aligned}$$

Hence, $\lambda_{rIR} = \frac{1}{1-\delta} - \lambda_{rC} - \lambda_{b^r}$ and $\frac{\theta - c(e^r)'}{1-\delta} + c(e^r)'\lambda_{b^r} = 0$. Rearranging further yields $\lambda_{rC} = \lambda_{rIR}\frac{\beta\delta}{(1-\beta\delta)}$, $\lambda_{vIRP} = \lambda_{rIR}\frac{\beta\delta^2(1-\beta)}{(1-\delta)(1-\beta\delta)}$, and $\frac{\beta\delta^2(1-\beta)}{(1-\delta)(1-\beta\delta)}\left(\theta - c(e^v)'\right) = 0$. Furthermore, (rC), (rIR) and (vIRP) all bind simultaneously, which follows from the same arguments as in the prove to Proposition 1.

Hence, if $b^r = 0$, e^r is above the level given by $\theta - c(e^r)' = 0$. Plugging binding constraints into utilities gives the desired values. □

Proof to Corollary 3

Proof. This follows immediately by Proposition 3 and application of the Implicit Function Theorem with $c(e^{FB}) + \overline{u} - \beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) = 0$ and $c(e^r) + \overline{u} - \beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) = 0$. □

Proof to Corollary 4

Proof. This proof is analogous to the proof to Proposition 2 as long as the limited liability does not bind. When the limited liability constraint becomes binding (for a smaller t, as we are doing a backward induction), the principal cannot further decrease b_t^r, but the only way for her to (then imperfectly) transfer from the agent to her is increasing e_t^r. Analogous to Corollary 2, this implies the claim. □

Proof to Proposition 4

Proof. For simplicity, we still assume that once the virtual contract has been selected, the agent cannot revert to the real contract in later periods. Then, the principal's objective is to maximize $e^r\theta - b^r$ in every period, subject to the following constraints.

$$b^r - c(e^r) - \overline{u} + \delta\left(b_1^v - kc(e_1^v) - \overline{u}\right) + \frac{\delta^2}{1-\delta}\left(b^v - kc(e^v) - \overline{u}\right) \geq 0 \tag{rIRA}$$

$$\left(b_1^v - kc(e_1^v) - \overline{u}\right) + \frac{\delta}{1-\delta}\left(b^v - kc(e^v) - \overline{u}\right) \geq 0 \tag{v1IRA}$$

$$\left(b^v - kc(e^v) - \overline{u}\right) \geq 0 \tag{vIRA}$$

$$\left(e_1^v\theta - b_1^v - \overline{\pi}\right) + \frac{\delta}{1-\delta}\left(e^v\theta - b^v - \overline{\pi}\right) \geq 0 \tag{v1IRP}$$

$$e^v\theta - b^v - \overline{\pi} \geq 0 \tag{vIRP}$$

$$b^r - c(e^r) - (b_1^v - c(e_1^v)) + \delta\left[(b_1^v - kc(e_1^v)) - (b^v - kc(e^v))\right] \geq 0 \tag{rC}$$
$$b_1^v - kc(e_1^v) - (b^r - c(e^r)) + \delta\left[(b^v - kc(e^v)) - (b_1^v - kc(e_1^v))\right] \geq 0 \tag{vC}$$

For later use, note that, if rC binds, plugging it into vC yields $(1-k)\,c(e_1^v) \geq 0$. Hence, vC is automatically satisfied if rC binds.

Now, v1IRA can be omitted because the principal cannot do worse than setting $b^r = c(e^r) + \overline{u}$. Thus, $b^r \leq c(e^r) + \overline{u}$, and v1IRA is implied by rIRA. Furthermore, it is without loss of generality to set $b^v = kc(e^v) + \overline{u}$. With $b^v > kc(e^v) + \overline{u}$, reducing b^v by a small ε (such that vIRA still holds) and increasing b_1^v by $\varepsilon\frac{\delta}{1-\delta}$ keeps all constraints (and the objective) unaffected, besides vIRP which is relaxed. Then, it is optimal to have e^v at its first-best level, hence it is characterized by $\theta - kc' = 0$. This maximizes the left-hand sides of v1IRP and vIRP, whereas the size of e^v is immaterial for the objective as well as all other constraints (taking $b^v = kc(e^v) + \overline{u}$ into account). Because $e^v\theta - kc(e^v) - \overline{u} - \overline{\pi} > 0$, vIRA and vIRP can in the following be omitted.

This yields the following Lagrange function

$$\begin{aligned}
L = & e^r\theta - b^r + \lambda_{rIRA}\left[b^r - c(e^r) - \overline{u} + \delta\left(b_1^v - kc(e_1^v) - \overline{u}\right)\right] \\
& + \lambda_{v1IRP}\left[e_1^v\theta - b_1^v - \overline{\pi} + \frac{\delta}{1-\delta}\left(e^v\theta - kc(e^v) - \overline{u} - \overline{\pi}\right)\right] \\
& + \lambda_{rC}\left[b^r - c(e^r) - (b_1^v - c(e_1^v)) + \delta\left(b_1^v - kc(e_1^v)\right) - \delta\overline{u}\right] \\
& + \lambda_{vC}\left[b_1^v - kc(e_1^v) - (b^r - c(e^r)) + \delta\left[\overline{u} - (b_1^v - kc(e_1^v))\right]\right]
\end{aligned}$$

$$\begin{aligned}
\frac{\partial L}{\partial e^r} = & \theta - \lambda_{rIRA}c'(e^r) - \lambda_{rC}c'(e^r) + \lambda_{vC}c'(e^r) = 0 \\
\frac{\partial L}{\partial b^r} = & -1 + \lambda_{rIRA} + \lambda_{rC} - \lambda_{vC} = 0 \\
\Rightarrow & \lambda_{rIRA} = 1 - \lambda_{rC} + \lambda_{vC} \\
\Rightarrow & \theta - c'(e^r) = 0
\end{aligned}$$

$$
\begin{aligned}
\frac{\partial L}{\partial e_1^v} =& - \lambda_{rIRA}\delta k c'(e_1^v) + \lambda_{v1IRP}\theta + \lambda_{rC} c'(e_1^v)\left(1-\delta k\right) - \lambda_{vC} k c'(e_1^v)\left(1-\delta\right) = 0 \\
\Rightarrow & c'(e_1^v)\left[\lambda_{rC} - \lambda_{vC} k - \delta k\right] + \lambda_{v1IRP}\theta = 0 \\
\frac{\partial L}{\partial b_1^v} =& \lambda_{rIRA}\delta - \lambda_{v1IRP} - \lambda_{rC}\left(1-\delta\right) + \lambda_{vC}\left(1-\delta\right) = 0 \\
\Rightarrow & \lambda_{rC} = \delta + \lambda_{vC} - \lambda_{v1IRP} \\
\Rightarrow & \lambda_{v1IRP} = -c'(e_1^v)\frac{\left(1-k\right)\left(\delta+\lambda_{vC}\right)}{\left(\theta - c'(e_1^v)\right)} > 0 \\
\Rightarrow & \lambda_{rC} = \left(\delta+\lambda_{vC}\right)\frac{\theta - kc'(e_1^v)}{\left(\theta - c'(e_1^v)\right)} \Rightarrow \theta - kc'(e_1^v) \leq 0 \\
\Rightarrow & \lambda_{rIRA} = 1 - \delta + \lambda_{v1IRP} > 0
\end{aligned}
$$

Hence, v1IRP and rIRA bind in a profit-maximizing equilibrium. In the following, we further distinguish between the two cases $\theta - \delta kc'(e_1^v) < 0$ (and $\lambda_{rC} > 0$) and $\theta - \delta kc'(e_1^v) = 0$ (and $\lambda_{rC} = 0$) and explore the conditions for each to hold.

(I) $\theta - \delta kc'(e_1^v) < 0$**; hence,** $e_1^v > e^v$**:** Then, $\lambda_{rC} > 0$, which implies $\lambda_{vC} = 0$ (rC and vC cannot bind simultaneously), and outcomes can be obtained by having (rIRA), (rC) and (v1IRP) constraints hold as equalities.

First, a binding (v1IRP) constraints yields $b_1^v = e_1^v\theta - \overline{\pi} + \frac{\delta}{1-\delta}\left(e^v\theta - kc(e^v) - \overline{u} - \overline{\pi}\right)$. Plugging this into (rC) delivers

$$
b^r - c(e^r) = \overline{u} - \delta c(e_1^v)\left(1-k\right)
$$

and

$$
b_1^v = e_1^v\theta - \overline{\pi} + \frac{\delta}{1-\delta}\left(e^v\theta - kc(e^v) - \overline{u} - \overline{\pi}\right),
$$

whereas plugging these values into (rIRA) gives $e_1^v\theta - c(e_1^v) - \overline{u} - \overline{\pi} + \frac{\delta}{1-\delta}\left(e^v\theta - kc(e^v) - \overline{u} - \overline{\pi}\right) = 0$, which determines e_1^v.

(II) $\theta - kc'(e_1^v) = 0$**; hence,** $e_1^v = e^v$**:** Then, (rC) is slack, whereas (rIRA) and (v1IRP) still bind. Taking into account $e_1^v = e^v$, the two constraints yield

$$
b^r - c(e^r) = \overline{u} - \frac{\delta}{1-\delta}\left(e^v\theta - kc(e^v) - \overline{u} - \overline{\pi}\right)
$$

and

$$
b_1^v = e_1^v\theta - \overline{\pi} + \frac{\delta}{1-\delta}\left(e^v\theta - kc(e^v) - \overline{u} - \overline{\pi}\right).
$$

To assess whether this case is feasible, we plug the derived values into (rC), which becomes

$$-\left(e^{v}\theta - kc(e^{v}) - \overline{u} - \overline{\pi}\right)\frac{1}{1-\delta} + c(e^{v})\left(1-k\right) \geq 0.$$

If this condition holds, results are characterized by case (II). Otherwise, case (I) determines outcomes.

Finally, note vC is also slack in case (II); a binding vC then would yield $\left(e^{v}\theta - kc(e^{v}) - \overline{u} - \overline{\pi}\right)\frac{1}{1-\delta} = 0$, which does not hold for of e^{v}. □

C.1 Labor Market Competition and Bargaining

Now, the (net) surplus in a given period t amounts to $e_t\theta - c(e_t) - \overline{u} - \overline{\pi}$. In order to characterize payments, we also have to specify what happens if bargaining fails. We assume that in this case, players do not enter the employment relationship and consume their respective outside options. Then, already taking into account that the agent expects to select the virtual contract in future periods, the agent is willing to accept any offer that gives him at least an equivalent to

$$\begin{aligned}&\overline{u} + (1-\alpha)\left(e^{r}\theta - c(e^{r}) - \overline{u} - \overline{\pi}\right)\\ &+\beta\delta\left\{\overline{u} + (1-\alpha)\left(e_1^{v}\theta - c(e_1^{v}) - \overline{u} - \overline{\pi}\right) + \frac{\delta}{(1-\delta)}\left[\overline{u} + (1-\alpha)\left(e^{v}\theta - c(e^{v}) - \overline{u} - \overline{\pi}\right)\right]\right\}.\end{aligned}$$

As before, the principal will optimally offer a contract menu that shifts a major part of the compensation into the virtual contract. Therefore, it remains optimal to promise the agent the full virtual surplus (from the second period of the virtual contract on), and consequently reduce real payments. Furthermore, $e^{r} = e^{v} = e^{FB}$. There, note that if the agent deviates from equilibrium effort, he is subsequently still pushed to his outside option, and (IC) constraints are "slacker" than in our main model with $\alpha = 1$. Hence, the principal's commitment allows her to rule out any attempt by the agent to renegotiate the agreement. Note that our results would be unaffected if we allowed to agent to obtain a share $1-\alpha$ of the surplus also after a deviation, for example because of alternative job opportunities.

In addition, the agent's utility in the first period of the virtual contract, u_1^{v}, must be sufficiently low for the agent to not go for the virtual contract. There, we set $e_1^{v} = e^{FB}$, which is without loss of generality because what matters in the first period of the virtual contract is the value u_1^v, not the exact specifications of e_1^v and w_1^v.

As before, the relevant constraints which pin down payments and utilites are (rIRA) and (rC):

$$\begin{aligned} & b^r - c(e^{FB}) + \beta\delta\left(b_1^v - c(e^{FB}) + \delta\frac{b^v - c(e^{FB})}{1-\delta}\right) \\ & \geq \overline{u} + (1-\alpha)\left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right) + \beta\delta\frac{\left[\overline{u} + (1-\alpha)\left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\right]}{1-\delta} \end{aligned} \tag{rIRA}$$

and

$$\begin{aligned} & b^r - c(e^{FB}) + \beta\delta\left(b_1^v - c(e^{FB}) + \delta\frac{b^v - c(e^{FB})}{1-\delta}\right) \\ & \geq b_1^v - c(e^{FB}) + \frac{\delta\beta}{1-\delta}\left(b^v - c(e^{FB})\right). \end{aligned} \tag{rC}$$

Both constraints will bind for the same reasons as in the main part. Thus, (rIRA) and (rC) pin down the values of the agent's real utility and his utility in the first period of the virtual contract:

$$u^r - \overline{u} = \left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\left((1-\alpha) - \frac{\alpha(1-\beta)\beta\delta^2}{1-\delta}\right)$$

and

$$u_1^v - \overline{u} = \left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\left((1-\alpha) - \frac{\alpha\beta\delta}{1-\delta}\right).$$

C.2 Moral Hazard

In this variant of the model, we show that the results of our basic model do not depend on the assumption that the agent's effort is contractible, but also hold under moral hazard. In this case, the naive agent is still mainly incentivized by rents provided by the virtual contract. More precisely, consider the following adjustment to our model setting: Effort e is the agent's private information. Output remains verifiable, but now amounts to $y \in \{0, \theta\}$, with $\text{Prob}(y = \theta) = e$. Hence, output is not deterministic anymore but can be either high or low. Effort determines the probability with which high output is realized. Moreover, in order to guarantee an interior solution we assume $\lim_{e\to 1} c' = \infty$. This implies that first-best effort e^{FB} is still characterized by $\theta - c'(e^{FB}) = 0$. For concreteness we also set $\overline{u} = \overline{\pi} = 0$.

First, note that the optimal contract for a sophisticated agent is the same as for an agent without a present bias: In every period, the agent receives the bonus if output has been high. This bonus is set such that the agent chooses first-best effort, i.e., $b = \theta$. Furthermore, the wage is used to extract the generated rent, hence $w = \overline{u} + c(e^{FB}) - e^{FB}\theta$. Then, the principal collects the full surplus, and the optimal contract is stationary. This is different from the repeated moral hazard settings in Rogerson (1985) or Spear and Srivastava (1987), where the optimal contract exhibits memory. This is driven by the agent's risk aversion, though, which does not make it optimal to effectively sell the firm to the agent in every period (as is the case with a risk neutral agent).[16]

The structure of the optimal contract menu for the naive agent is the same as without moral hazard, with the exception that the bonus is only paid if output has been high (note that because we can freely choose w^r, it is without loss of generality to set the bonus after a low output realization to zero).

Therefore, the naive agent's expected utility in a given period amounts to

$U^{rv} = w^r - c(e^r) + e^r b^r + \beta\delta U_1^v$.

Furthermore, the agent's virtual payoffs amount to

$U_1^v = w_1^v - c(e_1^v) + e_1^v b_1^v + \delta U^v$ and $U^v = \left(w^v - c(e^v) + e^v b^v\right)/(1-\delta)$.

Now, the principal must use wage payments in the real and first period of the virtual contract to finetune the arrangement, that is, extract real rents and in the end make the agent *not* go for the virtual contract. Without moral hazard, the principal could also set e_1^v in a way to prevent the agent from choosing the virtual contract. Now, effort is automatically pinned down by bonus payments and future rents. More precisely, the agent's real effort is given by

$$e^r \in \operatorname{argmax}\left[e^r b^r - c(e^r)\right]. \qquad \text{(rIC)}$$

[16] Even with a risk neutral agent the optimal contract might contain memory, namely if the agent were protected by limited liability. Then, a profit-maximizing contract would provide incentives via a combination of bonus payments and on-path termination threats. See Fong and Li (2017), who derive a profit-maximizing dynamic contract in case output is not verifiable.

Therefore, e^r is characterized by $b^r - c'(e^r) = 0$. Virtual effort levels are characterized by $b_1^v - c'(e_1^v) = 0$ and $b^v - c'(e^v) = 0$, respectively.

It is straightforward to show that the agent still receives the full virtual rent from the second virtual period on, and that it remains optimal to maximize this rent. Hence, $e^v = e^{FB}$ (that is, $b^v = c'(e^{FB})$, and w^v is set such that $u^v = e^{FB}\theta - c(e^{FB}) - \overline{\pi}$). Furthermore, b^r is set such that $e^r = e^{FB}$ as well (since we do not impose a limited liability constraint in this section, the bonus b^r can potentially be negative). The first virtual period is – as before – designed in a way that the agent actually does not select the virtual but sticks to the real contract.

As without moral hazard, it is optimal to have (rIRA) and (rC) constraints hold as equalities (furthermore, the respective solution will satisfy the (vC) constraint).

These constraints are $U^{rv} \geq \overline{u} + \beta\delta\frac{\overline{u}}{1-\delta}$ (rIRA) and $U^{rv} \geq u_1^v + \beta\delta U^v$ (rC), where $u_1^v = w_1^v + e_1^v b_1^v - c(e_1^v)$, and can be rewritten as

$$w^r - c(e^{FB}) + e^{FB}b^r + \beta\delta U_1^v \geq \overline{u} + \beta\delta\frac{\overline{u}}{1-\delta} \qquad \text{(rIRA)}$$

and

$$w^r - c(e^{FB}) + e^{FB}b^r + \beta\delta U_1^v \geq u_1^v + \beta\delta U^v. \qquad \text{(rC)}$$

Having both constraints bind yields
$u^r = \overline{u} - \delta^2\beta(1-\beta)\frac{\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)}{1-\delta}$ and $u_1^v = \overline{u} - \beta\delta\frac{\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)}{1-\delta}$.

Concluding, the agent still is exploited in case of moral hazard, with (qualitatively) equivalent comparative statics. Furthermore, the first period of the virtual contract can be set such that the agent never selects it in the end.

C.3 Partial Naiveté and Learning

So far we only considered the extreme cases of completely naive and fully sophisticated agents. Here, we relax this assumption and show that a partially naive agent receives exactly the same contract as a fully naive agent. A partially naive agent thinks that in any future period, he will discount the future with a factor $\hat{\beta} \in (\beta, 1)$. A fully sophisticated agent has $\hat{\beta} = \beta$, whereas a completely naive agent has $\hat{\beta} = 1$.

The principal's maximization problem is very similar to the problem when she faces a completely naive agent. (v1IC), (vIC), and (vC) are changed as these constraints involve the agent's expectations about his future self:

$$-c(e_1^v) + b_1^v + \hat{\beta}\frac{\delta}{1-\delta}\left(w^v + b^v - c(e^v) - \overline{u}\right) \geq 0, \tag{v1IC}$$

$$-c(e^v) + b^v + \hat{\beta}\frac{\delta}{1-\delta}\left(w^v - c(e^v) + b^v - \overline{u}\right) \geq 0, \tag{vIC}$$

and

$$\begin{aligned}(w_1^v + b_1^v - c(e_1^v)) + \hat{\beta}\frac{\delta}{(1-\delta)}\left(w^v + b^v - c(e^v)\right)\\ \geq (w^r + b^r - c(e^r)) + \hat{\beta}\delta\left[(w_1^v + b_1^v - c(e_1^v)) + \frac{\delta}{(1-\delta)}\left(w^v + b^v - c(e^v)\right)\right].\end{aligned} \tag{vC}$$

The analysis is analogous to the one with the naive agent and by the same arguments we can omit several constraints. Thus we are left with the following simplified problem, maximizing

$$\Pi^r = \frac{e^r\theta - b^r}{1-\delta},$$

subject to

$$b^r - c(e^r) - \overline{u} + \beta\delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}\left(b^v - c(e^v) - \overline{u}\right)\right] \geq 0, \tag{rIR}$$

$$\begin{aligned}-c(e^r) + b^r - \overline{u} + \beta\delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}\left(b^v - c(e^v) - \overline{u}\right)\right]\\ \geq (b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta\beta}{1-\delta}\left(b^v - c(e^v) - \overline{u}\right),\end{aligned} \tag{rC}$$

$$(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{(1-\delta)}\left(b^v - c(e^v) - \overline{u}\right)\hat{\beta}\frac{(1-\delta)}{\left(1-\hat{\beta}\delta\right)} \geq \frac{(b^r - c(e^r) - \overline{u})}{\left(1-\hat{\beta}\delta\right)}, \tag{vC}$$

$$e^v\theta - b^v \geq \overline{\pi}. \tag{vIRP}$$

Proposition 5. *The partially naive agent receives the same contract as the fully naive agent.*

We show that (vC) never binds, which implies that the principal optimally does not differentiate between a fully and a partially naive agent. The reason is exactly the same as in the case with a fully naive agent: Unless $\beta = \hat{\beta}$ or the agent is always kept at his outside option, (vC) and (rC) cannot bind simultaneously.

Proof. We show that (vC) does not bind.
Ignore the non-negativity constraints. (rIRP) binds. If not, one could increase b^v by ε and decrease b_1^v by $\frac{\delta\hat{\beta}}{1-\delta\hat{\beta}}\epsilon$. Then one could decrease b^r slightly without violating any constraint and thereby increasing the principal's profits.
(rC) binds: Note that (vC) can only bind simultaneously if $\beta = \hat{\beta}$. Assume $\beta < \hat{\beta}$ and (rC) does not bind. Then one could decrease b^r by ε (or increase e^r accordingly) and increase b_1^v by $\frac{\varepsilon}{\beta\delta}$ without violating any constraint, but increasing the principal's profits.
(rIR) binds: Otherwise one could decrease both b^r and b_1^v without violating any constraint, but increasing the principal's profits. □

Learning Take any learning process characterized by a sequence of beliefs where the agent starts with a belief $\hat{\beta}_1$ about his present bias. Whenever he fails to accept the virtual contract, he learns that his beliefs $\hat{\beta}_{t-1}$ about his present bias must have been wrong and adjusts his belief to $\hat{\beta}_t \in [\beta, \hat{\beta}_{t-1}]$.

Corollary 5. *Consider an arbitrary learning process in which $\hat{\beta}_t > \beta$ for all t, i.e., the agent never learns his true β. Then learning about the present bias does not affect the optimally offered contracts.*

This immediately follows from Proposition 5.

Corollary 6. *If the learning process allows the agent to learn his true β (or leads to belief $\hat{\beta} < \beta$) with a positive probability, then it is optimal to add the first-best contract to the menu of contracts.*

This result is straightforward: When adding the first-best contract to the menu of contracts, an agent who has become sophisticated will from then on choose the first-best contract. A (fully or partially) naive agent does not have an incentive to select the first-best contract as he is indifferent between the first-best contract and the contract intended for him.

C.4 Weaker Present Bias in the Monetary Domain

Augenblick, Niederle, and Sprenger (2015) find in an experiment that people have a stronger present bias in effort than in monetary choices. Our results are qualitatively unchanged if the agent has a weaker or no present bias in the monetary domain. Moreover, we find that having a weaker or no present bias in the monetary domain even hurts a naive present-biased agent.

Let β be the discount factor for all future periods as before but now restricted to effort. Let $\beta' \in (\beta, 1]$ be that discount factor for money and the outside option. Morever, let b' and e' be the bonus and effort levels in these situations. By the same arguments as in our basic model, the equivalents of the contraints (rIRA) and (rC) hold with equality, i.e.

$$b^{r\prime} = c(e^{r\prime}) + \overline{u} + \frac{\beta'\delta}{1-\delta}\overline{u} - \delta(\beta' b_1^{v\prime} - \beta c(e_1^{v\prime})) - \frac{\delta^2}{1-\delta}\left(\beta' b^{v\prime} - \beta c(e^{v\prime})\right) \tag{1}$$

and

$$b^{r\prime} = c(e^{r\prime}) + b_1^{v\prime} - c(e_1^{v\prime}) + \delta\beta'(c(e^{v\prime} - c(e_1^{v\prime}).$$

These equations together imply

$$b_1^{v\prime} - c(e_1^{v\prime}) = \overline{u} + \frac{\delta\beta}{1-\delta}c(e^{v\prime}) - \frac{\delta\beta'}{1-\delta}(b^{v\prime} - \overline{u}). \tag{2}$$

Unlike in our basic model, where solving for the optimal first-period virtual contract only pinned down the *difference* between b_1^v and $c(e_1^v)$, the optimal $b_1^{v\prime}$ and $c(e_1^{v\prime})$ are unique. As one can see from equation (1), the principal benefits from choosing them as high as possible. Consequently,

$$b_1^{v\prime} = \theta e_1^{v\prime} - \overline{\pi}.$$

Furthermore, as the maximal future per-period surplus from the agent's perspective is determined by maximizing $\beta'(\theta e - \overline{\pi} - \overline{u}) - \beta c(e)$, the optimal effort for the later periods of the virtual contract is higher than in our basic model, i.e., $e^{v\prime} > e^v$. Together with equation (2) we can conclude that

$$b^{r\prime} < b^r - (\beta' - \beta)\delta\left(\frac{\delta}{1-\delta}b^{v\prime} - \frac{1}{1-\delta}\overline{u} + b_1^{v\prime}\right) < b^r.$$

This implies that the principal can exploit a naive agent whose present bias is more intense in the effort domain than in the monetary domain more than an an agent whose present bias is equally intense in both domains. The reason for this is that an agent with a less severe present bias in the monetary domain discounts the high bonus of the future virtual contract less, which lets the future virtual contract appear even more attractive to him.

C.5 Unobservable Agent Types

Assume there are two types of agents, $i \in \{1, 2\}$ with different values β_i. For simplicity, we set $\hat{\beta}_1 = \hat{\beta}_2 = 1$. Relaxing this would not affect the results. Without loss of generality, assume $\beta_1 < \beta_2 \leq 1$. This allows the principal to screen agents as they value the future differentially despite having the same belief about their future β. Let s_1 be the share of agents in the population with β_1, and $1 - s_1$ the share of agents with β_2. The principal cannot observe the agent's type, but only knows the distribution of types. As we assume that the principal can fully commit to the long-term contract contingent on the history of the game, she can preclude an agent, after the initial contract choice, to switch from one contract to the contract intended for the other type.[17] Hence, if different menus are offered, we need to make sure that it is optimal for agent i to select the respective contract intended for type i, which we denote $\tilde{\boldsymbol{C}}_i$.

Now, let $\boldsymbol{C}_1$ and $\boldsymbol{C}_2$ be the profit-maximizing contracts derived in Subsection 5.2, with the slight modification that the virtual contract is only offered from $t = 1$ on, i.e., the agent can only choose the real contract in $t = 0$. If these menus were also offered for unobservable types, agent 2 would actually go for $\boldsymbol{C}_1$. This is because agent 1 is just indifferent between taking up the contract and choosing his outside option. As $\beta_1 < \beta_2$, agent 2 values the future benefits of the virtual contract more than agent 1 and hence expects to receive a higher utility from $\boldsymbol{C}_1$. Furthermore, since $\boldsymbol{C}_1$ is designed in such a way that agent 1 is just indifferent between the virtual contract and the real contract, agent 2 would actually choose the virtual contract after selecting $\boldsymbol{C}_1$.

Hence, either $\tilde{\boldsymbol{C}}_2$ must be constructed in a way that makes it optimal for agent 2 to choose it (keeping $\tilde{\boldsymbol{C}}_1 = \boldsymbol{C}_1$), or $\tilde{\boldsymbol{C}}_1$ must be made sufficiently unattractive for 2 (keeping $\tilde{\boldsymbol{C}}_2 = \boldsymbol{C}_2$). The main result is given in Proposition 6; how exactly the menus are adjusted is described in the proof. We say that an agent is fully exploited when he receives the same utility as in our main analysis with one type of naive present-biased agent. We say that an agent is not exploited when he receives the utility of his outside option in every period. We say that an agent is partially exploited if an agent receives an intermediate utility.

[17] Note that it turns out to be in fact optimal for the principal to preclude these switches. If it were optimal to let the agent switch the principal could have just amended the original contract by the respective components of the alternative contract.

Proposition 6. *For all* $\beta_1, \beta_2, \delta > 0$*, with* $\beta_1 < \beta_2 < 1$*, there exists a threshold* $\underline{s} \in \left(0, \frac{\beta_2 - \beta_1}{\beta_2}\right)$ *such that for all* $s_1 \leq \underline{s}$ *it is optimal to offer a menu of contracts such that agent* 1 *is not exploited and agent* 2 *is fully exploited.*
Furthermore, it is optimal to offer two different contracts to the agents which both exploit the agents, but only partially for all $s_1 \in [\underline{s}, 1 - \beta_1]$ *. For all* $s_1 \geq 1 - \beta_1$*, it is optimal to fully exploit agent* 1 *and to partially exploit agent* 2*.*

Therefore, one type is still exploited in exactly the same way as with observable types, the other type is less exploited in order to induce a separation. Note that we also show below that a separation is always strictly optimal for $\beta_2 < 1$. If $\beta_2 = 1$ and $s_1 > 1 - \beta_1$, the principal is indifferent between inducing a separation or just letting agent 2 select $\tilde{\boldsymbol{C}}_1 = \boldsymbol{C}_1$ who would then go for the virtual contract. If writing different menus was associated with some small costs, we would predict only one menu, with non-present-biased agents actually "making a career". Therefore, our results are qualitatively unaffected if types cannot be observed, only the extent of exploitation has to be reduced for one type.[18]

Proof to Proposition 6

Proof. We will first approach the solution to the principal's screening problem under the assumption that offering one menu of contracts for each agent is optimal. Then we will show that offering one menu of contracts for each agent is indeed optimal.

Separation by Menu Choice Assume the principal wants the agents to choose different contracts. To develop an idea about the structure of these contracts, take $\boldsymbol{C}_1$ and $\boldsymbol{C}_2$, the profit-maximizing contracts derived in the main part, with the slight modification that the virtual contract is only offered from $t = 1$ on. Hence, agents can only choose the real contract in $t = 0$. We will discuss how to optimally modify these contracts below.

Assume the principal offered $\boldsymbol{C}_1$ and $\boldsymbol{C}_2$. Then agent 2's expected utility level when choosing $\boldsymbol{C}_1$ was
$\tilde{U}_2^r = \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) (1 - \beta_1) (\beta_2 - \beta_1) > 0$. This is positive because of $\beta_2 > \beta_1$, hence agent 2 puts more weight on future utilities than agent 1 does, and since both agents did not expect to get a rent before. Furthermore, because the respective (rC) constraints have been binding before, agent 2 would actually go for the virtual contract.

When choosing $\boldsymbol{C}_2$ (and expecting to select the virtual contract in the future), agent 1 gets:

[18] Other papers analyze the screening of present-biased agents: First, Heidhues and Kőszegi (2010) analyze a screening problem between two agents where one is naive and the other is sophisticated. By contrast, we allow for two naive agents. Second, Yan (2011) does not look at naive types with differently strong present biases. Finally, Li, Yan, and Xiao (2014) and Galperti (2015) analyze setups where all agents are sophisticated.

$\tilde{U}_1^r = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_2)(\beta_1-\beta_2) < 0$. Hence, agent 1 would stick to $\boldsymbol{C}_1$.

Separation by menu choice involves giving at least one agent an expected rent (which will also materialize in a real rent compared to the case with symmetric information). The principal could either adjust $\boldsymbol{C}_2$ in a way that it becomes more attractive for agent 2 (without making it too attractive for agent 1), or adjust $\boldsymbol{C}_1$ in a way that it becomes less attractive for agent 2, or adjust both.

First, note that the principal is restricted in increasing 2's virtual surplus - simply because this is already made as attractive as feasible for agent 2, with the exception of the first period of the virtual contract. Hence, the principal has the following options to make $\boldsymbol{C}_2$ more attractive: She can include an additional payment in period $t = 0$ (which only consists of the real contract), which we denote X_2. Alternatively, she can increase the payments in the (first period of the) virtual contract in $t = 1$. Since she still wants agent 2 to actually choose the real contract in $t = 1$ and since the (rC) constraint has been binding before, she must increase the payments in the real contract in period $t = 1$ by the same amount. This amount is denoted by Y_2. Offering agent 2 a contract where she chooses C_2 but then takes the virtual contract is dominated by this contract. In that case, the agent would still receive Y_2, but additionally capture the rents from later periods in the virtual contract.

Finally, the principal could reduce agent 1's payoff from the virtual contract and instead increase the real payoff he receives in period $t = 0$. We denote this amount by Z_1. More precisely, an increase of 1's real contract by Z_1 goes hand in hand with a reduction of his virtual payments by an amount Z_{1v} (in order to keep the rIR constraint for agent 1 binding) and potentially with an increase of his real payments in later periods (in order to keep the (rC) constraint for agent 1 binding). Note that a decrease of the virtual surplus in $t = 1$ does not only affect the contract in $t = 1$, but also limits what contracts the principal can offer in any later period. The principal cannot always simply decrease the payment of the virtual contract starting in the next period. If Z_{1v} is large and only the payment of the virtual contract starting in the next period was reduced, the agent would plan to choose the real contract in the next period and to choose the virtual contract only in the period after that (which would violate (vC)). Reducing the payment of the real contract is not an option, because the agent would eventually rather quit than choose the real contract. So the principal has to decrease the payment of the later virtual contract and even increase the next period's real contract in order to make the agent choose the real contract in the next period. If the necessary reduction of the payment in the later virtual contract is large, even later contracts might have to be changed by the same logic. The larger Z_{1v}, the more later periods are affected. When we talk about Z_1, we mean the full set of these adjustments. However, we first assume that any costs of these additional adjustments after $t = 0$ are zero, solve the simplified problem, and take the actual costs into account thereafter.

Now, expected payoffs when choosing the intended contracts and when all other components remain unchangend are

$U_1^r = Z_1 - \beta_1\delta Z_{1v} = 0$ (hence $Z_{1v} = \frac{Z_1}{\beta_1\delta}$) and $U_2^r = X_2 + \beta_2\delta Y_2$.

When deviating and selecting the other agent's menu, an agent's expected payoffs (and expecting to pursue the virtual contracts there) are
$\tilde{U}_1^r = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_2)(\beta_1-\beta_2) + X_2 + \beta_1\delta Y_2$ and
$\tilde{U}_2^r = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_1)(\beta_2-\beta_1) + Z_1 - \frac{\beta_2}{\beta_1}Z_1$.
If separation by menu-choice is intended, each agent must have an incentive to choose his intended contract, i.e., the no-deviation (ND) constraints $U_i^r \geq \tilde{U}_i^r$ must hold. Plugging in the respective values, we get

$$X_2 + \beta_1\delta Y_2 \leq \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_2)(\beta_2-\beta_1). \qquad \text{(ND1)}$$

for agent 1 and

$$X_2 + \beta_2\delta Y_2 + Z_1\left(\frac{\beta_2-\beta_1}{\beta_1}\right) \geq \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_1)(\beta_2-\beta_1) \qquad \text{(ND2)}$$

for agent 2.

Compared to the situation where the principal can observe each agent's β_i, in our simplified problem, she has to make additional real expected payments of

$$K = s_1 Z_1 + (1-s_1)(X_2 + \delta Y_2).$$

The profit-maximizing set of menus of contracts that induces a separation by menu choice now minimizes these costs, subject to (ND1) and (ND2).

First of all, note that (ND2) must bind. Otherwise, any of the payments could be reduced, thereby also relaxing (ND1) and reducing the principal's real costs. Furthermore, note that when comparing X_2 and Y_2, the principal would ceteris paribus always prefer to use X_2, i.e., using X_2 is cheaper than using Y_2: A reduction of Y_2 by ε requires increasing X_2 by $\delta\beta_2\varepsilon$ in order to keep (ND2) satisfied. This adjustment would lead to a cost change of $\delta\beta_2\varepsilon - \delta\varepsilon < 0$. However, if (ND1) binds as well, then setting $X_2 > 0$ will also require $Y_2 > 0$ to make agent 1 indifferent between the contracts.

The following lemma provides a lower bound of the total effective costs of using Z_1, X_2 and Y_2 as a function of s_1. We make use of the fact that (ND2) binds and that costs are linear in payments.

Lemma C4. *The following use of Z_1, X_2 and Y_2 minimizes the cost of separation by menu choice if there were no costs due to Z_1 later than in $t = 0$:*

- $s_1 \leq \frac{(\beta_2-\beta_1)}{\beta_2}$

  $X_2 = Y_2 = 0$ *and* $Z_1 = \beta_1(1-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$*; costs are*

$$K = s_1\beta_1(1-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$$

- $\frac{(\beta_2-\beta_1)}{\beta_2} < s_1 \leq 1-\beta_1$

  $Y_2 = 0,\ X_2 = (1-\beta_2)(\beta_2-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$ *and*

  $Z_1 = \beta_1(\beta_2-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$*; costs are*

$$K = \left[s_1\beta_1 + (1-s_1)(1-\beta_2)\right](\beta_2-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$$

- $s_1 > 1-\beta_1$

  $Z_1 = 0,\ X_2 = (1-\beta_1-\beta_2)(\beta_2-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$ *and*

  $\delta Y_2 = (\beta_2-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$*;*

  costs are

$$K = (1-s_1)(2-\beta_1-\beta_2)(\beta_2-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right).$$

Proof. By Lagrange optimization, we minimize costs, subject to (ND) as well as non-negativity constraints. There, note that Z_1 and Y_2 cannot be negative: If Z_1 was negative, player 1's (rIR) constraint would not hold (and it will not be optimal to increase agent 1's virtual surplus, since this would further tighten the (ND2) constraint). If Y_2 were negative, player 2's (rIR) constraint would not hold in period $t = 1$. X_2 can be negative, but only if Y_2 is increased accordingly. Hence, the constraint $X_2 + \delta\beta_2 Y_2 \geq 0$ must hold as well.

This gives the Lagrange function

$L = -s_1 Z_1 - (1-s_1)(X_2 + \delta Y_2)$

$+\lambda_{ND1}\left[\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_2)(\beta_2-\beta_1) - X_2 - \beta_1\delta Y_2\right]$

$+\lambda_{ND2}\left[X_2 + \beta_2\delta Y_2 + Z_1\left(\frac{\beta_2-\beta_1}{\beta_1}\right) - \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_1)(\beta_2-\beta_1)\right]$

$+\mu_X(X_2 + \delta\beta_2 Y_2) + \mu_Y Y_2 + \mu_Z Z_1$

and first-order conditions

$\frac{\partial L}{\partial Z_1} = -s_1 + \lambda_{ND2}\left(\frac{\beta_2-\beta_1}{\beta_1}\right) + \mu_Z = 0,$

$\frac{\partial L}{\partial X_2} = -(1-s_1) - \lambda_{ND1} + \lambda_{ND2} + \mu_X = 0,$

$\frac{\partial L}{\partial Y_2} = -(1-s_1)\delta - \beta_1\delta\lambda_{ND1} + \beta_2\delta\lambda_{ND2} + \delta\beta_2\mu_X + \mu_Y = 0.$

We know that $\lambda_{ND2} > 0$, furthermore rearranging and substituting gives the three conditions

$\frac{\partial L}{\partial Z_1}: \lambda_{ND2} = \frac{\beta_1(s_1-\mu_Z)}{\beta_2-\beta_1}$, (I)

$\frac{\partial L}{\partial X_2}: \lambda_{ND1} = \frac{-(\beta_2-\beta_1)+s_1\beta_2}{\beta_2-\beta_1} - \frac{\beta_1\mu_Z}{\beta_2-\beta_1} + \mu_X$, (II)

$\frac{\partial L}{\partial Y_2}: s_1 - (1-\beta_1) - \beta_1\mu_Z + \frac{\mu_Y}{\delta} + \mu_X(\beta_2-\beta_1) = 0.$ (III)

Combining (II) and (III) implies that $\lambda_{ND1} = \frac{(1-s_1)(1-\beta_2)-\frac{\mu_Y}{\delta}}{\beta_2-\beta_1}$.

In the following, we just go through all potential cases and analyze whether they are feasible and if yes under which conditions.

1. $s_1 - (1-\beta_1) > 0$. Then, (III) implies that $\mu_Z > 0$, giving the following potential cases:

   (a) $\mu_Y > 0$: This is not feasible, since for $Y_2 = Z_1 = 0$, obtaining $X_2 = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_1)(\beta_2-\beta_1)$ from binding (ND2) and plugging it into (ND1) gives $\beta_1 \geq \beta_2$, which is ruled out by assumption.

   (b) $\mu_Y = 0$: Then, (ND1) binds as well. Obtaining $X_2 = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_1)(\beta_2-\beta_1) - \beta_2\delta Y_2$ from binding (ND2) and plugging it into (ND1) gives $Y_2 = \frac{\delta}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(\beta_2-\beta_1)$, implying that $X_2 = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_1-\beta_2)(\beta_2-\beta_1)$. Then, $X_2 + \delta\beta_2 Y_2 = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_1)(\beta_2-\beta_1) > 0$.

2. $s_1 - (1-\beta_1) < 0$. Then, (III) implies that either $\mu_X > 0$ or $\mu_Y > 0$, or both, giving the following potential cases:

   (a) $\mu_X > 0$, $\mu_Y > 0$: Since $X_2 = Y_2 = 0$, $\mu_Z = 0$, and a binding (ND2) constraint gives $Z_1 = \beta_1(1-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$. Furthermore, (II) implies that this is only feasible for $-(\beta_2-\beta_1) + s_1\beta_2 < 0$.

   (b) $\mu_X > 0$, $\mu_Y = 0$: Hence, (ND1) binds as well. Plugging $X_2 = -\delta\beta_2 Y_2$ into (ND2) gives $Z_1 = \beta_1(1-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$. Plugging $X_2 = -\delta\beta_2 Y_2$ into (ND1) gives $\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_2)(\beta_2-\beta_1) + \delta Y_2(\beta_2-\beta_1) = 0$, which is not feasible.

   (c) $\mu_X = 0$, $\mu_Y > 0$:

      i. $\mu_Z = 0$: (III) gives $\mu_Y = -\delta[s_1 - (1-\beta_1)]$, hence $\lambda_{ND1} = \frac{s_1\beta_2}{\beta_2-\beta_1} - 1 > 0$ if $-(\beta_2-\beta_1) + s_1\beta_2 > 0$. Then, binding (ND1) gives $X_2 = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_2)(\beta_2-\beta_1)$, and plugging this into (ND2) gives: $Z_1 = \beta_1(\beta_2-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$. (II) implies that this is only feasible for $-(\beta_2-\beta_1) + s_1\beta_2 \geq 0$.

      ii. $\mu_Z > 0$: Only using X_2 is not feasible (see 1.(a)).

□

By the above lemma, we obtain restrictions on $\underline{s}$ and the threshold $1-\beta_1$. Note that the menu offered for low s_1 is $\boldsymbol{C}_2$ and the first-best contract (intended to be chosen by agent

1).[19] It is easy to see that this is optimal even for the actual costs of Z_1[20] if s_1 is low enough: The costs of Z_1 occur for the fraction of agents of type 1, s_1, whereas X_2 and Y_2 have to be paid to the fraction of agents of type 2, $1-s_1$. As Z_1 can be substituted by X_2 and Y_2 in a linear way (to separate the types), there exists an $\underline{s}$ such that $X_2 = Y_2 = 0$ is optimal for all $s_1 \leq \underline{s}$. Hence, when there are only few agents of type 1, the principal does not exploit them, while fully exploiting the agents of type 2.
To see that it is optimal to alter both $\boldsymbol{C}_1$ and $\boldsymbol{C}_2$ compared to the case without screening for intermediate values of s_1, note that the lower bound for the cost of using Z_1 is equal to the actual costs if Z_1 is low enough and does not require to alter any contracts but the real contract at the begining of the game and the virtual contract starting in the following period. Hence, the principal should choose a positive Z_1 for $s_1 < 1-\beta_1$. At the same time, the principal should not only use Z_1 as the actual costs for using Z_1 would be strictly larger than the lower bound in this case.

Note that also (vIRA), (vIC), and (vC) are fulfilled, as they are not affected by the changes.[21] For (vrIRP) and (vIRP) observe that the compensation in the real contract in the first period and in the virtual contract's first period are not higher than the principal's surplus in these periods: $c(e^{FB}) + \overline{u} - \beta_2(1-\beta_2)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) + (1-\beta_2)\left(\beta_2 - \beta_1\right)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) \leq e^{FB}\theta - \overline{\pi}$ and
$c(e^{FB}) + \overline{u} - \beta_2\frac{\delta}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) + \left(\beta_2 - \beta_1\right)\frac{\delta}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) \leq e^{FB}\theta - \overline{\pi}$. For the real contract in the second period the compensation does not exceed the principal's surplus when $(\beta_2 - \beta_1 - \beta_2(1-\beta_2)\delta - \beta_2\delta)\frac{\delta}{1-\delta} \leq 1$, which is true.

Separation by Action Now we calculate the costs for the principal when she offers only one menu of contracts which is supposed to be chosen by both agents and show that these costs are higher than making agents choose different menus of contracts.

Instead of offering a menu of contracts for each type of agent, the principal could just let agent 2 choose 1's contract, taking into account that 2 would then go for the virtual contract and make a career. In this case, it is without loss of generality to assume that only the profit-maximizing menu for agent 1, $\boldsymbol{C}_1$, is offered. Note that it cannot be optimal to induce separation by menu choice and then let agent 2 actually choose the virtual contract

[19] To see this, note that $Z_1 = \beta_1\left(1-\beta_1\right)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$ comes with $Z_{1v} = \left(1-\beta_1\right)\frac{\delta}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$ which means that agent 1's perceived net utility from $t=1$ on is reduced to zero. At the same time, making the agent accept a real contract that puts his utility below his outside option requires a benefit in the future virtual contract such that the perceived combined utility of accepting the real contract and choosing the virtual contract in the future is at least (and optimally exactly) zero. As this must hold even if the virtual contract is discounted with $\beta\delta$ from the time when the agent accepts the real contract, this means that taking the future real contract and the virtual contract after that must have a positive net utility from agent's perspective in $t=0$. This is a contradiction to it being reduced to zero.

[20] In fact, the costs of increasing Z_1 are $s_1\left(Z_1 + \sum_{i=1}^{\infty}\max\{0, \frac{Z_1}{\delta^2(1-\beta_1)} - \frac{\beta_1}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)\left(1-\delta^i(1-\beta_1)\right)\}\right)$. To simplify the analysis, we use a cost function that does not take into account potential payments in later real contracts (due to a high Z_1) as a lower bound.

[21] (vC) is fulfilled by construction if Z_1 is large.

(unless $\beta_2 = 1$). Such a setting would give agent 2 a higher real rent than the one derived above because (ND) constraints would still have to hold.

If agent 2 is offered $\boldsymbol{C}_1$, his expected as well as real utility is

$$U_2^r = \tilde{U}_2^r = (1-\beta_1)(\beta_2-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right),$$

whereas agent 1 expects to get nothing.

Compared to the situation with symmetric information, though, the principal also foregoes the benefits from exploitation – because agent 2 not only select $\boldsymbol{C}_1$, but also goes for the virtual contract. Recall that under symmetric information, the net per-period profits the principal generates are $\pi^r - \overline{\pi} = \left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\left(1 + \beta(1-\beta)\frac{\delta^2}{1-\delta}\right)$. Hence, under symmetric information, the principal's total profits when dealing with agent 2 would be $\frac{\left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)}{1-\delta}\left(1 + \beta_2(1-\beta_2)\frac{\delta^2}{1-\delta}\right)$. If letting agent 2 choose $\boldsymbol{C}_1$, the principal's profits dealing with agent 2 are $(1+\delta)\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) + (2-\beta_1)\beta_1\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$.

Therefore, the principal's costs when letting agent 2 choose $\boldsymbol{C}_1$, taking into account that he then actually goes for the virtual contract, compared to the case of symmetric information (which also served as our benchmark above), are

$$\tilde{K} = (1-s_1)\left[1 + \frac{\beta_2(1-\beta_2)}{1-\delta} - (2-\beta_1)\beta_1\right]\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right).$$

Comparing these costs to the costs of separation by menu choice for large s_1 (which serve as an upper bound), the condition equivalent to separation by action being cheaper,

$$\left[1 + \frac{\beta_2(1-\beta_2)}{1-\delta} - (2-\beta_1)\beta_1\right] \leq (2-\beta_1-\beta_2)(\beta_2-\beta_1),$$

shows that separation by action can only be optimal if $\beta_2 = 1$.

For $\beta_2 = 1$, the cost expressions for separation by actions and for separation by menu choice are the same for $s_1 > 1-\beta_1$, so separation by action is (weakly) optimal. □

More than two Types Finally, we discuss what might happen with $n > 2$ different types? Let $\beta_1 < \beta_2 < ... < \beta_n \leq 1$. By a similar argument as for two types, there exists a threshold $\underline{s}^{prime}$ such that if $s_1 + ... + s_{n-1} < \underline{s}'$, i.e., if type n is very common, it is optimal to not exploit types 1 through $n-1$ and to fully exploit type n. We conjecture that this logic applies more generally: If types $1, ..., i$ are rare, where $i \in \{1, ..., n-2\}$, the principal will not exploit these types (because preventing higher types from choosing those exploitative menus would be more costly), but instead focus on how to optimally exploit types $i+1, ...n$. Similarly to the case of two types, if type n is rare, it will be optimal for the principal to increase the real payments in the first and second period, as well as the payments of the virtual contract's first period for the menu to be chosen by type n – because this is cheaper than letting type n select another type's menu and eventually go for the virtual contract. However, this does not allow the principal to fully exploit the other types: Because there are at least two further types, the same logic as in the case of just two different types applies, and they cannot be given the profit-maximizing menus derived in Subsection 5.2. We conjecture that for those types – as well as when type n is neither very common nor rare – the principal uses the same instruments as in the case of two types to partially exploit various types. An interesting difference to the case of two types might be the following: With $n > 2$ it is possible that a naive agent may be allowed to choose a virtual contract. If type $j \in \{2, ..., n-1\}$ is very rare, deterring type $j+1$ from choosing j's menu and deterring type j from choosing $j-1$'s menu would be more costly than letting j choose $j-1$'s menu and eventually the virtual contract.