

DISCUSSION PAPER SERIES

DP13227

LONG-TERM EMPLOYMENT RELATIONS WHEN AGENTS ARE PRESENT BIASED

Florian Englmaier, Matthias Fahn and Marco
Schwarz

INDUSTRIAL ORGANIZATION



LONG-TERM EMPLOYMENT RELATIONS WHEN AGENTS ARE PRESENT BIASED

Florian Englmaier, Matthias Fahn and Marco Schwarz

Discussion Paper DP13227
Published 08 October 2018
Submitted 02 October 2018

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Florian Englmaier, Matthias Fahn and Marco Schwarz

LONG-TERM EMPLOYMENT RELATIONS WHEN AGENTS ARE PRESENT BIASED

Abstract

We analyze how agents' present bias affects optimal contracting in an infinite-horizon 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. Moreover, under limited liability, implemented effort can be inefficiently high. With a finite time horizon, the degree of exploitation of agents decreases over the life-cycle. While the baseline model abstracts from moral hazard, we show that the result persists also when allowing for non-contractible effort.

JEL Classification: D03, D21, J31, M52

Keywords: employment relations, Dynamic Contracting, present bias

Florian Englmaier - Florian.Englmaier@econ.lmu.de
LMU München and CEPR

Matthias Fahn - matthias.fahn@jku.at
JKU Linz

Marco Schwarz - marco.schwarz@uibk.ac.at
University of Innsbruck

Acknowledgements

We are grateful to Dan Barron, Paul Heidhues, Botond Köszegi, 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, and at the 2015 CESifo conference on Behavioral Economics in Munich 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.

Long-Term Employment Relations When Agents are Present Biased*

Florian Englmaier[†] Matthias Fahn[‡] Marco A. Schwarz[§]

August 6, 2018

Abstract

We analyze how agents' present bias affects optimal contracting in an infinite-horizon 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. Moreover, under limited liability, implemented effort can be inefficiently high. With a finite time horizon, the degree of exploitation of agents decreases over the life-cycle. While the baseline model abstracts from moral hazard, we show that the result persists also when allowing for non-contractible effort.

JEL Codes: D03, D21, J31, M52

Keywords: employment relations, dynamic contracting, present bias

*We are grateful to Dan Barron, Paul Heidhues, 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, and at the 2015 CESifo conference on Behavioral Economics in Munich 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

Numerous studies have documented that a substantial fraction of the population suffers from a present bias, and that this bias results in problems of self control.¹ This present bias affects decision making in diverse domains like retirement savings, see Laibson (1997), health club attendance, see DellaVigna and Malmendier (2006), or credit card usage, see Shui and Ausubel (2005). All these settings have in common that immediate costs have to be traded off with delayed benefits and that their self-control problems leads people to make too “shortsighted” decisions. A domain that, so far, has not been studied with this focus but that also shares this feature – immediate costs, delayed rewards – is individuals’ career decisions in the labor market.

Deciding which career to take and where to work is one of the most important decisions people make. It is an inherently long-term decision, involving inter-temporal trade-offs. Moreover, for most employees, and this is well understood by firms, the prospect of career advancement is among the most important motivating factors. Survey evidence suggests that most young employees want to make a “career”, i.e., they strongly care about opportunities for career advancement and development that are offered by prospective employers² and the opportunities for career development are a top priority in job choice.³ Hence, many firms prominently advertise their development programs⁴ and attract and motivate employees with promotion prospects and career development plans.

However, there is also widespread frustration that these promotion and career prospects often turn out to be false hopes.⁵ And this frustration is not (only) caused by firms defaulting on their promises, or by a lack of talent: The popular management literature has abundantly referred to aspects of self-control problems as major reasons for why so many

¹See DellaVigna (2009) for an excellent survey.

²See, e.g., slide 28 of the “E&Y Studentenstudie 2014”, [https://www.ey.com/Publication/vwLUAssets/EY_-_Acht_von_zehn_Studenten_sind_zufrieden/\\$FILE/EY-studentenstudie-2014-werte-ziele-perspektiven-pr%C3%A4sentation.pdf](https://www.ey.com/Publication/vwLUAssets/EY_-_Acht_von_zehn_Studenten_sind_zufrieden/$FILE/EY-studentenstudie-2014-werte-ziele-perspektiven-pr%C3%A4sentation.pdf), last accessed, June 26, 2018. It states that 96% of the respondents regard these aspects as important.

³See, e.g., slide 23 of the “E&Y Absolventenbefragung 2013”, [http://www.ey.com/Publication/vwLUAssets/EY-Absolventenbefragung_2013/\\$FILE/EY-Absolventenbefragung-2013-Studie.pdf](http://www.ey.com/Publication/vwLUAssets/EY-Absolventenbefragung_2013/$FILE/EY-Absolventenbefragung-2013-Studie.pdf), last accessed, June 26, 2018. It states that this aspect was among the top five criteria for 72% of respondents.

⁴To be a “Great Place to work” the Harvard Business Review asks firms to “Provide employees with ongoing opportunities and incentives to learn, develop and grow [...]”; see <https://hbr.org/2011/09/the-twelve-attributes-of-a-tru.html>, last accessed January 28, 2015.

⁵The phenomenon has been documented by several studies in organizational psychology; see, e.g., Garavan and Morley (1997) or Rindfuss, Cooksey, and Sutterlin (1999).

employees fail to take career development options. For example, a "lack of persistence" has been identified as a reason for "Why People Fail in Their Careers"⁶, "laziness" has been found to be one of "11 Reasons for Career Failure"⁷, "you are not persistent enough" and "you fail because you're too distracted" are two of "10 Reasons Why You Fail"⁸, and "you can, but won't, do what's necessary" is among the "Top 8 Reasons You Are Failing at Your Job"⁹.

Here we offer an explanation of this phenomenon that builds on the above described core feature of career choices: A firm designing a profit-maximizing incentive scheme *endogenously* uses a cost/reward structure with immediate costs, but delayed benefits. Such a structure opens up room for time-inconsistency phenomena. In this paper we study how the design of optimal long-term employment contracts is affected by such a present bias. We analyze how agents' present bias affects contracting in an infinite-horizon employment setting, first abstracting from other agency problems such as moral hazard, adverse selection, or limited commitment. Naive agents are offered a menu of contracts, consisting of a *virtual* contract – which they plan to choose in the future – and a *real* contract which they end up choosing. This virtual contract allows the principal to exploit the agent, because a major part of the agent's compensation is shifted from the real into the virtual contract. If the agent is protected by limited-liability, implemented effort can be inefficiently high from a social planner's perspective. Moreover, it turns out that changes that in general should improve the agent's situation, like employment protection legislation, in fact hurt him in our setting. Finally, considering a finite-horizon version of our model reveals that the degree of exploitation of naive agents decreases over the life-cycle, i.e., older workers are *ceteris paribus* better off.

In the main part of the analysis, we assume a setting where a risk neutral principal and a risk neutral agent interact over an infinite time horizon. Whereas the principal discounts future profits exponentially, the agent discounts future utilities in a quasi-hyperbolic way. In any period, the agent can either work for the principal or not. If the agent works

⁶See <https://www.theglobeandmail.com/report-on-business/careers/management/sixteen-reasons-why-people-fail-in-their-careers/article621085/>, last accessed June 13, 2017.

⁷See <http://www.positionignition.com/blog/2013/4/11/11-reasons-for-career-failure.html>, last accessed June 13, 2017.

⁸See <https://markmanson.net/why-you-fail>, last accessed June 13, 2017.

⁹See <http://www.cbsnews.com/news/top-8-reasons-you-are-failing-at-your-job/>, last accessed June 13, 2017.

for the principal, he chooses costly effort. Effort generates a deterministic output which is consumed by the principal. The agent receives a fixed wage payment and a bonus. As effort today yields output today and effort and output are verifiable, a present bias per se does not represent an obstacle for efficient outcomes. A static contract could implement efficient allocations and allow the principal to extract the full surplus. Therefore, sophisticated present-biased agents, i.e., those who are aware of their future self-control problems, are effectively in the same situation as agents without any present bias. They do not suffer from their self-control problems because it is optimal to compensate them for their effort at the end of a given period. Any distortion will thus be due to the fact that the principal, who is aware of a naive agent's bias, can design a dynamic contract which makes use of the different priors principal and agent have concerning the agent's future time preferences.

A naive agent, who is not aware of his future self-control problems, is offered a menu of contracts, consisting of a virtual contract which the agent intends to choose in the future and a real contract which he ends up choosing. The virtual contract has a "probation phase" in the first period it is chosen. In this probation phase, the agent receives a low utility as he either has to work extra hard or receives only a very low compensation. This probation phase deters the present-biased agent from actually choosing the virtual contract as it entails immediate costs but only delayed gratification. After the probation phase, the virtual contract grants the agent very attractive compensation. From today's perspective, the prospect of this attractive compensation after the probation phase outweighs the low utility during the probation phase tomorrow and thus makes the agent willing to accept a lower compensation today. However, when tomorrow comes, the lower utility during the probation phase does not seem worth the future benefits anymore and the agent postpones taking up the virtual contract. Consequently, the agent always ends up choosing the real contract which gives him a utility below his outside option. Hence, in our setting, long-term contracts are strictly better for the principal than a sequence of spot contracts.¹⁰

To come back to the above example of career choice, we find that firms can exploit naive present-biased employees by offering them a lucrative "career" as compensation for (presumably) short-run sacrifices. However, the naive present-biased agent fails to take up this option as he incrementally postpones making an extra investment to get onto the

¹⁰This differs from standard results in contract theory where, in general, whenever actions taken by the agent in period t affect the output of period t only, a sequence of short-term contracts is as efficient as a long-term contract; see, e.g., Malcomson and Spinnewyn (1988) or Fudenberg, Holmstrom, and Milgrom (1990).

career trajectory. Nevertheless, the employee is willing to accept lower compensation today because of his misconceived effort (and career) expectations. As a consequence, the firm can provide incentives at substantially lower costs. When we analyze a version of our model with a finite horizon, (young) agents who are at the beginning of their working lives will be exploited more by the principal as they can be more easily lured by career prospects: They work weakly more (strictly so, if there is a binding limited liability constraint) 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.¹¹

Finally, note that how much the principal can exploit the agent with the real contract 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 from today's perspective – such as a higher “standard” discount factor or lower outside options – actually harms a naive present-biased agent. This gives rise to the additional interpretation of our model where a more stringent employment protection actually harms a naive present-biased agent: The principal's outside option (which can also be negative) includes the cost the latter has to bear when firing the agent. Therefore, higher firing costs (for example induced by a more stringent employment protection) increase the future virtual surplus and consequently reduce the amount the agent is paid in the real contract.

Moreover, we analyze a number of variations of the model that deliver additional insights on how dealing with present-biased employees might affect labor market outcomes. We let the agent be protected by limited liability and show that competition in the labor market can mitigate the extent of exploitation. Furthermore, we allow for a weaker present bias in the monetary domain, as well as for moral hazard, and partial naivete of the agent. We also allow the agent to learn about his present bias over time and consider the case of heterogeneous and unobservable agent types. None of these variations changes our findings qualitatively. They rather offer additional insights:

If the agent is protected by limited liability, the principal cannot always fully exploit the agent when letting him work at the efficient effort level. Instead, once limited liability becomes binding and reducing the payment is no longer possible because the agent already

¹¹This general pattern, overly hard work early on in an agent's working life and relatively less hard work towards the end of a career is also generated in models of career concerns; see Holmstrom (1999). However, there this pattern is driven by a signal-jamming logic.

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). Such an inefficiency is more pronounced for (*ceteris paribus*) more productive employment relationships. Furthermore, having limited liability in a finite time horizon would give rise to an effort level that gradually goes down over time.

Introducing competition on the labor market, which we model by allowing the agent to capture a varying share of the relationship surplus (in our benchmark model, the principal enjoys full bargaining power), does not change the qualitative structure of the optimal menu of contracts. However, we can show that in more competitive labor markets there is less scope to exploit agents, since more competition for workers increases their real rents and hence restricts the principal's ability to take advantage of an agent's present bias. Moreover, if the agent enjoys a considerable degree of bargaining power, he is not harmed anymore by a larger perceived virtual surplus. In this case, policies that reduce the principal's outside option, such as a more stringent employment protection, will eventually benefit the agent.

When we allow for a weaker present bias in the monetary domain, our results do not change qualitatively. Moreover, we find that a naive agent who has a present bias in the effort domain may benefit from a bias in the monetary domain because this would reduce the degree to which compensation could be shifted from the real into the virtual contract.

Under the assumption of unobservable agent types regarding the degree of the present bias, the principal generally still offers exploitative menus of contracts. In this case, only the extent of exploitation is lower than when types are observable in order to induce each agent to select the menu intended for his type.

Allowing for moral hazard, the existence of partially naive types, and the possibility of learning about one's type do not affect the structure of profit-maximizing contracts. Neither moral hazard nor bargaining change the basic structure of the optimal contract. A (partially) naive agent expects that in the future choosing the virtual contract will be strictly better than choosing the real contract, while in fact, tomorrow he will be just indifferent. For this decision, his true present bias is relevant, not his perception about its future extent. This result on partial naivete also implies that the offered contracts do not change when the agent is able to learn that he is time-inconsistent (unless he learns his bias perfectly and becomes fully sophisticated).

The paper is organized as follows. Following a literature overview in Section 2, Section 3

introduces our model. Section 4 deals with the benchmark cases of the principal facing a rational or a sophisticated agent. Section 5 sets up the principal’s problem when facing a naive time-inconsistent agent and presents the results of our basic model. Section 6 deals with a finite time horizon, limited liability, a weaker present bias in the monetary domain, moral hazard, competition and bargaining, partially naive agents, learning, and adverse selection. Section 7 concludes, Appendix A collects the figures, and Appendices B and C collect all proofs.

2 Related Literature

The paper relates to the literature on inconsistent time 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 inconsistent time preferences affect savings, the possibility of individuals having a present bias should be incorporated into growth models. Laibson (1997) shows that given people have inconsistent time preferences, 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 inconsistent time preferences for individuals, 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 time inconsistency and the latter are not.¹²

Besides numerous theoretical contributions, there also 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 et al. (2010, 2015) provide evidence from a field experiment with data entry workers that self-control problems at work are important. The study by Augenblick, Niederle, and Sprenger (2015) is particularly interesting for us as they document that subjects show a considerable present bias in effort (while they only find limited time inconsistency in monetary choices). This suggests that studying the role of present bias in the workplace and

¹²Further theoretical contributions like Harris and Laibson (2001) or Krusell and Smith (2003) show that hyperbolic discounting of individuals can have important implications for consumption-savings decisions.

in workers' careers is a particularly relevant and promising topic for research.

There is also a literature focusing on optimal contracting choices when agents exhibit time-inconsistencies. 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 problem 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 the slacking-off is unpredicted by a naive agent), however requiring the agent to "compensate" the principal for that. We relate to Gilpatric (2008) in the sense that agents deviate, post entering the contract, from their planned choices, and that the principal foresees that and can exploit it. However, as in Gilpatric (2008) principal and agent only interact once there is neither a way to interpret dynamics in his model nor is there scope to gain insights regarding our host of labor market relevant variations we study in Section 6.

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 time-inconsistency, a lesser degree of consumption smoothing is optimal. Similarly, Li, Yan, and Yu (2012) find that the principal might optimally sell a risk project to a risk-averse agent. However, both restrict the contracting space to consist of only one element, and hence do not allow the agent to revise his observable actions in future periods.

From a modeling perspective, but analyzing a rather different environment, the paper closest to ours is Heidhues and Kőszegi (2010). They analyze contracting with time-inconsistent consumers in a competitive credit market and find that naive consumers are attracted by contracts 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 the application, interpretation of the contract features, relevance of institutional details of the (labor market) application, and the specific interpretation of results and policy implications differ from this study.

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 naivete. 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.

Being naive about one’s future time preferences might also be perceived as a specific form of overconfidence. In a principal-agent setting with moral hazard, Gervais and Odean (2011) and de la Rosa (2011) find that, the principal potentially can incentivize the agent more cheaply as the agent overestimates the likelihood to succeed. In our model the agent’s bias makes it cheaper for the principal to incentivize the agent as well, however, the structure of the optimal contract is entirely different. Hoffman and Burks (2015) 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 - as Hoffman and Burks (2015) interpret it - the drivers’ overconfidence as drivers with a larger bias are less likely to quit.

3 Model Setup

Technology & Per-Period Utilities There is one risk neutral principal (“she”) and one risk neutral agent (“he”). We consider an infinite time horizon where time is discrete and periods are denoted by $t = 0, 1, 2, \dots$. In any period t , the agent can either work for the principal or not, which is described by $d_t \in \{0, 1\}$. If $d_t = 1$, the agent works for the principal and 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 $\lim_{e_t \rightarrow \infty} c' = \infty$. Effort e_t generates a deterministic output $y(e_t) = e_t\theta$ which is consumed by the principal. There, $\theta > 0$ captures the value of the agent’s effort for the principal.

Furthermore, the agent receives a fixed wage payment w_t and a bonus $b_t(e_t)$. Note that we do not impose a limited liability constraint on the agent, hence the payments can (and under some instances will) be negative (implying payments from agent to principal). We consider the case of limited liability below, in Section 6.2. The agent's payoff in period t when $d_t = 1$ is

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

whereas the principal receives

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

If $d_t = 0$, i.e., the agent does not work for the principal in period t , he receives his outside option \bar{u} . The principal's outside option in this case is denoted by $\bar{\pi}$ (where \bar{u} and $\bar{\pi}$ can also be negative).

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

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

For the remainder of the paper, we assume $\theta e^{FB} - c(e^{FB}) - \bar{u} - \bar{\pi} > 0$, 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 future utilities in a quasi-hyperbolic way according to Phelps and Pollak (1968) and Laibson (1997): While current utilities are not discounted, future (period t) utility is 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. Furthermore, as is standard in these kinds of models, the present-biased agent is also more impatient than the principal, i.e., there is no normalization of δ . Concerning the agent's belief about his future preferences for instant gratification, we follow O'Donoghue and Rabin (2001) and their description of partial naivete. 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 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 persistent present bias. In Section 6.6 we explicitly allow for partial naivete, i.e., $\hat{\beta} \in (\beta, 1)$ and show that the outcome in this case is *exactly* the same as with a fully naive agent. Furthermore, in Section 6.7 we also allow for learning in the sense that $\hat{\beta}$ decreases over time and show that results are robust as long as learning is not perfect. Finally, note that in the main part of the analysis we assume the agent to be equally present-biased regarding money and effort but show in Section 6.4 that our results also hold if the present bias is weaker in the monetary domain.

Perceptions Here, 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 predicts 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 Section 6.8 we explicitly allow for unobservable types and derive optimal screening contracts.

Contractability, Timing, and Histories To isolate the effect of the agent’s present bias on the structure of the employment relationship, we abstract from any other potential agency problem. Hence, the agent’s effort as well as wage and bonus payments are verifiable and the principal can commit to long-term contracts.¹³ The principal’s commitment is limited by the firm value, though, and she can always escape her obligations by declaring bankruptcy and proceeding to consume her outside option $\bar{\pi}$ in every subsequent period.¹⁴

¹³This strong commitment could be endogenized in a setting with many agents where the principal’s behavior can be observed by everyone, and where she cares about her reputation to keep her promises.

¹⁴Alternatively, we can assume that the principal can fire the agent at some cost – for example reflecting the degree of employment protection – implying that $\bar{\pi}$ could also be negative. Bankruptcy or firing costs are then given by $-\frac{\bar{\pi}}{1-\delta}$.

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 about his future choices, she actually does not want to grant the agent the opportunity to commit to a long-term contract. Moreover, if the agent is either not present-biased or sophisticated about his present bias, any long-term commitment on his side will not increase profits. Hence, we can restrict attention to situations where the agent is free to leave in any period.¹⁵

In the first period of the game, in $t = 0$, the principal makes a take-it-or-leave-it long-term contract offer, denoted by \mathbf{C} , to the agent. This offer consists of a menu of contracts (with finitely many elements) for every future period, contingent on any possible history. Each of these contracts contains a fixed wage payment and a bonus for every potential effort level. The principal is fully committed to this long-term contract and could only walk away from her obligations by declaring bankruptcy.

The menu of contracts offered in period t is denoted by C_t . It has I_t elements, where a single element is indexed $i_t \in \{0, \dots, I_t\}$, i.e., $C_t = \{C_t^{i_t}\}_{i_t=1}^{I_t} = \{w_t^{i_t}, b_t^{i_t}(e_t)\}_{i_t=1}^{I_t}$, where $w_t^{i_t} \in (-\infty, \infty)$ for each i_t , and $b_t^{i_t}(e_t) \in (-\infty, \infty)$ for each i_t and each e_t . At the beginning of every period t , the principal first decides whether to declare bankruptcy or whether to offer C_t . This choice is denoted by $d_t^P \in \{0, 1\}$, where $d_t^P = 0$ indicates a bankruptcy and implies that $C_t = \{\emptyset\}$.

Declaring bankruptcy is an irreversible decision (hence, $d_{t+1}^P = 0$ automatically follows from $d_t^P = 0$), inducing principal and agent to consume their outside utilities $\bar{\pi}$ and \bar{u} in every subsequent period. Given $d_t^P = 1$, the agent chooses whether to work for the principal or not, hence selects $d_t \in \{0, 1\}$. If $d_t = 0$, principal and agent consume their outside utilities in the respective period. If $d_t = 1$, the agent selects one element out of C_t , where this choice is denoted by $\hat{i}_t \in I_t$. Then, the agent receives $w_t^{\hat{i}_t}$ and makes his effort choice, triggering the (automatically enforced) bonus payment $b_t^{\hat{i}_t}(e_t)$. Finally, the principal consumes the output $e_t\theta$ and the game moves on to the next period.

The publicly observable events during periods $t \geq 0$ are $h_t = \{d_t^P, C_t, d_t, \hat{i}_t, w_t, e_t, b_t(e_t)\}$. The history of publicly observable events at the beginning of period t is $h^t = \cup_{\tau=1}^{t-1} h_\tau \cup \{\mathbf{C}\}$, with $h^0 = \{\emptyset\}$. Furthermore, H^t is the set of histories until t and $H = \cup_t H^t$ the set of histories.

¹⁵Again, any costs for the agent to leave his current occupation could be captured by an appropriate choice of \bar{u} .

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 predicts 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 \mathbf{C} . In every period $t \geq 0$, \mathbf{C} maps the history $h^t \in H^t$ into an offered menu of contracts, C_t . Following any history $h^t \in H^t$, $t \geq 0$, σ_P also determines whether the principal follows \mathbf{C} or terminates the contract by declaring bankruptcy.

An agent’s strategy is denoted by σ_A and – given $h^t \cup \{d_t^P\} \cup \{\mathbf{C}\}$ – determines d_t and eventually \hat{i}_t and e_t .

Real and Virtual Contract Without loss of generality, we can restrict C_t to either consist of one or two elements, depending on the agent’s naivete. If the agent is sophisticated or not present-biased, he does not have misperceptions concerning his future behavior, and C_t consists of exactly one element. Given that a profit-maximizing contract is included in the menu, adding additional contracts will have no (or even adverse) effects on profits.

If the agent is (partially) naive, the principal finds it optimal to let C_t consist of exactly two elements: the element that the agent believes to choose in the future, and the element the agent is actually going to choose (which is correctly anticipated by the principal).¹⁶ We call the former the virtual contract and describe the respective components (as well as the effort level the agent expects to choose when selecting the virtual contract) with a superscript “ v ”. The latter is called real contract and its components (as well as the effort level the agent chooses when selecting the real contract) 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.

Payoffs The agent’s actually realized utility stream at the beginning of any period t in a setting where the principal never declares bankruptcy is

$$U_t^r = d_t^r (b_t^r + w_t^r - c(e_t^r)) + (1 - d_t^r) \bar{u} + \beta \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} [d_\tau^r (b_\tau^r + w_\tau^r - c(e_\tau^r)) + (1 - d_\tau^r) \bar{u}].$$

¹⁶Similar features of an „optimal“ exploitative contract have previously been derived in Eliaz and Spiegel (2006), or Heidhues and Kőszegi (2010).

For naive types, this real utility may be different from their perceived payoff, namely if he expects to choose the virtual contract in the future. A naive agent's perceived payoff – consisting of real current and virtual future payoffs – is indicated with the superscript “ rv ”.

$$U_t^{rv} = d_t^r (b_t^r + w_t^r - c(e_t^r)) + (1 - d_t^r) \bar{u} + \beta \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} [d_\tau^v (b_\tau^v + w_\tau^v - c(e_\tau^v)) + (1 - d_\tau^v) \bar{u}].$$

The principal's payoff in any period t in a setting where she never declares bankruptcy is

$$\Pi_t^r = \sum_{\tau=t}^{\infty} \delta^{\tau-t} [d_\tau^r (e_\tau^r \theta - b_\tau^r - w_\tau^r) + (1 - d_\tau^r) \bar{\pi}].$$

However, a naive agent wanting to choose the virtual contract in the future expects the principal to maximize

$$\Pi_t^{rv} = d_t^r (e_t^r \theta - b_t^r - w_t^r) + (1 - d_t^r) \bar{\pi} + \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} [d_\tau^v (e_\tau^v \theta - b_\tau^v - w_\tau^v) + (1 - d_\tau^v) \bar{\pi}].$$

For later use, we further define perceived future virtual payoffs,

$$U_t^v = \sum_{\tau=t}^{\infty} \delta^{\tau-t} [d_\tau^v (b_\tau^v + w_\tau^v - c(e_\tau^v)) + (1 - d_\tau^v) \bar{u}]$$

for the agent, and for the principal:

$$\Pi_t^v = \sum_{\tau=t}^{\infty} \delta^{\tau-t} [d_\tau^v (e_\tau^v \theta - b_\tau^v - w_\tau^v) + (1 - d_\tau^v) \bar{\pi}].$$

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 \mathbf{C} in order to maximize Π_0^r , and d_t^P in order to maximize Π_t^r , taking into account the agent's actual, i.e., “real”, behavior. The (partially) naive agent, though, expects the principal to choose d_t^P in order to maximize Π_t^{rv} or Π_t^v , respectively.

The (partially) naive agent makes optimal choices given her current preferences and her perceptions of future behavior, i.e., expecting to have different time preferences in the future. Hence, he chooses d_t , \hat{i}_t and e_t in order to maximize U_t^{rv} , given that an agent believes to choose the virtual contract rather than the real contract in the future.

4 Benchmarks: Time-Consistent and Sophisticated Agent

We first derive two benchmarks, profit-maximizing contracts with time-consistent and with sophisticated present-biased agents.

Time-consistent agent Consider a time-consistent agent, i.e., an agent who has $\beta = \hat{\beta} = 1$. Since the agent's effort is verifiable, effectively there is no agency problem that must be addressed. In this case, the principal does not need to make use of her ability to make long-term commitments. Instead, she can use a series of spot contracts in order to make the agent choose surplus-maximizing effort in every period, and collect the whole surplus for herself.

One possibility to generate such an outcome is to offer the following contract in every period t :

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

The agent always accepts such a contract and exerts effort e^{FB} . Since the principal extracts the whole surplus, she has no incentives to declare bankruptcy. The optimal menu of contracts with a time-consistent agent hence maximizes the surplus and holds the agent down to his outside option.

Sophisticated present-biased agent A sophisticated present-biased agent (i.e., $\hat{\beta} = \beta$), is aware of his future preferences and hence of the choices he is going to make in the future. Then, as with a time-consistent agent, it is sufficient to let \mathbf{C} consist of only one element in every period. Generally, the principal has the opportunity to reward the agent for effort with a bonus paid at the end of the period, or with the promise of a higher continuation payoff. Since $\beta < 1$, though, the agent effectively discounts the future at a higher rate than the principal does. Therefore, it is optimal for the principal to as above offer the following contract in every period t :

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

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 this case, the agent would simply walk away after this period. Concluding, time-inconsistencies have no impact on the optimal contract if the agent is sophisticated about his present bias. This is because the production technology is effectively static in the sense that there are no technological linkages across periods, because the agent can immediately be compensated for his effort, and because effort is verifiable.

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 whose $\hat{\beta} = 1$. In Subsection 6.6, we show that the results derived here remain unaffected if $\hat{\beta} \in (\beta, 1)$. Note that as with a sophisticated agent, the possibility of writing formal spot contracts to motivate the agent implies that the existence of a present bias does not automatically trigger inefficiencies. The same spot contract offered to a sophisticated agent could also be offered to naive present-biased agents, yielding exactly the same outcome. However, now the principal can design a menu of long-term contracts which can exploit the naive agent’s misperception of his future behavior. She is able to include elements into \mathbf{C} that seem optimal for the agent from the perspective of earlier periods – but that are not chosen once the agent actually faces the respective choice. As discussed above, we can without loss of generality restrict \mathbf{C} to consist of exactly two elements in every period: The contract the agent actually chooses and the contract the agent had planned to choose from the perspective of earlier periods. We call the former contract the *real contract* and add the superscript “ r ” to all its components and the latter contract the *virtual contract* and add the superscript “ v ” to all its components. Both contracts can be contingent on the full history of the game, h^t :

$$C(h^t) = \{C^r(h^t), C^v(h^t)\} = \{(w^r(h^t), b^r(e_t^r, h^t)), (w^v(h^t), b^v(e_t^v, h^t))\}.$$

Hence, \mathbf{C} also gives values $U^r(h^t)$, $U^{rv}(h^t)$, $U^v(h^t)$, $\Pi^r(h^t)$, $\Pi^{rv}(h^t)$ and $\Pi^v(h^t)$ for every history h^t .

In the following, we focus on contracts that, on the equilibrium path, have $d_t^r = d_t^v = d_t^P = 1$ in all periods t . Hence, the principal never declares bankruptcy, and the agent always accepts the employment offer. Any (temporary or permanent) termination of the relationship can never be optimal – simply because the principal could always include the optimal contract for time-consistent and sophisticated agent in the long-term contract offered to the naive agent, which the latter would accept if the alternative was $d_t = 0$.

The elements of the long-term contract \mathbf{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 (compared to reject all of them for the respective period and consume his outside option instead). 2)

Individual rationality constraints for the principal (IRP) that keep her from declaring bankruptcy. 3) **Incentive compatibility constraints** for the agent (IC) that induce him to select the principal's preferred effort level. These constraints must also hold for the virtual contracts, i.e., for future histories that never materialize. Finally, 4) **Selection constraints** ensure that the agent keeps **choosing** the real contract in every period, while still **intending** to choose the virtual contract in all future periods.

Individual rationality constraints for the agent For every history h^t , it must be optimal for the agent to accept the real contract (expecting to choose equilibrium effort e_t^r and to select the virtual contract in all future periods), compared to rejecting any contract and consuming \bar{u} in the respective period:

$$w^r(h^t) + b^r(h^t, e_t^r) - c(e^r) + \beta\delta U^v(h^t \cup \{d_t = 1\}) \geq \bar{u} + \beta\delta U^v(h^t \cup \{d_t = 0\}). \quad (\text{rIRA})$$

Note that when stating $U^v(\cdot)$, we only include the elements that are relevant for the respective constraint (and do the same when describing other constraints), assuming that all other elements are chosen as prescribed by play on the equilibrium path.

Furthermore, \mathbf{C} has to be such that the agent expects to accept a contract in all future periods, i.e., for all h^t ,

$$U^v(h^t \cup \{d_t = 1\}) \geq U^v(h^t \cup \{d_t = 0\}) \quad (\text{vIRA})$$

has to hold.¹⁷ However, an individual rationality constraint is not needed for U^r because the agent does not expect to select the real contract in future periods. I.e.,

$$w^r(h^t) + b^r(h^t, e_t^r) - c(e^r) + \beta\delta U^r(h^t \cup \{d_t = 1\}) \geq \bar{u} + \beta\delta U^r(h^t \cup \{d_t = 0\})$$

does *not* have to hold. In fact, this constraint is generally violated and the agent effectively will be exploited and receive less than his outside option.

¹⁷Note that vIRA also implies $U^v(h^t) \geq \frac{\bar{u}}{1-\delta}$ for all histories.

Individual rationality constraints for the principal For every history h^t , the following constraints must be satisfied for the principal:

$$\Pi^{rv}(h^t) \geq \frac{\bar{\pi}}{1 - \delta} \quad (\text{vrIRP})$$

$$\Pi^v(h^t) \geq \frac{\bar{\pi}}{1 - \delta} \quad (\text{vIRP})$$

If either of these constraints was not satisfied, the agent would expect the principal to not honor her obligations in the virtual contract and instead shut down. A real IR constraint for the principal, $\Pi^r(h^t) \geq \bar{\pi}/(1 - \delta)$, is not required because of the agent's misperceptions about his own behavior. However, it is obvious that shutting down will never be optimal for the principal because she could at any point offer the optimal contract for the time-consistent or sophisticated agent and thereby collect a rent that is strictly positive.

Incentive compatibility constraints It has to be in the agent's interest to choose equilibrium effort e_t^r , given the compensation he receives today, and given his expectation of future (virtual) payoffs. This gives rise to a real incentive compatibility constraint,

$$-c(e_t^r) + b^r(h^t, e_t^r) + \beta\delta U^v(h^t \cup \{e_t^r\}) \geq -c(\tilde{e}_t^r) + b^r(h^t, \tilde{e}_t^r) + \beta\delta U^v(h^t \cup \{\tilde{e}_t^r\}), \quad (\text{rIC})$$

which has to hold for all histories h^t and effort levels \tilde{e}_t in a given contract menu.

The agent also has to expect to select on-path effort levels in the virtual contract, which for all histories h^t and effort levels \tilde{e}_t gives

$$-c(e_t^v) + b^v(h^t, e_t^v) + \beta\delta U^v(h^t \cup \{e_t^v\}) \geq -c(\tilde{e}_t^r) + b^r(h^t, \tilde{e}_t^r) + \beta\delta U^v(h^t \cup \{\tilde{e}_t^r\}). \quad (\text{vIC})$$

Note that because effort is verifiable, one could simply assume infinitely negative payments in case the agent does not choose equilibrium effort. We take the current general formulation because it allows us to easily extend our setup to limited liability and moral hazard.

Selection constraints Finally, the agent has to select the real contract in every period, however expect to select the virtual contract in the future. For every history h^t , this yields

the constraints

$$w^r(h^t)+b^r(h^t, e^r)-c(e_t^r)+\beta\delta U^v\left(h^t\cup\{\hat{i}_t=r\}\right)\geq w^v(h^t)+b^v(h^t, e^v)-c(e_t^v)+\beta\delta U^v\left(h^t\cup\{\hat{i}_t=v\}\right) \quad (\text{rC})$$

and

$$w^v(h^t)+b^v(h^t, e^v)-c(e_t^v)+\delta U^v\left(h^t\cup\{\hat{i}_t=v\}\right)\geq w^r(h^t)+b^r(h^t, e^r)-c(e_t^r)+\delta U^v\left(h^t\cup\{\hat{i}_t=r\}\right). \quad (\text{vC})$$

Note that due to naivete, β does not feature in (vC).

Objective The principal's objective is to offer a long-term menu of contracts $\mathcal{C} = \{C^r(h^t), C^v(h^t)\}$ for all $h^t \in H$ that maximizes $\Pi^r(h^0)$, subject to (rIRA), (vIRA), (rIRP), (vIRP), (rIC), (vIC), (rC) and (vC) that must hold for all histories. In the following, we first simplify the problem and then characterize an optimal long-term menu of contracts, \mathcal{C} .

5.2 Profit-Maximizing Contract

In this section, we fully characterize a profit-maximizing menu of contracts \mathcal{C} .

As we show in Appendix B, the real contract as well as the virtual contract (except for the first period of it) are stationary. 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 e^r, w^r, b^r , the virtual contract of e_1^v, w_1^v, b_1^v for the first and e^v, w^v, b^v for all subsequent periods after it has been selected. To exploit the agent's present bias and keep him below his outside option, the virtual contract must be different in the first period it is chosen: In order to extract rents from the agent in the real contract, the principal shifts as much as possible of the compensation promised to the agent 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 to be sufficiently unattractive for the agent – but sufficiently attractive to expect him to still opt for C^v in future periods. In Appendix B (Lemma B6), we further show that $-c(e^r) + b^r < \bar{u}$, which implies that the naive agent is exploited by the principal. This is because principal and agent have non-common priors concerning the

agents future time preferences. 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 therefore focus on the exact 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).

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

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

The first bullet point of Lemma 1 indicates that the principal promises the whole virtual rent to the agent (binding (vIRP)), with the exception of the first virtual period which is constructed to be just sufficiently unattractive to be never selected by the agent (binding (rC)). Here, we also see that if the principal was able to credibly make arbitrary promises and not able to shut down (i.e., if she did not face (IR) constraints), we would not have an 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. Furthermore, even though the agent expects to get a rent in the future, his expected 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.

Finally, the real bonus b^r captures the link between real and virtual compensation. The first elements, $c(e^{FB}) + \bar{u}$, capture the agent's effort costs and outside option, hence would constitute his "fair" compensation. The last term characterizes the principal's savings 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,

which is supposed to serve as the reward for *today's* effort. If this value is rather high, b^r can actually be negative - in this case, the agent has to make a payment to the principal in every period in order to keep the option of receiving future virtual rents. Note that, as we show below in Section 6.2, even in the arguably more realistic case of limited liability the agent would not be better off. In such a setting, the principal will respond to a binding limited liability constraint by requiring real effort to be above e^{FB} . While this leaves the agent's real utility unaffected, it entails an allocative inefficiency and allows the principal to extract only a lower rent from the agent compared to the case absent the binding limited liability constraint.

Real payoffs of principal and agent are characterized in Proposition 1.

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

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

and

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

respectively.

The principal uses the promise of a virtual future surplus (which we interpret as the perceived attractiveness of making a career) to keep the agent below his outside option while still accepting the contract - which is perceived to be optimal by the agent as he expects to earn a rent in the future from choosing the virtual contract - and working for the principal in every period.

Proposition 1 also 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). Put differently, agents with a medium inclination to procrastinate can be most effectively exploited. Hence, these types 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

Consequently, our model predicts that firms benefit from offering long-term contracts – as compared to a sequence of spot contracts – to agents with an intermediate time-inconsistency, while there are no benefits from offering long-term contracts to other agents. If there were some additional costs of long-term contracts, e.g., because they are more complicated to write, then agents with an intermediate degree of time-inconsistency would be more likely to receive long-term contracts than other agents.

Finally, we take a brief look at the first period of the virtual contract, which is constructed as something like a probation phase in order 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. 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 Section 6.4. Using the respective binding constraints gives

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

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 would adjust the required level of e_1^v accordingly in order to keep the agent from selecting the virtual contract.

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

Note that a high relationship surplus, as well as a high long-run discount factor δ , is bad for the naive agent. 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. The agent also benefits if the principal’s outside option is higher because then, she can only commit to

lower future benefits for the agent. As the agent expects to earn less in the future, he wants to earn more today and can be exploited less. Somewhat counter-intuitively, this argument implies that employment protection, associated with increased layoff costs (corresponding to a reduced outside option), might hurt the agent. This makes it easier for the principal to commit to future employment and hence improves the scope for exploitation.

Corollary 1. *The agent's utility is increasing in $\bar{\pi}$, the principal's outside option.*

6 Variations of the Model

In this section, we look at variations of the original model. Besides serving as a robustness test, those 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 when we consider a model with a finite time horizon.

Proposition 2. *For a finite time horizon T , a profit-maximizing contract has the following features for all periods $t \leq T$:*

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

Note that for $T \rightarrow \infty$, this expression approaches the result of the infinite-horizon case. The more periods there are left, the larger are the total future benefits the principal can (“virtually”) promise to the agent. Hence, the worker is willing to accept lower payments today, although he does not actually choose the virtual contract in the future.

At later stages of a career, i.e., in periods closer to retirement, less future benefits can be promised to the agent. In fact, this protects the agent as he cannot be exploited as much.

In combination with a limited liability constraint, which we will analyze in Section 6.2, younger workers are also more likely to work harder – because a binding limited liability constraint induces the principal to optimally implement an inefficiently high effort level in order to exploit the agent as much as possible. This finding broadly resembles actual labor market patterns where we find age earning profiles with relatively better paid but less productive older workers.

Corollary 2. *For a finite time horizon T , the optimal b_t^r is decreasing in $T - t$. That is, the principal can exploit the agent more if he is in the early stages of his career.*

6.2 Limited Liability

We have shown that without a lower bound on payments, b^r can actually be negative, indicating effective payments from agent to 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 ($b^r \geq 0$). Interestingly (and different from moral hazard problems with limited liability where the agent generally gets a rent), the agent is not better off than before but receives exactly the same level of real utility. This is because the principal optimally responds to a binding limited liability constraint by letting the agent work harder and setting e^r above e^{FB} to extract additional rents. Hence, we might face a situation where the present-biased agent works harder than a sophisticated or time-consistent one.

Summing up, the profit-maximizing contract when payments have to be non-negative is characterized in Proposition 3.

Proposition 3. *Assume $b, w \geq 0$ in every period t . Then, a profit-maximizing contract \mathbf{C} has*

$$e^r \geq e^{FB}, e^v = e^{FB} \text{ and } b^r = \max \left\{ 0, c(e^{FB}) + \bar{u} - \beta(1 - \beta) \frac{\delta^2}{1 - \delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) \right\}.$$

Moreover,

- if $c(e^{FB}) + \bar{u} - \beta(1 - \beta) \frac{\delta^2}{1 - \delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) < 0$ holds, the limited liability constraint for the agent's real bonus binds. Then $e^r > e^{FB}$ and is chosen such that $b^r = c(e^r) + \bar{u} - \beta(1 - \beta) \frac{\delta^2}{1 - \delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) = 0$.
- If $c(e^{FB}) + \bar{u} - \beta(1 - \beta) \frac{\delta^2}{1 - \delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) \geq 0$, the limited liability constraints do not bind and hence do not affect \mathbf{C} .

- In either case, the agent's real payoff is $u^r - \bar{u} = -\beta(1-\beta)\frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u})$.

The agent's real payoff is unaffected by him being protected by limited liability. Hence, the full burden of the inefficient outcome is borne by the principal: Because an effort above e^{FB} reduces total surplus, she can extract less from the agent than when payments are not restricted. While under a binding limited liability constraint for the agent's real bonus, the principal's net profit is $\pi^r - \bar{\pi} = e^r\theta - \bar{\pi}$, it is $\pi^r - \bar{\pi} = (e^{FB}\theta - c(e^{FB}) - \bar{u} - \bar{\pi}) \left(1 + \beta(1-\beta)\frac{\delta^2}{1-\delta}\right)$ if no limited liability constraints bind.

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 $\bar{u}, \bar{\pi}$, there exists a $\underline{\theta}$ such that for all $\theta > \underline{\theta}$, $e^r > e^{FB}$. Moreover, for all θ , $\frac{de^r}{d\theta} \geq 0$.*

A more productive employment relationship (characterized by a higher value of θ) generally triggers larger inefficiencies. Besides making it more likely that the limited liability constraint binds, increasing the productivity of an employment relationship then further exacerbates the inefficiencies caused by excessive real effort levels.

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

6.3 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. Then the agent accepts any contract that – in his expectation – at least gives him his outside option. However, in many labor market settings firms compete for agents. This will generally allow agents to extract a share of the relationship surplus. In this variant of the model, we explore whether our results hold up if the principal does not have full bargaining power. We assume that players (potentially) bargain about how to allocate the relationship surplus at the beginning of every period. 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 in expectation leaves him with at least $1 - \alpha$ of his “present-biased view” of the total relationship surplus (note that this is equivalent to assuming that players bargain about the per-period relationship surplus). We interpret α as a proxy for the competitiveness of the labor market. A higher α , i.e., a higher share of the surplus for the principal, corresponds to less competition for agents.

We show that unless $\alpha = 0$, i.e., the agent has full bargaining power, 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. Therefore, competition on the labor market helps agents with self-control problems. Still, competition does not cause our results to disappear. This is only the case if competition for the agent’s labor is perfect and frictionless.

Importantly, because the principal can commit to long-term contracts, she is able to back-load the agent’s compensation and hence can commit *not* to renegotiate the virtual contract (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). However, because the agent can *not* commit not to renegotiate an agreement, a front-loading of the agent’s compensation is not feasible – although it might be optimal under some instances given the agent’s lower effective discount factor.

Now, the (net) surplus in a given period t amounts to $e_t\theta - c(e_t) - \bar{u} - \bar{\pi}$. 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 an employment relationship in the given period 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 $\bar{u} + (1 - \alpha)(e^r\theta - c(e^r) - \bar{u} - \bar{\pi}) + \beta\delta \{\bar{u} + (1 - \alpha)(e_1^v\theta - c(e_1^v) - \bar{u} - \bar{\pi}) + \delta [\bar{u} + (1 - \alpha)(e^v\theta - c(e^v) - \bar{u} - \bar{\pi})] / (1 - \delta)\}$.

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, $e^r = e^v = e^{FB}$. In addition, the agent’s

utility in the first period of the virtual contract, u_1^v , must be sufficiently low such that the agent will eventually 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 .

Still, the relevant constraints which pin down payments and utilities are (rIRA) and (rC). We also have to take into account that the agent's real (IC) constraint looks different because once he deviates from selecting equilibrium effort, he cannot be punished by receiving his outside option from then on. Instead, he would be able to negotiate a new contract, expecting a share $1 - \alpha$ of the relationship surplus. Therefore, when deviating, the agent only sacrifices this period's bonus payment and the option to select the virtual contract in the subsequent period. However, it can be shown that the (rIC) constraint is automatically implied by the (rIRA) constraint.

This yields the following (relevant) constraints:

$$\begin{aligned} w^r + b^r - c(e^{FB}) + \beta\delta \left(w_1^v + b_1^v - c(e^{FB}) + \delta \frac{w^v + b^v - c(e^{FB})}{1 - \delta} \right) \\ \geq \bar{u} + (1 - \alpha) (e^{FB}\theta - c(e^{FB}) - \bar{u} - \bar{\pi}) + \beta\delta \frac{[\bar{u} + (1 - \alpha) (e^{FB}\theta - c(e^{FB}) - \bar{u} - \bar{\pi})]}{1 - \delta} \end{aligned} \quad (\text{rIRA})$$

and

$$\begin{aligned} w^r + b^r - c(e^{FB}) + \beta\delta \left(w_1^v + b_1^v - c(e^{FB}) + \delta \frac{w^v + b^v - c(e^{FB})}{1 - \delta} \right) \\ \geq w_1^v + b_1^v - c(e^{FB}) + \frac{\delta\beta}{1 - \delta} (w^v + b^v - c(e^{FB})). \end{aligned} \quad (\text{rC})$$

Those will bind for the same reasons as above and pin down the values of the agent's real utility, as well as his utility in the first period of the virtual contract:

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

and

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

Therefore, the agent receives less than his fair share, which would amount to $\bar{u} + (1 - \alpha)(e^{FB}\theta - c(e^{FB}) - \bar{u} - \bar{\pi})$ in every period. Furthermore, the agent's real payoff is decreasing in α , which also means that the agent benefits from more competition on the labor market. 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.

Finally, the above analysis delivers one further interesting takeaway: Recall that in our benchmark model (which corresponds to $\alpha = 1$), the agent is better off if the principal's outside option is higher. This comparative static reverses if α is small (more precisely, if $\alpha \leq (1 - \delta) / [1 - \delta + (1 - \beta)\beta\delta^2]$). Then, the agent generally 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. If his bargaining power is high, the consequences of more restrictive employment protection laws rather are positive for the agent. This implies that even the *qualitative* effects of prominent labor market regulations crucially depend on the competitiveness of the labor market.

6.4 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. Moreover, 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 constraints (rIRA) and (rC) hold with equality, i.e.

$$b^{r'} = c(e^{r'}) + \bar{u} + \frac{\beta'\delta}{1-\delta}\bar{u} - \delta(\beta'b_1^{v'} - \beta c(e_1^{v'})) - \frac{\delta^2}{1-\delta}(\beta'b^{v'} - \beta c(e^{v'})) \quad (1)$$

and

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

These equations together imply

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

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 $b_1^{v'}$ and $c(e_1^{v'})$ are unique. As one can see from equation (1), the principal benefits from choosing them as high as possible. Consequently,

$$b_1^{v'} = \theta e_1^{v'} - \bar{\pi}.$$

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

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

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 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.

6.5 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_t is the agent's private information. Output remains verifiable, but now amounts to $y_t \in \{0, \theta\}$, with $\text{Prob}(y_t = \theta) = e_t$. 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 $\lim_{e_t \rightarrow 1} c' = \infty$. This implies that

first-best effort e^{FB} is still characterized by $\theta - c'(e^{FB}) = 0$. For concreteness we also set $\bar{u} = \bar{\pi} = 0$.

First, note that the optimal contract for a sophisticated agent is the same as for a time-consistent agent: 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 $w = \bar{u} + c(e^{FB}) - e^{FB}\theta$. 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).¹⁸

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

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

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

$$e^r \in \operatorname{argmax} [e^r b^r - c(e^r)]. \tag{rIC}$$

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

¹⁸Even 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.

$b_1^v - c'(e_1^v) = 0$ 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 $u^v = e^{FB}\theta - c(e^{FB}) - \bar{\pi}$). 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 $U^{rv} \geq \bar{u} + \beta\delta\frac{\bar{u}}{1-\delta}$ (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

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

and

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

Having both constraints bind yields

$$w^r = \bar{u} - \delta^2\beta(1-\beta)\frac{(e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u})}{1-\delta} \quad \text{and} \quad u_1^v = \bar{u} - \beta\delta\frac{(e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u})}{1-\delta}.$$

Concluding, the agent still is exploited in case a moral hazard is present, 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.

6.6 Partial Naivete

So far we only considered the extreme cases of completely naive and fully sophisticated agents. Now 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:

$$-c(e_1^v) + b_1^v + \hat{\beta} \frac{\delta}{1-\delta} (w^v + b^v - c(e^v) - \bar{u}) \geq 0, \quad (\text{v1IC})$$

$$-c(e^v) + b^v + \hat{\beta} \frac{\delta}{1-\delta} (w^v - c(e^v) + b^v - \bar{u}) \geq 0, \quad (\text{vIC})$$

and

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

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

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

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

$$(b_1^v - c(e_1^v) - \bar{u}) + \frac{\delta}{(1-\delta)} (b^v - c(e^v) - \bar{u}) \hat{\beta} \frac{(1-\delta)}{(1-\hat{\beta}\delta)} \geq \frac{(b^r - c(e^r) - \bar{u})}{(1-\hat{\beta}\delta)}, \quad (\text{vC})$$

$$e^v \theta - b^v \geq \bar{\pi}. \quad (\text{vIRP})$$

We show in Appendix B 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.

Proposition 4. *The partially naive agent receives the same contract as the fully naive agent.*

A (partially) naive agent expects that in the future choosing the virtual contract is strictly better than choosing the real contract. In fact, tomorrow he will be just indifferent, though. For this decision, his true β is relevant but his belief $\hat{\beta}$ is not. 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 the naive agents, no matter whether they are fully or only partially naive, and the first-best contract for the sophisticated and the rational agents.

Such a discontinuity in the form of contracts is a common feature among papers 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).

6.7 Learning

In our main model, we assume that the agent fails over and over again: In every period he selects the real contract, although before he planned to take up the virtual contract in the current period. This assumption seems quite strong. One might expect that at least after a couple of failed attempts to actually select the virtual contract, the agent should realize that he has a present bias (for fully naive agents) or that his present bias is stronger than he thought (for partially naive agents). In the following, we show that even if the agent becomes aware of this, as long as he does not become fully sophisticated, the principal will offer exactly the same contract.

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 4. *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 4.

If the agent has the chance to learn his true β , he will choose the first-best contract from then on, which can always be added to the menu of contracts by the principal, but will only be chosen by a sophisticated agent.

Corollary 5. *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.

6.8 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 it is not clear whether the principal necessarily knows an agent's type, i.e., his exact degree of the present bias. We show 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 one. 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.

Formally, 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.¹⁹ 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 $\tilde{\mathbf{C}}_i$.

Now, let \mathbf{C}_1 and \mathbf{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 \mathbf{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 \mathbf{C}_1 . Furthermore, since \mathbf{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 \mathbf{C}_1 .

Hence, either $\tilde{\mathbf{C}}_2$ must be constructed in a way that makes it optimal for agent 2 to choose it (keeping $\tilde{\mathbf{C}}_1 = \mathbf{C}_1$), or $\tilde{\mathbf{C}}_1$ must be made sufficiently unattractive for 2 (keeping $\tilde{\mathbf{C}}_2 = \mathbf{C}_2$). How exactly the menus are adjusted is described in the Appendix, the main result is given in Proposition 5. 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 intermediate utility.

Proposition 5. *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

¹⁹Note 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.

in Appendix C 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 $\tilde{\mathbf{C}}_1 = \mathbf{C}_1$ 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 time-consistent 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.²⁰

7 Conclusion

In this paper we have shown how a principal should optimally contract with a present-biased agent in a long-term relationship absent moral hazard. The principal can take advantage of a naive present-biased agent by promising career advancement and career prospects that involve small short term losses. 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 actually gets to enjoy the generous benefits from the virtual contract.

These findings are robust to imposing limited liability on the agent, taking moral hazard into account, varying the degree of labor market competition, letting the agent have differently strong present biases in different domains, considering the case of partial naivete, allowing the agent to learn about his bias, and considering a finite time horizon.

A number of our results appear at odds with usual findings or basic intuition in models with time-inconsistent agents. First, in our model, the time-inconsistent agent might work harder than his time-consistent counterpart, while in many other settings previously studied in the literature he appears more lethargic and lazy. This is driven by the limited liability constraint. As the agent cannot receive a negative compensation, the only way to exploit him further is to let him work inefficiently hard. Note, however, that – very much in

²⁰Other papers analyze screening time-inconsistent agents: First, Heidhues and Köszegi (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.

line with the standard intuition – the agent very much procrastinates as he indefinitely postpones the investment to take up the virtual contract and actually “make a career”.

Second, a higher outside option for the principal may actually benefit the agent. The reason for this is that a higher outside option makes it harder for the principal to credibly commit to employing the agent in the future and providing him with generous benefits. Therefore, the promised future benefit must be lower, making the virtual contract appear less desirable. Thus the agent can be exploited less today. This only holds, though, if competition for agents is weak. If the market for agents is competitive, those also benefit from a higher (perceived) future surplus, that is, a lower outside option of the principal.

Third, our results indicate that the effect of employment protection regulations on agents depends on the competitiveness of labor markets. Inverting the preceding argument, employment protection increases the principal’s commitment to future employment, allowing her to credibly promise a higher compensation in the future – which she will never have to pay as the agent will always choose the real over the virtual contract. If competition for agents is weak, the agent is then willing to accept a lower payment or work harder today. On the other hand, if competition for agents is intense, those also benefit from an increased commitment of the principal. Then, the positive effects of more competition and more restrictive employment protection on an agent’s real utility complement each other.

References

- 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.
- 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.
- (2008): “Consumer Optimism and Price Discrimination,” *Theoretical Economics*, 3, 459–497.
- FAHN, M. (2017): “Minimum Wages and Relational Contracts,” *Journal of Law, Economics, and Organization*, 33(2), 301–331.
- FONG, Y., AND J. LI (2017): “Relational Contracts, Limited Liability, and Employment Dynamics,” *Journal of Economic Theory*, 169, 270–293.
- FUDENBERG, D., B. HOLMSTROM, AND P. MILGROM (1990): “Short-Term Contracts and Long Term Agency Relationships,” *Journal of Economic Theory*, 51, 1–31.
- GALPERTI, S. (2015): “Commitment, Flexibility, and Optimal Screening of Time Inconsistency,” *Econometrica*, 83(4), 1425–1465.
- GARAVAN, T. N., AND M. MORLEY (1997): “The socialization of high-potential graduates into the organization,” *Journal of Managerial Psychology*, 12(2), 118–137.
- GERVAIS, SIMON, J. H., 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.
- HARRIS, C., AND D. LAIBSON (2001): “Dynamic Choices of Hyperbolic Consumers,” *Econometrica*, 69(4), 935–957.
- HEIDHUES, P., AND B. KŐSZEGI (2010): “Exploiting Naivete 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*.
- HOLMSTROM, B. (1999): “Managerial Incentives Problems - A Dynamic Perspective,” *Review of Economic Studies*, 66, 169–182.
- 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 (2003): “Consumption–Savings Decisions with Quasi–Geometric Discounting,” *Econometrica*, 71(1).
- LAIBSON, D. (1997): “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, 112(2), 443–477.
- 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*.
- MALCOMSON, J., AND F. SPINNEWYN (1988): “The Multiperiod Principal-Agent Problem,” *Review of Economic Studies*, 55, 391–408.
- 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.
- RINDFUSS, R. R., E. C. COOKSEY, AND R. L. SUTTERLIN (1999): “Young Adult Occupational Achievement: Early Expectations Versus Behavioral Reality,” *Work and Occupations*, 26(2), 220 – 263.
- 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*.
- 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.

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. *Without loss of generality, the following holds for all histories:*

- $U^v(h^t \cup \{d_t = 0\}) = \frac{\bar{u}}{1-\delta}$
- *If the agent selects effort $\tilde{e}_t \neq e_t^v$ in the virtual contract, then $b^v(h^t, \tilde{e}_t) \leq 0$ and $U^v(h^t \cup \{\tilde{e}_t\}) = \frac{\bar{u}}{1-\delta}$*
- *If the agent selects $\tilde{e}_t \neq e_t^r$ in the real contract, then $b^r(h^t, \tilde{e}_t) \leq 0$ and $U^v(h^t \cup \{\tilde{e}_t\}) = \frac{\bar{u}}{1-\delta}$.*

Proof. The optimality of $b^v(h^t, \tilde{e}_t) \leq 0$ for $\tilde{e}_t \neq e_t^v$ and $b^r(h^t, \tilde{e}_t) \leq 0$ for $\tilde{e}_t \neq e_t^r$ is straightforward because lower values of $b^r(h^t, \tilde{e}_t)$ and $b^v(h^t, \tilde{e}_t)$ relax (IC) constraints.

Concerning continuation values after the agent either chose $d_t = 0$ or did not exert the equilibrium effort level, assume that, for example, there is a history h^t with $U^v(h^t \cup \{d_t = 0\}) > \frac{\bar{u}}{1-\delta}$. Replace the contract following the history $h^t \cup \{d_t = 0\}$ with the following stationary contract for all subsequent histories: $e^r = e^v = e^{FB}$, $w^r = w^v = \bar{u}$, $b^r = b^v = c(e^{FB})$. This relaxes (IR) and (IC) constraints for history h^t , and keeps all off-equilibrium constraints following the history $h^t \cup \{d_t = 0\}$ satisfied. The same can be done for all other cases. \square

The next Lemma also simplifies the analysis.

Lemma B2. *Without loss of generality, we can set $w^r(h^t) = w^v(h^t) = 0$ for all histories h^t .*

Proof. Assume $w^v(h^t) > 0$. Reducing $w^v(h^t)$ by ε and increasing $b^v(h^t)$ by $\varepsilon > 0$ does not tighten any constraint, but relaxes (vIC). Assume $w^r(h^t) > 0$. Reducing $w^r(h^t)$ by $\varepsilon > 0$ and increasing $b^r(h^t)$ by ε does not tighten any constraint but relaxes (rIC). \square

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

Now, the remaining constraints are

$$b^r(h^t, e_t^r) - c(e_t^r) + \beta\delta U^v(h^t \cup \{d_t = 1\}) \geq \bar{u} + \beta\delta \frac{\bar{u}}{1-\delta}. \quad (\text{rIRA})$$

$$U^v(h^t \cup \{d_t = 1\}) \geq \frac{\bar{u}}{1-\delta} \quad (\text{vIRA})$$

$$\Pi^{rv}(h^t) \geq \frac{\bar{\pi}}{1-\delta} \quad (\text{vrIRP})$$

$$\Pi^v(h^t) \geq \frac{\bar{\pi}}{1 - \delta} \quad (\text{vIRP})$$

$$b^r(h^t, e_t^r) - c(e_t^r) + \beta \delta U^v \left(h^t \cup \left\{ \hat{i}_t = r \right\} \right) \geq b^v(h^t, e_t^v) - c(e^v) + \beta \delta U^v \left(h^t \cup \left\{ \hat{i}_t = v \right\} \right) \quad (\text{rC})$$

$$b^v(h^t, e_t^v) - c(e^v) + \delta U^v \left(h^t \cup \left\{ \hat{i}_t = v \right\} \right) \geq b^r(h^t, e_t^r) - c(e^r) + \delta U^v \left(h^t \cup \left\{ \hat{i}_t = r \right\} \right) \quad (\text{vC})$$

In the next two Lemmas, we derive the structure of \mathbf{C} and show that the real contract is stationary, as well as the virtual contract with the exception of the first period where it is chosen.

Lemma B3. $b^r(h^t, e^r) - c(e_t^r)$ is the same for all histories h^t .

Proof. Assume there are two histories $h^{\bar{t}}$ and $h^{\underline{t}}$ such that the agent's real net payoff differs for both histories, and without loss of generality assume that $b^r(h^{\bar{t}}, e_{\bar{t}}^r) - c(e_{\bar{t}}^r) > b^r(h^{\underline{t}}, e_{\underline{t}}^r) - c(e_{\underline{t}}^r)$. Then, replace $b^r(h^{\bar{t}}, e_{\bar{t}}^r)$ by $c(e_{\bar{t}}^r) + b^r(h^{\underline{t}}, e_{\bar{t}}^r) - c(e_{\bar{t}}^r)$, as well as the virtual contract following the history $h^{\bar{t}}$ by the virtual contract following the history $h^{\underline{t}}$. This increases the principal's real profits, without affecting any constraint for any other history – also not in earlier periods, because there the agent does not expect to choose the real contract in any future period. Hence, the agent's incentives in earlier periods are not affected by what happens if he chooses the real contract in a future periods. \square

Therefore, the real contract is history-independent, and we can omit dependence on time and histories when describing the elements of C^r .

In a next step, we show that the virtual contract can be independent of calendar time in a sense that without loss of generality, its components are only contingent on the duration of the virtual contract, no matter when it has been initiated. Denote C_τ^v as the τ 's period of the virtual contract ($\tau \geq 1$), independent of the remaining components of the history of the game.

Lemma B4. Without loss of generality, the virtual contract $C^v(h^t)$ is of the form C_τ^v for all histories of the game. Furthermore, it is stationary for all $\tau \geq 2$.

Proof. First, we show that the agent's expected continuation utility after choosing the real contract in any period, $U^v \left(h^t \cup \left\{ \hat{i}_t = r \right\} \right)$ or, if using the form C_τ^v , $U_1^v(h^t)$, is the same for all histories. To the contrary, assume there are two equilibrium histories $h^{\bar{t}}$ and $h^{\underline{t}}$ where the agent's respective future virtual utility differs, i.e., assume $U_1^v(h^{\bar{t}}) \neq U_1^v(h^{\underline{t}})$ and without loss of generality $U_1^v(h^{\bar{t}}) > U_1^v(h^{\underline{t}})$. Replace all components of the history following $h^{\underline{t}}$ with the components of the history following $h^{\bar{t}}$. Then, all constraints following $h^{\bar{t}}$ remain satisfied, and $b^r - c(e^r)$ can be reduced (before, rIRA was slack at the history $h^{\bar{t}}$).

Hence, $U_1^v(h^t)$ is the same for all histories h^t , i.e., we can write U_1^v . To prove that the virtual contract can be stationary for all $\tau > 2$, note that the respective elements of C_τ^v are only relevant for the constraints (vIRA), (vIRP) and (vC). Now take the history $h^{\tilde{t}}$ where the value $U_2^v(\cdot)$ assumes its highest value. If there are histories with a lower U_2^v , replace the respective elements of the virtual contract with ones determining $U_2^v(h^{\tilde{t}})$, implying that also $U_2^v(\cdot)$ can be independent of calendar time. Finally, take the per-period value of U_2^v , $(1-\delta)U_2^v$, and set all per-period utilities for the virtual contract for $\tau \geq 2$ equal to $(1-\delta)U_2^v$. This is clearly feasible and violates no constraint. \square

These two Lemmas prove Lemma B5:

Lemma B5. *The real contract is stationary, hence independent of the history of the game. The virtual contract is independent of calendar time and is of the form C_τ^v , where τ counts the virtual contract's number of periods after it has first been chosen. Furthermore, the virtual contract is stationary for all $\tau \geq 2$.*

Hence, in the following, we omit time- and history-dependence when describing the elements of C^r . Regarding C_τ^v , we also omit dependence on τ for all $\tau \geq 2$. For $\tau = 1$, we keep the subscript "1". Therefore, the remaining constraints are

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

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

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

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

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

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

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

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

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

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

Lemma B6. *If the agent is naive and has $\beta \in (0, 1)$, then in the profit-maximizing menu of contracts, $-c(e^r) + b^r < \bar{u}$.*

Proof. First, assume $-c(e^r) + b^r > \bar{u}$. Then, (vC) and (v1IRA) imply that $(b_1^v - c(e_1^v) - \bar{u}) + \frac{\delta}{(1-\delta)} (b^v - c(e^v) - \bar{u})$ must be strictly positive as well. Change \mathbf{C} in the following way: Set $b^r - c(e^r) - \bar{u} = b_1^v - c(e_1^v) - \bar{u} = (b^v - c(e^v) - \bar{u}) = 0$, 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 $-c(e^r) + b^r = \bar{u}$. Change \mathbf{C} in the following way: First, set $(b_1^v - c(e_1^v) - \bar{u}) = (b^v - c(e^v) - \bar{u}) = 0$ 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 . \square

$-c(e^r) + b^r < \bar{u}$ 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: $e^r\theta - b^r - \bar{\pi}$ will be strictly positive in a profit-maximizing equilibrium, and (v1IRP) yields $(e_1^v\theta - b_1^v) + \frac{\delta}{1-\delta} (e^v\theta - b^v) \geq \frac{\bar{\pi}}{1-\delta}$.

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

$$(b^v - c(e^v) - \bar{u}) \geq 0 \tag{vIRA}$$

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

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

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

Lemma B7. *(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:

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

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

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. \square

C Proofs to Lemmas and Propositions from the Main Part

Proof to Lemma 1 and Proposition 1

Proof. Given the remaining constraints, the principal's maximization problem gives rise to

the following Lagrange function:

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

with first-order conditions

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

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

$$\frac{\beta \delta^2 (1 - \beta)}{(1 - \delta)(1 - \beta \delta)} (\theta - c(e^v)') = 0, \text{ i.e., } e^v = e^{FB}.$$

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 . \square

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: $e_T^v = e_{T-1}^v = e^{FB}$, $b_{T-1}^v = c(e^{FB}) + \bar{u} - \beta \delta (e^{FB} \theta - c(e^{FB}) - \bar{\pi} - \bar{u})$, $b_T^v = e^{FB} \theta - \bar{\pi}$, 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 $T-2$: $e_{T-2}^r = e^{FB}$, $b_{T-2}^r = c(e^{FB}) + \bar{u} - \beta(1-\beta)\delta^2 (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u})$. 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 $b_t^v = c(e^{FB}) + \bar{u} - \beta\delta \sum_{j=0}^{T-t-1} \delta^j (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u})$. This in turn allows the principal to exploit the agent by setting $b_t^r = c(e_t^r) + \bar{u} - \beta(1-\beta)\delta^2 \sum_{j=0}^{T-t-2} \delta^j (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u})$. \square

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:

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

with first-order conditions:

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

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 $\frac{\beta\delta^2(1-\beta)}{(1-\delta)(1-\beta\delta)} (\theta - c(e^v)') = 0$. 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. \square

Proof to Corollary 3

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

Proof to Proposition 4

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 $\frac{\delta\hat{\beta}}{1 - \delta\hat{\beta}}\varepsilon$. 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. \square

Proof to Proposition 5

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 \mathbf{C}_1 and \mathbf{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 \mathbf{C}_1 and \mathbf{C}_2 . Then agent 2's expected utility level when choosing \mathbf{C}_1 was

$\tilde{U}_2^r = \frac{\delta^2}{1 - \delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) (1 - \beta_1) (\beta_2 - \beta_1) > 0$. 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 agent 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 \mathbf{C}_2 (and expecting to select the virtual contract in the future), agent 1 gets:

$\tilde{U}_1^r = \frac{\delta^2}{1 - \delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) (1 - \beta_2) (\beta_1 - \beta_2) < 0$. Hence, agent 1 would stick to \mathbf{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 \mathbf{C}_2 in a way that it becomes more attractive for agent 2 (without making it too attractive for agent 1), or adjust \mathbf{C}_1 in a way that it becomes less attractive for agent 2.

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 where it is expected to be chosen. Hence, the principal has the following opportunities to make \mathbf{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. Then, 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 unchanged are

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

When deviating and selecting the other agent's menu, an agent's expected payoffs (and expecting to pursue the virtual contracts there) are

$$\begin{aligned} \tilde{U}_1^r &= \frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) (1 - \beta_2) (\beta_1 - \beta_2) + X_2 + \beta_1 \delta Y_2 \text{ and} \\ \tilde{U}_2^r &= \frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) (1 - \beta_1) (\beta_2 - \beta_1) + Z_1 - \frac{\beta_2}{\beta_1} Z_1. \end{aligned}$$

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

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

for agent 1 and

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

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 C8. *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 $Z_1 = \beta_1 (1 - \beta_1) \frac{\delta^2}{1 - \delta} (e^{FB} \theta - c(e^{FB}) - \bar{\pi} - \bar{u})$; costs are

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

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

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

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

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

- $s_1 > 1 - \beta_1$

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

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

costs are

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

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

$$\begin{aligned} L = & -s_1 Z_1 - (1 - s_1) (X_2 + \delta Y_2) \\ & + \lambda_{ND1} \left[\frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) (1 - \beta_2) (\beta_2 - \beta_1) - X_2 - \beta_1 \delta Y_2 \right] \\ & + \lambda_{ND2} \left[X_2 + \beta_2 \delta Y_2 + Z_1 \left(\frac{\beta_2 - \beta_1}{\beta_1} \right) - \frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) (1 - \beta_1) (\beta_2 - \beta_1) \right] \\ & + \mu_X (X_2 + \delta\beta_2 Y_2) + \mu_Y Y_2 + \mu_Z Z_1 \end{aligned}$$

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

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

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

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

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

$X_2 = \frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) (1 - \beta_1) (\beta_2 - \beta_1)$ 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

$X_2 = \frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) (1 - \beta_1) (\beta_2 - \beta_1) - \beta_2 \delta Y_2$ from binding (ND2) and plugging it into (ND1) gives

$Y_2 = \frac{\delta}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) (\beta_2 - \beta_1)$, implying that

$X_2 = \frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) (1 - \beta_1 - \beta_2) (\beta_2 - \beta_1)$. Then,

$X_2 + \delta \beta_2 Y_2 = \frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) (1 - \beta_1) (\beta_2 - \beta_1) > 0$.

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 $Z_1 = \beta_1 (1 - \beta_1) \frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u})$. 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 $Z_1 = \beta_1 (1 - \beta_1) \frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u})$. Plugging $X_2 = -\delta \beta_2 Y_2$ into (ND1) gives

$\frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) (1 - \beta_2) (\beta_2 - \beta_1) + \delta Y_2 (\beta_2 - \beta_1) = 0$, 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

$X_2 = \frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) (1 - \beta_2) (\beta_2 - \beta_1)$, and plugging this into

(ND2) gives: $Z_1 = \beta_1 (\beta_2 - \beta_1) \frac{\delta^2}{1-\delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u})$. (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 \mathbf{C}_2 and the first-best contract (intended to be chosen by agent

1).²¹ It is easy to see that this is optimal even for the actual costs of Z_1 ²² 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 \mathbf{C}_1 and \mathbf{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 beginning 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.²³ 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: $c(e^{FB}) + \bar{u} - \beta_2(1 - \beta_2)\frac{\delta^2}{1-\delta}(e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) + (1 - \beta_2)(\beta_2 - \beta_1)\frac{\delta^2}{1-\delta}(e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) \leq e^{FB}\theta - \bar{\pi}$ and $c(e^{FB}) + \bar{u} - \beta_2\frac{\delta}{1-\delta}(e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) + (\beta_2 - \beta_1)\frac{\delta}{1-\delta}(e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) \leq e^{FB}\theta - \bar{\pi}$. 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, \mathbf{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

²¹To see this, note that $Z_1 = \beta_1(1 - \beta_1)\frac{\delta^2}{1-\delta}(e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u})$ comes with $Z_{1v} = (1 - \beta_1)\frac{\delta}{1-\delta}(e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u})$ 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.

²²In fact, the costs of increasing Z_1 are $s_1 \left(Z_1 + \sum_{i=1}^{\infty} \max\left\{0, \frac{Z_1}{\delta^2(1-\beta_1)} - \frac{\beta_1}{1-\delta}(e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u})(1 - \delta^i(1 - \beta_1))\right\} \right)$. 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.

²³(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 \mathbf{C}_1 , his expected as well as real utility is

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

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 \mathbf{C}_1 , but also goes for the virtual contract. Recall that under symmetric information, the net per-period profits the principal generates are $\pi^r - \bar{\pi} = (e^{FB}\theta - c(e^{FB}) - \bar{u} - \bar{\pi}) \left(1 + \beta(1 - \beta) \frac{\delta^2}{1 - \delta}\right)$. Hence, under symmetric information, the principal's total profits when dealing with agent 2 would be $\frac{(e^{FB}\theta - c(e^{FB}) - \bar{u} - \bar{\pi})}{1 - \delta} \left(1 + \beta_2(1 - \beta_2) \frac{\delta^2}{1 - \delta}\right)$. If letting agent 2 choose \mathbf{C}_1 , the principal's profits dealing with agent 2 are $(1 + \delta) (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u}) + (2 - \beta_1) \beta_1 \frac{\delta^2}{1 - \delta} (e^{FB}\theta - c(e^{FB}) - \bar{\pi} - \bar{u})$.

Therefore, the principal's costs when letting agent 2 choose \mathbf{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

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

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. \square