

THÈSE de DOCTORAT

de l'UNIVERSITÉ TOULOUSE CAPITOLE

Présentée et soutenue par

Monsieur Javier GONZALEZ MORIN

Le 15 juin 2026

Essais sur la théorie des incitations

École doctorale : **TSE Toulouse Sciences Economiques**

Spécialité : **Sciences Economiques - Toulouse**

Unité de recherche : **TSE-R - Toulouse School of Economics -
Recherche**

Thèse dirigée par Monsieur David MARTIMORT

Composition du jury

Rapporteur : M. William Bentley MACLEOD

Rapporteuse : Mme Malin Elisabeth ARVE

Examineur : M. Fabrice COLLARD

Directeur de thèse : M. David MARTIMORT

Président : M. Jean-Charles ROCHET



UNIVERSITÉ
TOULOUSE
CAPITOLE



Essays in the Theory of Incentives

Ph.D. Thesis

Javier González-Morín ¹

Toulouse School of Economics

April 2026

Jury Members

Jean-Charles Rochet, *Chair*

Malin Arve, *Referee*

Bentley MacLeod, *Referee*

Fabrice Collard, *Examiner*

David Martimort, *Advisor*

¹Address: 1, Esplanade de l'Université, 31080 Toulouse, France. E-mail: javier.gonzalez-morin@tse-fr.eu

ACKNOWLEDGMENTS

I would like to thank all those who helped and supported me throughout my PhD. The names mentioned below are only part of a much longer list, and I am grateful to many others who contributed, in different ways, to these years and to this thesis.

I would like to thank my advisor, David Martimort, for his guidance, support, and patience throughout my PhD. I am grateful for his advice and for the discussions that helped shape this thesis. I also want to thank Fabrice Collard, Bertrand Gobillard, Christian Hellwig, Johannes Hörner, Thomas Mariotti, Nour Meddahi, François Salanié, Anna Sanktjohanser, Sara Shahanaghi, Alex Smolin, Jean Tirole, Stéphane Villeneuve, Alex Winter, and Takuro Yamashita. A special thanks also to the professors in Salamanca who inspired and motivated me to pursue this goal, especially José Manuel Cascón, Miguel Ángel Malo, and Emma Moreno.

I am also grateful to my colleagues and fellow PhD students at TSE: Amirreza, Dario, Emil, Giovanni, Guillem, Hanlin, Jingyi, Juan, Kourosh, Liam, María, Pablo, Pascal, Pedram, Sergi, Tanja, Titouan, Valentina, Wen-Hao, and Xin. Sharing these years with them made the PhD a more stimulating and enjoyable experience.

Finally, I would like to thank my family and friends for their constant support and encouragement, especially Cristina. I am deeply grateful for everything they have done for me throughout these years.

SUMMARY

In this PhD thesis, I study dynamic incentives, private information, and commitment in three essays in economic theory. The common thread is that incentives, contracts, and policy promises depend not only on current outcomes, but also on how actions shape future states, future information, and future credibility. The three chapters analyze optimal contracts or policy promises when effort has intertemporal effects, when ex post signals affect screening, and when policymakers cannot fully commit.

In Chapter 1, *Exhaustion and Burnout: a Principal-Agent Approach*, I study a dynamic moral hazard problem in which effort increases current output but also raises the probability of future exhaustion. I model this as a continuous-time principal-agent problem with hidden action, limited liability, and a productivity state that switches between a productive regime and an exhausted one. The key insight is that exhaustion is harmful for future profits, yet informative about past effort, since harder work makes it more likely. The principal can therefore use protection during low-productivity spells as an incentive device. When exhaustion is more informative than success, the optimal contract relies more on payments during downturns; when success is more informative, it relies mainly on success-contingent bonuses. I also show that, when the productivity state is privately observed by the agent, the contract can still induce both effort and truthful revelation. The chapter thus shows that technologically harmful interruptions can nonetheless strengthen incentives.

In Chapter 2, *On the (Relative) Merits of Money Burning for Optimal Contracting*, coauthored with David Martimort, we study a screening problem in which an informed agent reports private information and the principal later receives an ex post private signal correlated with that information. A central contribution of the chapter is to incorporate money burning explicitly as an instrument of contractual design, alongside standard transfers. We characterize the set of incentive-feasible allocations and analyze when the first best can be implemented with transfers alone and when money burning becomes necessary. The analysis shows that the role of money burning depends on two key features of the informational environment: the timing at which the signal is learned and its quality. In particular, money burning is useful only when the signal is sufficiently informative, and its desirability differs sharply depending on whether the signal is observed before or after production. The chapter therefore shows that wasteful instruments may become part of the optimal contract because informational constraints make them valuable screening devices.

In Chapter 3, *Forward Guidance with Limited Commitment and Reputation*, coauthored with Emil Mortensen, we study monetary policy at the zero lower bound when the central bank cannot perfectly commit to future policy. In the standard New Keynesian model, forward guidance can have implausibly large effects on current output, giving rise to the forward guidance puzzle. We show that this puzzle is mitigated once private agents recognize that future policymakers may not wish to honor previously announced promises. We build a model with limited commitment and endogenous reputation, in which the central banker may be more or less forward-looking. In this environment, forward guidance remains useful, but its effects are bounded because longer promises are harder to sustain credibly. The chapter therefore provides a more disciplined theory of forward guidance by making credibility an endogenous equilibrium object.

RÉSUMÉ

Dans cette thèse de doctorat, j'étudie les incitations dynamiques, l'information privée et le commitment à travers trois essais de théorie économique. Le fil conducteur est que les contrats et les promesses de politique dépendent non seulement des résultats présents, mais aussi de la manière dont les actions influencent les états futurs, l'information future et la crédibilité. Les trois chapitres analysent la conception optimale des contrats ou des promesses de politique lorsque l'effort a des effets intertemporels, lorsque des signaux ex post interagissent avec la sélection adverse, et lorsque les autorités ne peuvent pas s'engager parfaitement.

Dans le chapitre 1, *Exhaustion and Burnout: a Principal-Agent Approach*, j'étudie un problème d'aléa moral dynamique dans lequel l'effort augmente la production courante mais accroît aussi la probabilité d'un épuisement futur. Je modélise cette situation comme un problème principal-agent en temps continu avec action cachée, responsabilité limitée, et une productivité qui alterne entre un régime productif et un régime d'épuisement. L'idée centrale est que l'épuisement nuit aux profits futurs tout en étant informatif sur l'effort passé. Le principal peut donc utiliser la protection pendant les périodes de faible productivité comme instrument incitatif. Lorsque l'épuisement est plus informatif que le succès, le contrat optimal repose davantage sur les paiements pendant les interruptions ; lorsque le succès est plus informatif, il privilégie les primes de performance. Je montre aussi que, lorsque l'état de productivité est privé à l'agent, le contrat peut encore induire l'effort et la révélation sincère de l'état.

Dans le chapitre 2, *On the (Relative) Merits of Money Burning for Optimal Contracting*, coécrit avec David Martimort, nous étudions un problème de sélection adverse dans lequel un agent informé déclare une information privée, puis le principal reçoit un signal privé ex post corrélé à cette information. Une contribution centrale du chapitre consiste à introduire explicitement le *money burning* comme instrument contractuel, en plus des transferts usuels. Nous caractérisons les allocations compatibles avec les contraintes d'incitation et analysons dans quels cas le premier best peut être mis en œuvre avec les seuls transferts, et dans quels cas le *money burning* devient nécessaire. Son rôle dépend de deux dimensions de l'environnement informationnel : le moment où le signal est observé et sa qualité informative. En particulier, le *money burning* n'est utile que lorsque le signal est suffisamment informatif, et son intérêt diffère selon que le signal est observé avant ou après la production.

Dans le chapitre 3, *Forward Guidance with Limited Commitment and Reputation*, coécrit avec Emil Mortensen, nous étudions la politique monétaire à la borne inférieure zéro lorsque la banque centrale ne peut pas s'engager parfaitement sur sa politique future. Dans le modèle keynésien standard, la *forward guidance* peut avoir des effets excessifs sur l'activité courante, ce qui donne lieu au *forward guidance puzzle*. Nous montrons que ce puzzle est atténué lorsque les agents privés comprennent que les décideurs futurs peuvent ne pas respecter les promesses passées. Nous développons un modèle avec commitment limité et réputation endogène. Dans cet environnement, la *forward guidance* reste utile, mais ses effets sont bornés, car les promesses longues sont plus difficiles à rendre crédibles.

Contents

Acknowledgments	1
Summary	2
Résumé	3
1 Exhaustion and Burnout: a Principal-Agent Approach	7
Abstract	7
1.1 Introduction	7
1.2 Two-Period Model	11
1.3 Fully-fledged Model	14
1.4 Optimal Contracting	19
1.5 Extensions and Robustness	31
1.6 Conclusion	36
Appendices	39
1.A Main Proofs	39
1.B First-best effort profile	56
2 On the (Relative) Merits of Money Burning for Optimal Contracting	57
Abstract	57
2.1 Introduction	57
2.2 The Model	61
2.3 Benchmarks	64
2.4 Incentive-Feasible Allocations	66
2.5 When the Signal Is Learned before Production	68

2.6	When the Signal Is Learned after Production	75
Appendices		83
2.A	Main Proofs	83
2.B	A Revelation Principle	92
2.C	Participation Constraints	94
3	Forward Guidance with Limited Commitment and Reputation	96
	Abstract	96
3.1	Introduction	96
3.2	Related Literature	98
3.3	A Model of Forward Guidance	100
3.4	Solution to the Forward Guidance Puzzle	106
3.5	The Effect of Forward Guidance	108
3.6	Relation to Discounted Euler Equation	116
3.7	Extension: Forward Guidance and Quantitative Easing	118
3.8	Conclusion	119
Appendices		123
3.A	Main Proofs	123
3.B	Interest Rate under Optimal Discretion	126
3.C	Equilibrium Description	127
3.D	Forward guidance and the variance of the demand shock	131

List of Figures

1.1	Productivity process	15
1.2	Value functions in the exhaustion-dominant case	22
1.3	Sample paths in the exhaustion dominant-case.	26
1.4	Value functions in the success-dominant case	27
1.5	Sample paths in the success-dominant case	30
1.6	Value functions with private state	33
1.7	Sample paths with private state	34
2.1	Transfers to implement first-best allocation	66
2.2	Money burning to implement first-best allocation	69
2.3	Output schemes in the regular case	72
2.4	Signal-dependent distortions and money burning thresholds.	74
2.5	Optimal contract of Proposition 2.6.	79
3.1	Mitigation of forward guidance effect for $\mu_0(\beta) = 0.6$	109
3.2	Reputation and forward guidance	111
3.3	Reputation convergence with $\mu_0(\beta) = 0.5$	112
3.4	A simulation of reputation and inflation with $\tau = 3$ and $\mu_0(\beta) = 0.5$	113
3.5	Forward guidance and the persistence of the natural rate of interest	114
3.6	Discounted Euler equation model	117
3.7	Forward guidance with government debt	119
3.8	Forward guidance and the standard deviation of the innovation to the demand shock	131
3.9	Nominal interest rates with medium realizations of the demand shock.	132

List of Tables

3.1	Parameter values	110
3.2	Welfare losses relative to ZLB with no FG	115
3.3	Recession at the ZLB with no forward guidance	116
3.4	Interest rates under discretionary policy with no forward guidance.	118

Chapter 1

Exhaustion and Burnout: a Principal-Agent Approach

Javier González-Morín

ABSTRACT

I study dynamic contracts where effort increases output but also raises the risk of exhaustion, leading to temporary downturns in productivity. Although exhaustion is a bad signal about future profits, it is a good signal of the agent’s effort. These downturns can thus be used as incentive devices: the principal commits to compensate the agent when productivity collapses, while successes play the conventional role of performance pay. Optimal contracts thus blend rewards with protection. My results show that negative intertemporal effects, although technologically harmful, can be leveraged to strengthen incentives in dynamic moral hazard settings. When downturns are private information of the agent, the contracts still provide payments during those periods to induce truthful revelation.

1.1. INTRODUCTION

Can peak performance be sustained indefinitely? Across many settings, periods of intensive workload are followed by costly slowdowns. When teams push, output rises in the moment but the risk of incidents increases—an injury at a worksite, a safety stand-down, a machine failure, or fatigue-driven errors that take time to unwind. Evidence from call centres, hospitals, and emergency services shows that heavier workloads and within-shift fatigue reduce performance, raise errors and delays, and can increase accident risk ([Collewet and Sauermann, 2017](#); [Kc and Terwiesch, 2009](#); [Bavafa and Jónasson, 2024](#); [Ibanez and Toffel, 2020](#)).

In practice, many workplaces build recovery into operations—micro-breaks on assembly lines, shift caps and handovers in hospitals, “time-out” protocols after safety events, and paid sick leave. Reviews and

meta-analyses show that brief resets, rest, and time off tend to improve safety and, in many settings, performance (Tucker, 2003; Albulescu et al., 2022; see also Vander Weerdt et al., 2023, on paid sick leave and firm outcomes). This finding motivates a clear economic question: How should long-term relationships be structured to induce strong performance and manage such interruptions?

In this paper I study dynamic contracts in environments where exhaustion—linking today’s effort to tomorrow’s risk of depletion—is a first-order feature of production. These downturns result in foregone output, yet they can be leveraged to sharpen incentives. The key idea is simple: the principal commits ex ante to protect the worker when production must pause, and that protection serves as a reward for past effort. Because working harder today both raises output and makes a future interruption more likely, promised support during interruptions becomes more valuable ex ante, strengthening the incentive to exert effort in the first place.

Specifically, I consider a continuous-time principal–agent relationship with hidden effort. The agent is protected by limited liability, and both players are risk-neutral. The agent’s productivity is modelled as a controlled Markov chain with two regimes—productive and exhausted—so that both exhaustion and recovery arrive with positive intensity. Effort raises the arrival rate of successes in the productive regime and, crucially, shifts the transition intensities between regimes by increasing the hazard of exhaustion and reducing the rate of recovery. The principal observes successes and the current productivity regime, but not effort. A contract specifies the compensation that is given to the agent conditional on the public history. The optimal contract is designed to maximize the principal’s profit.

My main results are as follows: (1) when the technology is *exhaustion-dominant*—extra effort raises the hazard of exhaustion more than the rate of successes—the optimal contract loads incentives on exhaustion events by committing to pay during low-productivity spells; (2) when the technology is *success-dominant*—effort boosts successes more than it increases exhaustion—the optimal contract loads incentives on successes, committing to large success-contingent bonuses sized to compensate both for effort and for the time spent in low productivity; and (3) when productivity is privately observed, the contract commits to pay on both successes and during downturns, using these two contingencies together to induce high effort and elicit the agent’s productivity state.

Effort has two observable consequences in this environment: it raises contemporaneous successes when production runs smoothly, and it also makes a pause more likely by pushing the system towards a low-productivity spell. Because both traces come from the same effort, a downturn is bad news about future output yet good evidence of past effort. The contract exploits this double signal in the spirit of the classic Informativeness Principle (Holmström, 1979): it places most of the incentive weight on the history that effort makes most revealing, while still using both sides to shape continuation promises. In this sense, the technology itself—how effort tilts success rates versus transition hazards—determines the contract’s shape: which event becomes decisive, where incentive weight is concentrated, and how histories govern continuation.

If an exhaustion event is more informative about effort than a success is—the exhaustion-dominant case—the relationship unfolds as follows. Initially, the contract runs a probationary phase of fixed length. If neither a success nor an exhaustion event arrives within this window, the principal fires the agent and pays no compensation. If a switch into low productivity occurs before probation ends, the

agent passes probation and the relationship moves to a stationary phase with no further termination risk. From then on, compensation is tied to the low-productivity spells. The principal pays a lump sum at each exhaustion event and a constant flow while the spell lasts. This schedule resembles accident insurance—an onset payment plus a per-period indemnity—but here it functions as an incentive device: it rewards past effort by putting transfers on the paths that higher effort makes more likely. During probation, successes are weaker evidence than a switch, so they are not paid directly. Instead, each success extends the probation window, raising continuation value, lowering the chance of termination, and increasing the likelihood of reaching the stationary phase.

If a success is more informative about effort than an exhaustion event is—the success-dominant case—the relationship also begins with a fixed-length probation. A first success ends probation and moves the relationship into a stationary phase with no further firing risk. From then on, pay is tied to successes through sizable success-contingent bonuses. If no success arrives before probation ends but the agent switches to the low-productivity state, the principal treats this as mild evidence of effort: probation is paused while productivity is low (with no firing risk) and, upon recovery, resumes with a granted extension. Because bonuses are not paid while productivity is low, they are chosen *ex ante* to offset expected time spent there. Their level therefore rises with both the probability of switching into low productivity and the expected share of time spent there, resembling a compensating premium.

While accidents and injuries are readily verifiable, depletion and burnout typically are not. To capture this, I consider the setting in which productivity is privately observed by the agent, so that the principal must elicit the state along the relationship. Because the state is reported by the agent, the contract should not reward nor punish the report itself—doing so would invite over- or under-reporting. The design keeps the worker through reported low-productivity spells and specifies transfers on two contingencies: publicly verifiable success events and time spent in reported low productivity. Success-contingent bonuses are set to satisfy the effort incentive constraint, since higher effort increases the arrival rate of successes and therefore the agent’s expected continuation value. Payments during low-productivity periods are set to satisfy the truth-telling constraints, making honest revelation weakly optimal at every date—there is no gain from concealing a downturn, and no profit from falsely declaring one. Calibrated jointly, these schedules implement high effort and truthful reporting while avoiding any direct reward or penalty tied to the act of reporting.

This paper contributes to the continuous-time dynamic moral-hazard literature using recursive promised-utility methods, with [Sannikov \(2008\)](#) as the canonical reference. The setup is a close variant of [Biais et al. \(2010\)](#)—risk-neutral agents under limited liability with Poisson timing—with successes as the arrival events instead of failures. My main addition is a dynamic technological externality of effort, i.e., exhaustion and recovery, so the productivity state is endogenous and effort-controlled. The work by [Bernard \(2024, 2025\)](#) is related methodologically, characterizing efficient payoff frontiers under Poisson information. However, while his framework studies strategic interactions between multiple players in repeated or stochastic games, my paper focuses on a single principal designing incentive-compatible contracts—that is, on the upper boundary of the implementable set rather than the equilibrium payoff set sustained by mutual incentives. Related is [Zhu \(2013\)](#), where shirking provides private utility and may be deliberately permitted. In my model, low effort is technological recovery rather than indirect pay,

and thus not a wage substitute. Because the relationship features two observables (performance events and regime transitions), my design results are related to the Informativeness Principle [Holmström \(1979\)](#): the more revealing observable bears most of the incentive load, while the other manages promises.

A related line of work studies the intertemporal technological links created by effort. [Fernandes and Phelan \(2000\)](#) develop recursive machinery for intertemporal links in repeated agency (although tractability for hidden-action moral hazard is limited); [Sannikov \(2014\)](#) analyses long-run incentives when effort has cumulative, positive effects; [DeMarzo and Sannikov \(2017\)](#) study effort-driven learning that separates beliefs; and [Martimort and Straub \(2016\)](#) generate dynamics via an irreversibility constraint that ties actions across periods. My contribution differs in that effort depletes—rather than builds—the state, and I study how an optimal contract should allocate incentives across two observables when effort both improves the current performance and increases the risk of tomorrow’s interruption.

Close in spirit are [Halac and Prat \(2016\)](#) and [Board and Meyer-ter-Vehn \(2022\)](#), which feature an endogenous state directly controlled by effort (attention, reputation). These papers study how effort builds the state and shapes information or reputation dynamics; in contrast, I instead study how effort erodes the state and derive the contract’s allocation of rewards across performance events versus depletion episodes.

A separate branch of the literature models fatigue and recovery as economic states. [Bechtold et al. \(1984\)](#), provide early work–rest scheduling in operations with productivity decay, showing how planned pauses can raise output. [Nocetti \(2008\)](#) models fatigue as a stock that depresses marginal productivity, yielding intertemporal substitution in effort and lower average output—evidence for the mechanism that hard work today can hinder performance later. [Dragone \(2009\)](#) studies a renewable effort stock in which rest rebuilds capacity, characterizing optimal effort–rest cycles. [Baucells and Zhao \(2019\)](#) axiomatize the accumulation and decay of fatigue and derive optimal pacing rules in single-agent settings, and [Sáez-Martí \(2025\)](#) embeds breaks and remote work in a principal–agent framework with schedule control. My contribution differs in that I place an effort-controlled depletion state inside an optimal contracting problem with hidden action and two observables: effort affects transition hazards, and the contract endogenously allocates bonuses vs. depletion-side payments to provide incentives—using “insurance” during exhaustion as an incentive instrument rather than only for risk sharing.

The results in this paper connects to firms’ use of policies such as paid sick leave when adopted on a voluntary basis. Indeed, before San Francisco’s mandate, 73% of firms already provided paid sick leave [Colla et al. \(2014\)](#), and nationally 79–81% of workers have employer-provided access even without a federal mandate ([U.S. BLS, 2024](#)). Conceptually, these practices fall under non-wage amenities—including paid leave/paid recovery—which are standard components of compensation ([Mas, 2025](#)) and account for a sizable share of total compensation ([Maideu-Morera, 2024](#)). This matters for measurement: the premium predicted by my model may appear in bonuses or in paid-recovery benefits, rather than in posted wage rates.

The result also relates to the literature on compensating differentials: success-contingent bonuses play the role of the premium for exposure to downtime and should compress when exposure falls. This aligns with [Lavetti \(2023\)](#), who reviews credible evidence that compensation co-moves with job risk,

and with [Guardado and Ziebarth \(2019\)](#), who document positive premia for exogenous workplace risk. Recent work debates the private provision of disability insurance ([Seibold et al., 2025](#); [Fischer et al., 2024](#); [Cabral and Cullen, 2019](#)). I contribute to this debate by providing a theoretical framework that predicts when employers embed disability-like protection in employment contracts and when they do not.

The remainder of the paper proceeds as follows. Section 2 presents a two-period benchmark to identify the key mechanisms. Section 3 lays out the continuous-time model. Section 4 characterizes optimal contracts in the two informativeness regimes (Propositions 2 and 3) and their dynamic evolution (Corollaries 1 and 2). Section 5 discusses some modelling assumptions and how modifying those assumptions would affect the optimal contract. Section 5 concludes. The Appendix contains proofs and omitted derivations.

1.2. TWO-PERIOD MODEL

I first present a simplified version of the model that allows me to highlight important features of extenuation in optimal contracts. High effort in the first period leads to a higher probability of depletion in the second period, when the agent is known to be unable to perform well. The principal anticipates these events and therefore combines bonuses (payments for performance) and insurance (payments when the agent is known to be depleted) in the optimal contract.

There are two players: a principal and an agent. The principal can hire the agent to run a project over two periods, indexed by $t \in \{1, 2\}$. The two players are risk neutral and discount payoffs with common factor $\delta = 1$. The agent is protected by limited liability.

If the agent is hired, he decides how much effort $a_t \in \{0, 1\}$ to exert in each period t . Effort is costly for the agent, and entails a disutility equal to ca_t with $c > 0$. In every period, the project can result in either a success or a failure. Each success has a value of $v > 0$ for the principal, and each failure has a value of 0. The probability of a success is affected by the effort exerted by the agent, and by a productivity state variable $\theta_t \in \{H, L\}$. The probability of a success in period t in state θ_t is

$$\mathbb{P}(\text{success}_t | a_t, \theta_t) = \begin{cases} \pi_0 & \text{if } \theta_t = L \\ \pi_0 + (\pi_1 - \pi_0)a_t & \text{if } \theta_t = H \end{cases}$$

with $0 < \pi_0 < \pi_1 < 1$. In state H the likelihood of a success is higher the more effort the agent exerts, while effort is useless in state L .¹ A success can still be realized in state L ; however the probability of such an event is indistinguishable from the probability of a success if a productive agent exerts no effort at all.

The state is commonly known to be $\theta_1 = H$ in period $t = 1$. Note that it is common knowledge that in period $t = 1$ both successes and failures realize on-path with positive probability regardless of the effort exerted by the agent. Effort exerted in period $t = 1$ triggers a stochastic transition of the productivity

¹This situation is clearly an extreme case, since one could consider that productivity in state L is reduced to a smaller, though still positive, value. The results remain unchanged as long as the productivity of effort in state L is small enough.

state variable θ_t . In particular:

$$\begin{aligned}\mathbb{P}(\theta_2 = L|a_1) &= \phi_0 + (\phi_1 - \phi_0)a_1, \\ \mathbb{P}(\theta_2 = H|a_1) &= 1 - \phi_0 - (\phi_1 - \phi_0)a_1.\end{aligned}$$

where $0 < \phi_0 < \phi_1 < 1$. The more effort that is exerted in period $t = 1$, the more likely the agent is to become unproductive in period $t = 2$, but neither of the two realisations of the state in period $t = 2$ perfectly reveals the effort in period $t = 1$. The productivity state is public in every period.

The principal offers a long-term contract that specifies transfers contingent on the public events. Effort is privately chosen by the agent; thus the principal can write a contract contingent on the successes and failures and the realisation of the productivity state in period $t = 2$. Thus, a contract is a tuple $(w_1, b_1, w_2^{FL}, w_2^{SL}, w_2^{SH}, b_2^{SH}, w_2^{FH}, b_2^{FH}) \in \mathbb{R}_+^8$, where the notation w_1 stands for the base payment that the principal commits to give to the agent in period $t = 1$ regardless of the success/failure realisation, and b_1 stands for the additional success-contingent bonus paid in period $t = 1$. The contract is history-dependent, such that $w_2^{j\theta}$ is the base payment given in period $t = 2$ if a success or failure, indexed by $j \in \{S, F\}$, was observed in period 1 and the productivity state $\theta \in \{H, L\}$ is realized in period 2, whereas $b_2^{j\theta}$ is the additional success-contingent bonus paid in period 2 if the combination $j\theta$ was observed. It is not necessary to specify bonuses in period 2 if the state is $\theta_2 = L$, since it is common knowledge that the agent deserves no credit for any realized success. In addition, as both players are risk-neutral and share a common discount factor, it is well known that there is no loss of generality when only terminal payments are considered, as any payment in period 1 can always be rescaled by a factor $\frac{1}{\delta}$ and entails no benefit or cost for either the principal or the agent. However, I will consider the optimal contract that front-loads transfers as much as possible.² Since the agent is protected by limited liability, all the payments must be nonnegative.

The timing of events is as follows:

Period 0

- 0.1 The principal proposes a contract $\mathcal{C}$.
- 0.2 The agent accepts/rejects $\mathcal{C}$. If the contract is rejected, both the principal and the agent obtain their outside option; otherwise:

Period 1

- 1.1 If $\mathcal{C}$ is accepted, then the agent chooses a_1 privately.
- 1.2 A success/failure $j \in \{F, S\}$ is realized; payment w_1 or $w_1 + b_1$ is provided.

Period 2

- 2.1 The state $\theta_2 \in \{L, H\}$ is publicly realized.
- 2.2 The agent chooses a_2 privately.

²Payments will be front-loaded if the discount factor of the agent, δ_A , is smaller than the discount factor of the principal, δ_P . This contract can be seen as the limit of the optimal contract as $\delta_A \uparrow \delta_P$.

2.3 A success/failure is realized; payment w_2^{ji} or $w_2^{ji} + b_2^{ji}$ is provided.

The contract can induce profiles of effort of the form $(a_1, a_2^{SH}, a_2^{FH}) \in \{0, 1\}^3$. Since the addition of fatigue effects is detrimental for the implementation of effort, I assume that the optimal contract absent fatigue effects, *i.e.*, when $\phi_1 = \phi_0 = 0$, induces maximal effort in all periods. This amounts to assume that:

$$\frac{\pi_1 - \pi_0}{c} v \geq 1 + \frac{\pi_0}{\pi_1 - \pi_0} \frac{1}{1 - \pi_1}$$

In particular, this assumption implies that inducing effort in one period only has a positive value even though the agent earns some limited liability rent. Since there is discounting and the agent is known to be in the productive state in period 1, the principal would prefer to induce effort in period 1 only rather than in period 2 only. The optimal contract thus trades-off how much effort to induce in period 2 versus the size of the limited liability rents that are left to the agent.

PROPOSITION 1.1. *Any optimal contract induces $a_1 = 1$ and features $w_2^{SL} > 0$. In addition, there exist cutoffs $0 < \underline{\phi} < \bar{\phi} < 1$ such that:*

1. *If $\phi_1 \leq \underline{\phi}$, the optimal contract induces $a_t = 1$ if and only if $\theta_t = H$; and the only positive transfers are as follows:*

$$w_2^{SL} = \frac{c}{\pi_1 - \pi_0} \frac{\pi_1 - \pi_0 [1 - (\phi_1 - \phi_0)]}{\pi_1 \phi_1 - \pi_0 \phi_0} > 0$$

$$b_2^{SH} = b_2^{FH} = \frac{c}{\pi_1 - \pi_0}$$

2. *If $\phi_1 \in [\underline{\phi}, \bar{\phi}]$, the optimal contract induces $(a_1, a_2^{SH}, a_2^{FH}) = (1, 1, 0)$; and the only positive transfers are as follows:*

$$w_2^{SL} = \frac{\pi_0 c}{\pi_1 - \pi_0} + \frac{(1 - \pi_0) c}{\pi_1 \phi_1 - \pi_0 \phi_0} > 0$$

$$b_2^{SH} = \frac{c}{\pi_1 - \pi_0}$$

3. *Assume $\phi_1 \geq \bar{\phi}$ the optimal contract induces $(a_1, a_2^{SH}, a_2^{FH}) = (1, 0, 0)$; the only positive transfer is as follows:*

$$w_2^{SL} = \frac{c}{\phi_1 \pi_1 - \phi_0 \pi_0} > 0$$

The optimal contract is essentially³ unique except for the cases $\phi = \underline{\phi}$ and $\phi = \bar{\phi}$.

Since period 2 is the last period, the optimal way of inducing high effort is to use the one-shot optimal contract, which involves a zero base payment $w_2^{jH} = 0$ and positive bonus $b_2^{jH} = \frac{c}{\pi_1 - \pi_0}$. Since successes realize with positive probability of at least π_0 when the agent shirks, the bonus exceeds the actual flow cost of effort, thus leaving the agent some limited liability rent $\frac{\pi_0 c}{\pi_1 - \pi_0} > 0$. The incentives to exert effort

³Had I allowed bonuses in period $t = 2$ when $\theta_2 = L$ one could set $b_2^{SL} = \frac{w_2^{SL}}{\pi_0}$ and use any convex combination of the two. All those are different forms of insurance, since the agent is known to be unproductive and exerts no effort. The contract is therefore unique up to equivalent forms of insurance.

in period 1 are affected by various channels. First, the agent faces the static component of the trade-off, which accounts for the flow cost of effort and the bonus obtained in period 1 after a success is realized, b_1 . Second, the agent anticipates that, if he is productive again in period 2, he will obtain a limited liability rent when he is asked to work in that period. However, if he works in period 1, the state will be more likely to be $\theta_2 = L$, and part of that rent will be lost in expectation. Reducing that rent can make it easier for the principal to induce high effort in period 1. Third, the principal can commit to reward the agent when the state is $\theta_2 = L$ to compensate for the lost rent. Proposition 1 states that the principal should always reward the agent when the state $\theta_2 = L$ is realized after a success in period 1, to combine both signals about effort. The static reward b_1 should be zero since it provides low-powered incentives relative to w_2^{SL} . When exhaustion effects are low, $\phi_1 < \underline{\phi}$, the principal finds it optimal to induce maximal effort in period 2 when $\theta_2 = H$, since the expected rent that the agent misses when $\theta_2 = L$ is not significant enough. As exhaustion effects become more severe, $\phi_1 \in (\underline{\phi}, \bar{\phi})$, the optimal contract dictates that effort in period 2 should be induced only after a success in period 1 is realized and $\theta_2 = H$. The effect of this downsizing is twofold: the agent obtains a rent in period 2 only if he had a success in period 1; and the rent that the agent loses when the state is $\theta_2 = L$ is reduced. Both cases induce higher effort in period 1. Finally, when exhaustion effects are more significant, $\phi_1 > \bar{\phi}$, the principal decides to not induce effort in period 2 at all. In expectation, the agent becomes too worried about the rent he will lose if he becomes exhausted that it would become too costly to induce effort in period 1.

Although extenuation is technologically harmful, the principal leverages the low-productivity event to strengthen incentives. The optimal contract therefore uses *insurance as an incentive device*: it rewards past effort precisely when the agent is known to be depleted, even though both parties are risk-neutral. This situation contrasts with standard hidden action models in which insurance is motivated by risk sharing with a risk-averse agent: under first-best case one would fully insure, whereas under the moral hazard case only partial insurance is provided to preserve incentives by making pay depend on outcomes. Here, by contrast, transfers in the low state are not about smoothing risk; they are the very instrument that implements effort, because paying when $\theta = L$ credibly rewards past performance without distorting current effort (effort is unproductive in L).

In the remainder of the paper I adopt a continuous-time setting to capture two economically salient features missing from two periods: *recovery*—after depletion the match can return to high productivity; and *recurrence*—depletion and recovery can repeat, making insurance a reusable incentive and turning continuation versus termination into a genuine economic margin.

1.3. FULLY-FLEDGED MODEL

There are two players: a principal and an agent. The principal can hire the agent to run a project over time. Time is continuous and is indexed by $t \in [0, +\infty)$. The agent must choose a process $a = (a_t)_{t \geq 0}$, where $a_t \in [0, 1]$ is the effort exerted at each instant t . The cost of exerting effort is captured by a flow $ca_t dt$. Effort generates (stochastically) some profits over time. In particular, a success of value $v > 0$ for the principal is generated at time t with intensity $p_t(a_t, \theta_t) dt$. The state $\theta_t \in \{L, H\}$ determines

the productivity of effort at time t . In particular:

$$p_t(a_t, \theta_t) = \begin{cases} \pi_0 & \text{if } \theta_t = L, \\ \pi_0 + (\pi_1 - \pi_0)a_t & \text{if } \theta_t = H, \end{cases}$$

where $0 < \pi_0 < \pi_1$. That is, effort increases the likelihood of a success if the state is H , but is completely useless if the state is L . The state θ_t evolves over time as a controlled Markov chain: if the state is H at time t , the state jumps to L with instantaneous probability $[\phi_0 + (\phi_1 - \phi_0)a_t]dt$, with $\phi_1 > \phi_0 > 0$. If the state is L at time t , the state jumps to H with instantaneous probability $[\gamma_0 + (\gamma_1 - \gamma_0)a_t]dt$, with $\gamma_0 > \gamma_1 > 0$. That is, the transitions from H to L are more likely to occur with higher levels of effort, while the transitions from L to H are more likely to occur with lower levels of effort.




Figure 1.1: Productivity process as a Markov chain. Extenuation increases with effort ($\phi_1 > \phi_0$), and recovery decreases with effort ($\gamma_1 < \gamma_0$).

Let $(X_t)_{t \geq 0}$ denote the process of accumulated successes. This process is nonnegative and nondecreasing, with stochastic upwards jumps of size $dX_t = 1$. This process takes an initial value $X_{t-} = 0$. Let τ_H be the stopping time at which the state jumps from H to L , and let τ_L be the stopping time at which the system jumps from L to H . Let the associated jump processes be denoted as $(J_t^\theta)_t$. The state $\theta \in \{L, H\}$, jointly with an effort profile a , defines a probability measure $\mathbb{P}^{a, \theta}$ over the random variables X_t , τ_L , and τ_H . The associated expectation operator is denoted by $\mathbb{E}$.

I denote by $(\mathcal{F}_t)_t$ the filtration generated by the processes $(X_t, \theta_t)_t$, whose paths are publicly observed. This filtration satisfies the usual assumptions.⁴ A contract is an $\mathcal{F}_t$ -adapted process $(Y_t)_t$ that represents the cumulative transfers paid to the agent up to time t . Because of limited liability, this process is nonnegative and nondecreasing: $Y_{t-} = 0$, $Y_t \geq 0$ and $dY_t \geq 0$. The effort process $a = (a_t)_t$ is $\mathcal{F}_t$ -predictable; that is, it is a function of all the information up to t^- .

Both principal and agent are risk-neutral and discount the future at the common rate $r > 0$.⁵ Fixing a

⁴That is, the filtration is complete (all $\mathbb{P}^{a, \theta}$ -null sets are in $\mathcal{F}_0$) and right-continuous, i.e., $\mathcal{F}_t = \bigcap_{s > t} \mathcal{F}_s$ for all $t \geq 0$.

⁵In continuous time environments with risk-averse players it is quite common to assume different discount rates. Letting the agent be weakly more patient than the principal usually ends up with the nonexistence of an optimal contract, since the principal would prefer to delay transfers infinitely and compensate the agent with increases in the continuation payoff only. However, in this paper I find optimal contracts where, at sufficiently high continuation values, increasing the promise is as costly as an immediate transfer. Thus, there is no strict incentive to delay transfers thereafter.

pair (Y, a) , the expected payoffs for the principal and the agent are respectively⁶

$$\begin{aligned} V(Y, a, \theta_0) &= \mathbb{E} \left[\int_0^\infty e^{-rt} (v dX_t - dY_t) \right], \\ U(Y, a, \theta_0) &= \mathbb{E} \left[\int_0^\infty e^{-rt} (dY_t - ca_t dt) \right]. \end{aligned}$$

1.3.1. First-best benchmark

I start by characterizing the benchmark allocation in which effort is directly chosen and paid for by the principal. This case corresponds to the *first-best* benchmark, without incentive or liability constraints. It provides an upper bound on attainable welfare and a useful reference point for the constrained problems that will be studied later.

The first-best effort profile $a^* = (a_t^*)_t$ solves

$$W(\theta_0) = \sup_a \mathbb{E} \left[\int_0^\infty e^{-rt} (dX_t - ca_t dt) \right],$$

Because the problem is Markovian in the state θ_t , the value functions admit the following Hamilton–Jacobi–Bellman (HJB) characterization:

$$\begin{aligned} rW(H) &= \sup_{a_H \in [0,1]} \left\{ v(\pi_0 + (\pi_1 - \pi_0)a_H) - ca_H + [\phi_0 + (\phi_1 - \phi_0)a_H] [W(L) - W(H)] \right\}, \\ rW(L) &= \sup_{a_L \in [0,1]} \left\{ v\pi_0 - ca_L + [\gamma_0 + (\gamma_1 - \gamma_0)a_L] [W(H) - W(L)] \right\}. \end{aligned} \tag{P^{fb}}$$

Each objective combines three components: (i) the expected instantaneous output, $\mathbb{E}[v dX_t | a_t, H] = v[\pi_0 + (\pi_1 - \pi_0)a_t]dt$ in state H and $\mathbb{E}[v dX_t | a_t, L] = v\pi_0 dt$ in state L ; (ii) the contemporaneous flow cost $ca_t dt$; and (iii) the expected continuation value from a state switch, weighted by its intensity $\phi_0 + (\phi_1 - \phi_0)a_H$ when moving to L and $\gamma_0 + (\gamma_1 - \gamma_0)a_L$ when moving to H .

Under the minimal assumption $(\pi_1 - \pi_0)v > c$, to ensure that effort in the productive state is profitable from a static point of view, the efficient effort profile is bang-bang:

$$a_t^* = \begin{cases} 1 & \text{if } \theta_t = H, \\ 0 & \text{if } \theta_t = L. \end{cases}$$

Details can be found in the Appendix. When productivity is low ($\theta_t = L$), effort has no effect on the success rate: it only increases costs and slows recovery (reducing the transition rate from γ_0 to γ_1). Hence the optimal choice is no effort. When productivity is high ($\theta_t = H$), the marginal return to effort is strictly positive regardless of any dynamic concerns. Hence maximal effort is always optimal, allowing the decision maker to exploit the productive state fully while it lasts, before the process eventually

⁶Integrals are of Lebesgue-Stieltjes type, where $\int_s^t$ denotes $\int_{[s,t]}$ and $\int_s^{t-}$ denotes $\int_{[s,t)}$.

transitions to the low state L .

The values derived from being in states H and L are respectively as follows:

$$W(H) = \frac{r + \gamma_0}{r + \gamma_0 + \phi_1} \frac{\pi_1 v - c}{r} + \frac{\phi_1}{r + \gamma_0 + \phi_1} \frac{\pi_0 v}{r} > 0,$$

$$W(L) = \frac{\gamma_0}{r + \gamma_0 + \phi_1} \frac{\pi_1 v - c}{r} + \frac{r + \phi_1}{r + \gamma_0 + \phi_1} \frac{\pi_0 v}{r} \in (0, W(H)).$$

These expressions can be read as convex combinations of two benchmark payoffs: (i) the expected discounted net flow under high productivity with maximal effort, $(\pi_1 v - c)/r$; and (ii) the expected discounted flow under low productivity, $\pi_0 v/r$. The weights on these benchmarks reflect both transition intensities and discounting. Higher recovery (γ_0) and lower extenuation (ϕ_1) shift the weight towards a high payoff. As the discount rate $r \rightarrow 0$, the weights converge to $(\frac{\gamma_0}{\gamma_0 + \phi_1}, \frac{\phi_1}{\gamma_0 + \phi_1})$, which is the stationary distribution of the Markov chain θ_t when controlled by the effort process a^* .

The first-best solution derived above serves as a reference point for the constrained problems studied next. In particular, it will be useful to work with the principal's net payoff once she commits to deliver a continuation utility U to the agent, while otherwise retaining the first-best allocation. I therefore define

$$V_H^{fb}(U) = W(H) - U,$$

$$V_L^{fb}(U) = W(L) - U,$$

which represent the principal's value in states H and L , respectively, when incentive constraints are relaxed but the agent is guaranteed U .

1.3.2. Incentive Compatibility

Having characterized the first-best allocation, I now turn to the central constraint of the problem: how to design transfers so that the agent voluntarily implements the efficient effort profile a^* derived above. Since effort is privately chosen and costly, the contract must align the agent's incentives with the principal's objectives in every productivity state. My focus is therefore not on arbitrary effort processes, but specifically on making a^* incentive compatible until the eventual event of termination. In Section 1.4.3, I show that this effort process is indeed the optimal one to induce when the principal seeks to maximize her profits.

To formalize this approach, fix a contract Y and an effort profile a , which is a best reply to contract Y . Let $(U_t^a)_t$ denote the agent's continuation utility at time t . This process is defined as follows:

$$U_t^a = \mathbb{E} \left[\int_t^\infty e^{-r(s-t)} (dY_s - ca_s ds) \mid \mathcal{F}_t \right]$$

$$= \sup_{\hat{a}} \mathbb{E} \left[\int_t^\infty e^{-r(s-t)} (dY_s - c\hat{a}_s ds) \mid \mathcal{F}_t \right],$$

where the state θ_t is part of the history up to time t through $\mathcal{F}_t$. To capture how contracts provide incentives, it is convenient to introduce the so-called *sensitivity processes*. Intuitively, these processes

measure how strongly the agent's continuation utility reacts to observable shocks. Since the payoff-relevant shocks differ across states, I allow the sensitivities to depend on the current productivity state.⁷ That is, I denote by β_t^θ the sensitivity of the agent's continuation value to successes in state θ at time t , and by χ_t^θ the sensitivity to transitions out of state θ at time t . These sensitivities summarize how the agent's future payoff changes after each type of event, and thus provide a compact way to express the marginal incentives for effort. When the state is $\theta_t = H$, reducing effort from $a_t^* = 1$ to $\hat{a}_t = 1 - \varepsilon$ reduces (i) the intensity of a success by $(\pi_1 - \pi_0)\varepsilon$; (ii) the intensity of a jump to state L by $(\phi_1 - \phi_0)\varepsilon$; and (iii) the cost of effort borne by the agent by $c\varepsilon$. Shirking is profitable if the reduction in effort cost overcomes the loss of continuation value from the reduced intensities of the shocks: $\beta_t^H(\pi_1 - \pi_0)\varepsilon + \chi_t^H(\phi_1 - \phi_0)\varepsilon < c\varepsilon$. When the state is $\theta_t = L$, increasing effort by ε does not affect the intensity of successes, but reduces the intensity of jumps to the H state by $(\gamma_1 - \gamma_0)\varepsilon < 0$. Increasing effort is profitable if the expected change in continuation value overcomes the added effort cost: $\chi_t^L(\gamma_1 - \gamma_0)\varepsilon > c\varepsilon$. The next lemma formalizes this intuition and characterizes the evolution of the continuation value $U_t \equiv U_t^{a^*}$ on the equilibrium path:

LEMMA 1.1. *The efficient effort process a^* is a best reply to contract Y if and only if the sensitivity processes satisfy at all t :*

$$\beta_t^H(\pi_1 - \pi_0) + \chi_t^H(\phi_1 - \phi_0) \geq c, \quad \chi_t^L(\gamma_1 - \gamma_0) \leq c. \quad (1.3.1)$$

The agent's continuation utility then evolves according to

$$dU_t = \begin{cases} rU_t dt + cdt - dY_t + \beta_t^H dM_t + \chi_t^H dN_t^H & \text{if } \theta_t = H, \\ rU_t dt - dY_t + \beta_t^L dM_t + \chi_t^L dN_t^L & \text{if } \theta_t = L, \end{cases} \quad (1.3.2)$$

where $dM_t = dX_t - \pi_1 \mathbb{1}\{\theta_t = H\} dt - \pi_0 \mathbb{1}\{\theta_t = L\} dt$, $dN_t^H = dJ_t^H - \phi_1 dt$, and $dN_t^L = dJ_t^L - \gamma_0 dt$ are $\mathbb{P}^{a^*, \theta}$ -martingales.

Relative to a benchmark without exhaustion, incentives can now be provided through two distinct channels: successes and state transitions. In the productive state H , the incentive compatibility condition in equation (1.3.1) requires that the combined sensitivity to both types of shocks is large enough to cover the flow cost of effort. The principal may therefore rely on either success-contingent rewards (β_t^H) exhaustion-contingent rewards (χ_t^H), or a mixture of both to implement effort.

In the unproductive state L , successes are uninformative about effort, and the only relevant sensitivity is that to transitions, χ_t^L . Since $\gamma_1 < \gamma_0$, the inequality in equation (1.3.1) becomes a negative lower bound for χ_t^L . Any nonnegative χ_t^L thus satisfies it, so incentive compatibility in L is easy to satisfy in principle. This observation reflects the fact that the principal and agent are more naturally aligned in the L state: both prefer the agent to rest. However, the contract must still be designed carefully. If transfers during L are too attractive, the agent could have an incentive to exert positive effort precisely to *delay recovery* and extend the period of receiving those benefits.

⁷This notation is redundant, since the sensitivities are also $\mathcal{F}_t$ -adapted, so they should include all the information up to time t , including the state θ_t . However, since the sensitivities at time t depend explicitly on the state θ_t , I find it useful for the sake of exposition to emphasize the state in the notation.

Equation (1.3.2) describes how the agent's continuation utility evolves on the equilibrium path. The drift terms reflect discounting and effort costs: in state H , the continuation value must increase at rate c to compensate for the cost of exerting effort, whereas in state L , no such term appears. The martingale terms, weighted by the sensitivities β_t^θ and χ_t^θ , capture the random adjustments in continuation value following realized successes and transitions.

Finally, the $-dY_t$ term reflects the fact that transfers are part of the continuation promise itself. Delivering dY_t to the agent at time t amounts to paying out a portion of the promised utility U_t . For the same continuation value U_t , a higher transfer today necessarily reduces the amount of utility that remains to be delivered in the future. For this reason, current transfers appear as a negative drift in the law of motion for U_t .

1.4. OPTIMAL CONTRACTING

This section characterizes how the principal exploits the available shocks to design optimal contracts that maximize her profits. My approach is to first derive the HJB equations heuristically, assuming differentiability of the value functions so that they can be manipulated in the usual way. I then verify that the candidate analytical solutions indeed satisfy the HJB equations and are differentiable, which justifies the heuristic derivation. I separately study the two possible regimes, depending on whether transitions or successes are the more informative signals of effort.

The principal's problem is to design a schedule of transfers—and thereby the associated continuation sensitivities—that implements the efficient effort profile a^* until termination while maximising her expected payoff. Transfers can depend on the current productivity state $\theta_t \in \{H, L\}$ and on observed shocks. Formally, I can represent a contract by the transfer rule

$$dY_t^\theta = w_t^\theta dt + b_t^\theta dX_t + p_t^\theta dJ_t^\theta,$$

where w_t^θ is a flow transfer, b_t^θ is a success-contingent bonus, and p_t^θ is a transition-contingent lump-sum transfer.⁸ Given this transfer scheme, the agent's continuation utility evolves according to

$$dU_t = \begin{cases} (rU_t + c - w_t^H - \pi_1\beta_t^H - \phi_1\chi_t^H) dt + (\beta_t^H - b_t^H) dX_t + (\chi_t^H - p_t^H) dJ_t^H & \text{if } \theta_t = H, \\ (rU_t - w_t^L - \pi_0\beta_t^L - \gamma_0\chi_t^L) dt + (\beta_t^L - b_t^L) dX_t + (\chi_t^L - p_t^L) dJ_t^L & \text{if } \theta_t = L. \end{cases} \quad (1.4.1)$$

The principal chooses transfers $(w_t^\theta, b_t^\theta, p_t^\theta)$ and sensitivities $(\beta_t^\theta, \chi_t^\theta)$ to maximize her continuation value, subject to incentive compatibility (1.3.1) and limited liability. To see how the dynamic programming equations are derived, consider first the state to be H . At the discount rate r , the expected flow of the principal's value when arriving at time t equals

$$r V_H(U_{t-}). \quad (1.4.2)$$

⁸Again, the state notation is redundant since Y is $\mathcal{F}$ -adapted, but useful since the transfers will explicitly depend on the current state θ_t .

This flow can be decomposed into two parts. First, the principal obtains the net benefits from current events. A success arrives at rate π_1 , yielding v but requiring that the bonus b_t^H be paid; a transition $H \rightarrow L$ arrives at rate ϕ_1 , requiring that the lump-sum p_t^H be paid; and the flow w_t^H must be paid continuously. The expected current payoff is therefore

$$\pi_1(v - b_t^H) - \phi_1 p_t^H - w_t^H. \quad (1.4.3)$$

Second, the principal must account for how her continuation value evolves. Using (1.4.1), three forces matter: the deterministic drift of U_t , the jumps after successes, and the jumps after transitions. This account gives

$$\begin{aligned} & \underbrace{V'_H(U_{t-}) [rU_{t-} + c - \pi_1\beta_t^H - \phi_1\chi_t^H - w_t^H]}_{\text{deterministic drift}} \\ & + \underbrace{\pi_1 \left(V_H(U_{t-} + \beta_t^H - b_t^H) - V_H(U_{t-}) \right)}_{\text{jump after success}} \\ & + \underbrace{\phi_1 \left(V_L(U_{t-} + \chi_t^H - p_t^H) - V_H(U_{t-}) \right)}_{\text{jump after transition}}. \end{aligned} \quad (1.4.4)$$

Adding (1.4.3) and (1.4.4) and identifying the result with (1.4.2) yields the HJB equation in state H :

$$\begin{aligned} r V_H(U) = & \sup_{\beta^H, \chi^H, b^H, p^H, w^H} \left\{ \pi_1(v - b^H) - \phi_1 p^H - w^H \right. \\ & + V'_H(U) [rU + c - \pi_1\beta^H - \phi_1\chi^H - w^H] \\ & + \pi_1 [V_H(U + \beta^H - b^H) - V_H(U)] \\ & \left. + \phi_1 [V_L(U + \chi^H - p^H) - V_H(U)] \right\}. \end{aligned} \quad (1.4.5)$$

An analogous argument for $\theta_t = L$ gives the following:

$$\begin{aligned} r V_L(U) = & \sup_{\beta^L, \chi^L, b^L, p^L, w^L} \left\{ \pi_0(v - b^L) - \gamma_0 p^L - w^L \right. \\ & + V'_L(U) [rU - \pi_0\beta^L - \gamma_0\chi^L - w^L] \\ & + \pi_0 [V_L(U + \beta^L - b^L) - V_L(U)] \\ & \left. + \gamma_0 [V_H(U + \chi^L - p^L) - V_L(U)] \right\}. \end{aligned} \quad (1.4.6)$$

These HJB equations hold subject to the limited liability requirement that transfers are nonnegative at each instant, $dY_t^\theta \geq 0$. This requirement does not strictly imply that each component $(w^\theta, b^\theta, p^\theta)$ must be nonnegative, since only the total transfer matters. For example, one could in principle allow a negative bonus $b^\theta < 0$ as long as the continuous payment w^θ is large enough that $w^\theta + b^\theta \geq 0$ whenever a success occurs; similarly for p^θ in the case of transitions. However, in the absence of shocks, it must be that $dY_t^\theta = w^\theta dt \geq 0$, which implies $w^\theta \geq 0$. Moreover, since successes and transitions are

independent Poisson shocks (conditional on the exerted effort level a_t), the probability that they occur simultaneously is of order $o(dt^2)$, so contracts that rely on offsetting a negative b^θ with a positive p^θ (or *viceversa*) are irrelevant.

The boundary conditions are given by $V_H(0) = V_L(0) = \frac{\pi_0 v}{r}$. Intuitively, when the promised continuation value to the agent reaches zero, the principal cannot incentivize any positive effort from that point on. The agent is then left with her outside option, which corresponds to exerting no effort forever. From that point onwards, successes occur only at the baseline rate π_0 , yielding a perpetual flow of expected value $\pi_0 v$ to the principal, discounted at rate r . This outcome explains why both value functions collapse to $\frac{\pi_0 v}{r}$ at $U = 0$.

1.4.1. Exhaustion-Dominant Technology

I first consider the case in which state transitions provide more informative signals about effort than successes, i.e.,

$$\frac{\phi_1}{\phi_0} > \frac{\pi_1}{\pi_0}.$$

In this regime, the principal relies mostly on transitions to provide incentives. The resulting contract is simple: in the high state (H), effort is induced through exhaustion-contingent payments; in the low state (L), both parties prefer to rest, so the principal simply keeps the agent's continuation value stable until the state returns to H . Termination can occur only in the high state, when the continuation value drifts downwards in the absence of good news. The next proposition characterizes the optimal contract in this regime.

PROPOSITION 1.2. Assume that $\frac{\phi_1}{\phi_0} > \frac{\pi_1}{\pi_0}$, $r \geq \phi_0 \left(1 + \frac{\pi_1 - \pi_0}{\phi_1 - \phi_0}\right)$ and

$$\frac{\pi_1 - \pi_0}{c} v > 1 + \frac{\phi_0}{\phi_1 - \phi_0} \frac{r + \gamma_0 + \phi_1}{r + \gamma_0},$$

and let $U^* \equiv \frac{c}{r} \frac{\phi_0}{\phi_1 - \phi_0}$. There exists an optimal contract given by:

$$\begin{aligned} \beta^H(U_t) &= C_1 \cdot (U^* - U_t) & b^H(U_t) &= 0, \\ \chi^H(U_t) &= \frac{c}{\phi_1 - \phi_0} - \beta^H(U_t) \frac{\pi_1 - \pi_0}{\phi_1 - \phi_0} & p^H(U_t) &= \chi^H(U_t) + U_t - U^*, \\ \chi^L(U_t) &= \begin{cases} \frac{r}{\gamma_0} U_t, & U_t \leq \frac{\gamma_0}{r + \gamma_0} U^*, \\ U^* - U_t, & \frac{\gamma_0}{r + \gamma_0} U^* < U_t \leq U^*, \end{cases} & p^L(U_t) &= 0, \\ w^L(U_t) &= r U^* \cdot \mathbb{1}\{U_t = U^*\}, & \beta^L(U_t) &= b^L(U_t) = w^H(U_t) = 0, \end{aligned}$$

where $C_1 = \left(1 - \left(\frac{\phi_1 \pi_0 - \phi_0 \pi_1}{\pi_1 (\phi_1 - \phi_0)}\right)^{\frac{1}{\alpha - 1}}\right) \in (0, 1)$ and $\alpha > 1$ is an implicit function of parameters $(\pi_0, \pi_1, \phi_0, \phi_1, r)$.

The associated value functions are shown in Figure 1.2. Under the regularity assumptions required for




Figure 1.2: Value functions in the optimal contract under the exhaustion-dominant case: $\frac{\phi_1}{\phi_0} > \frac{\pi_1}{\pi_0}$.

Proposition 1.2, these value functions can be solved in closed form:

$$V_H(U) = V_H^{fb}(U) - \frac{[(\pi_1 - \pi_0)v - c](r + \gamma_0)}{(r + \gamma_0 + \phi_1)r} \left(1 - \frac{U}{U^*}\right)^\alpha,$$

$$V_L(U) = V_L^{fb}(U) - \frac{[(\pi_1 - \pi_0)v - c]\gamma_0}{(r + \gamma_0 + \phi_1)r} \left(1 - \frac{U}{U^*} \frac{r + \gamma_0}{\gamma_0}\right)^\alpha \mathbb{1}\left\{0 \leq U < \frac{\gamma_0}{r + \gamma_0} U^*\right\}.$$

These value functions are concave and single peaked at interiors values of the agent's continuation payoff U , ensuring that the principal indeed prefers to initiate the contract. The optimal contract then starts at one of the two peaks, depending on the initial state θ_0 . It then moves along the curve while the same state remains, and jumps between curves when transitions occur.

When the productivity state is L , the objectives of the principal and the agent are naturally aligned, since both prefer the agent to rest. Incentive compatibility constraints are therefore slack. The principal can simply perfectly align her marginal value pre- and post transition from L to H , leading to $V_H'(U + \chi^L(U)) = V_L'(U)$. In addition, when this functional equation in U is solved, one obtains that the sensitivity $\chi^L(U)$ is also chosen to stabilize the continuation value: for low U_t , $\chi^L(U)$ is chosen so that the deterministic drift is exactly zero ($rU_t - \gamma_0\chi^L(U) = 0$), keeping U_t constant in the absence of shocks. For larger U_t , $\chi^L(U)$ is capped so that the continuation value never exceeds U^* , although this value will never be part of the equilibrium path under this contract. As a result, the continuation value never drifts to zero and the contract is never terminated in L . Transfers are deferred until U_t reaches U^* , at which point the principal delivers it smoothly with a flow $w^L(U^*) = rU^*$. The precise form of this delivery is not unique: once the continuation value is fixed at U^* , the principal can equivalently use flow transfers, bonuses, or transition-contingent payments to deliver the same utility, as long as they sum to rU^* .

When the productivity state is H , the incentive constraint binds. The contract relies primarily on

transitions as the informative signal that ultimately triggers payments, while successes are used only to extend the probationary phase. Accordingly, the principal disregards successes as a payment instrument ($b^H(U) = 0$) but maintains a positive sensitivity $\beta^H(U) = C_1(U^* - U)$ so that a success increases the promise without any transfer. Incentives are provided through exhaustion-contingent adjustments: the principal sets $\chi^H(U)$ to satisfy incentive compatibility, and commits to lump-sum transfers $p^H(U)$ to ensure continuation values never exceed U^* . These transfers are increasing in U_t , and once $U_t = U^*$, the principal maintains it with constant payments $p^H(U^*) > 0$. Additionally, once the continuation value is U^* , the marginal benefit of increasing the continuation value and the marginal benefit of increasing the lump-sum transfer $p^H(U)$ are equal: $-\phi_1 V'_L(U^*) = -\phi_1$. Thus, there is no incentive for the principal to delay transfers once the contract enters the stationary phase. Termination may occur in H during the probationary phase: absent transitions or successes, the continuation value drifts downwards deterministically,

$$\begin{aligned} dU_t &= (rU_t + c - \phi_1 \chi^H(U_t) - \pi_1 \beta^H(U_t)) dt \\ &= \underbrace{\left(r + C_1 \frac{\phi_1 \pi_0 - \phi_0 \pi_1}{\phi_1 - \phi_0} \right)}_{\equiv \kappa > 0} (U_t - U^*) dt < 0, \end{aligned}$$

and eventually reaches zero. A success during probation extends the relationship but never pays on impact: the jump $U_t \mapsto C_1 U^* + (1 - C_1) U_t$ postpones the termination time by a constant $\Delta T > 0$ that does not depend on U_t . In contrast, only a transition $H \rightarrow L$ moves the promise all the way to U^* and issues an immediate payment $p^H(U) > 0$, after which the relationship remains in the mature stationary phase.

This dual use of signals can be understood from the concavity of the principal's value in the promise U . Concavity means that the marginal cost of raising U increases with the level of U . Relying on a single instrument would force large jumps in continuation values whenever an informative event occurs, making the path more volatile and pushing U too close to zero, risking premature termination. By combining both signals, the principal smooths these adjustments: transitions provide the main incentive device and eventually trigger payments, while successes, although less informative, extend the probationary phase and stabilize the path of U at lower cost. More specifically, consider the optimality conditions associated with the HJB system (1.4.5)-(1.4.6). Since there is a shadow cost associated with incentive compatibility constraints, the optimal contract must include a wedge between the principal's marginal payoff pre- and post-success: $V'_H(U^*) < V'_H(U + \beta^H(U)) < V'_H(U)$. The concavity of V_H in U thus delivers $\beta^H(U) \in (0, U^* - U)$.

The assumption $r \geq \phi_0 \left(1 + \frac{\pi_1 - \pi_0}{\phi_1 - \phi_0} C_1 \right)$ guarantees that the transfer p^H needed to cap the continuation value at U^* is nonnegative. This condition means that discounting is strong enough that the principal prefers to pay the agent immediately after an (informative) transition rather than postponing all the rewards. If discounting is weak, each upwards jump in continuation value after a transition $H \rightarrow L$ may fall short of U^* . Since the continuation value then drifts downwards once the state returns to H , the agent can reach the mature phase only if a sufficiently rapid sequence of positive shocks arrives before U_t reaches zero.

A further implication is that, provided the contract is not terminated by the lack of good news, it eventually implements the first-best, in the sense that the principal's value coincides with the first-best surplus net of the agent's promised utility U . In the current regime, this outcome occurs precisely when a transition pushes the continuation value to the threshold U^* . From that point on, incentive compatibility is enforced at minimal cost, the threat of termination is no longer needed, and the principal sustains efficient effort forever with constant transfers. When the state is L , the first-best is effectively attained as soon as the promised utility is such that, after a transition to H , the agent's continuation value jumps exactly to U^* . From that moment on, both parties behave as if in the stationary first-best phase, with efficient effort permanently implemented. The long-run dynamics therefore admit only two possibilities: either the contract is terminated if no sufficiently informative shocks arrive, or it converges to maturity at U^* , where incentive compatibility ceases to bind and only promise-keeping matters.

For these dynamics to be well defined, certain parameter restrictions are needed. In particular, Proposition 1.2 assumes that $\frac{\pi_1 - \pi_0}{c} v > 1 + \frac{\phi_0}{\phi_1 - \phi_0} \frac{r + \gamma_0 + \phi_1}{r + \gamma_0}$. This restriction ensures that both $V_H(U)$ and $V_L(U)$ are strictly increasing at $U = 0$. Intuitively, it guarantees that initiating the contract in either state is worthwhile for the principal: even when the promised utility is arbitrarily small, the marginal surplus created by a success outweighs the incentive and transition costs on the right-hand side. If the inequality fails, $V_H(U)$ and $V_L(U)$ are decreasing everywhere, and the optimal policy is immediate termination (*i.e.*, no contract is offered at inception). If the recovery rate γ_0 is low, non-initiation becomes more attractive because the prospects of returning quickly to high productivity are weak. The role of transitions is captured by the ratio $\frac{\phi_1}{\phi_0}$. When this ratio is high, transitions are highly informative and the principal can provide incentives efficiently. However, if the ratio falls, transitions lose their power as incentive devices and the condition is more difficult to satisfy, making immediate termination more likely. Holding the ratio constant while reducing ϕ_0 also increases the effective cost of incentive provision: although a lower ϕ_0 means fewer harmful transitions on average, it also makes the informative transitions rarer, such that reaching the mature phase through them becomes less likely before the promise drifts to zero. At the limit $\phi_0 \rightarrow 0$, the contract converges to a stationary contract without termination (provided it remains profitable).

The overall structure of the contract can also be understood by analogy with insurance-type arrangements. Rather than rewarding routine successes, which here carry less information about effort, the contract makes the agent's payoff contingent on the most informative adverse shocks: compensation is triggered by transitions to the low state, while successes only prolong the relationship without paying on impact. In my model this situation reflects the fact that negative shocks are the most informative signals of hidden effort, and the contract is therefore structured around them.

The dynamic implications of this optimal contract are captured in the following corollaries. They show how the relationship unfolds once it starts, depending on whether the initial productivity state is H or L : how the continuation value evolves, when termination may occur, and how payments are structured along the path.

COROLLARY 1.1. *Assume $\theta_0 = H$. The contract begins with a probationary phase. The agent is*

promised an initial continuation value strictly between zero and U^* ,

$$\hat{U}_0^H = U^* \left[1 - \left(\frac{U^*}{\alpha} \frac{r}{r + \gamma_0} \frac{r + \gamma_0 + \phi_1}{(\pi_1 - \pi_0)v - c} \right)^{\frac{1}{\alpha-1}} \right] \in (0, U^*),$$

and the continuation value drifts down deterministically as long as no shock occurs. If no shock is observed before

$$T^\varnothing = -\frac{1}{\kappa} \log \left(1 - \frac{\hat{U}_0^H}{U^*} \right),$$

the continuation value reaches $U_{T^\varnothing} = 0$, and the contract is terminated.

If there is any success $dX_t = 1$ during the probatory phase, such phase is extended by some constant $\Delta T > 0$, independent of the current continuation value.

If instead a transition occurs at some $t < T^\varnothing$, the agent's continuation value jumps to the stationary level U^* and remains there forever (the contract is never terminated from then on). At that instant, the principal pays a lump sum $p^H(U_t) > 0$, and from that point, the contract enters a stationary phase:

- While $\theta_t = L$, the principal pays a flow $w^L(U^*) = rU^* = \frac{\phi_0 c}{\phi_1 - \phi_0}$;
- At every $H \rightarrow L$ transition, a lump sum $p^H(U^*) = \frac{c}{\phi_1 - \phi_0}$ is paid;

If instead the contract begins in the low state, the picture is simpler: no transfers are paid while the agent rests, the promise is kept constant, and termination never occurs. The continuation value changes only when a transition to H arrives, at which point it jumps to the probationary level associated with starting in the high state. From then on the path follows exactly the dynamics in Corollary 1.1: successes may prolong probation but never trigger payments, while only a transition resets the promise to U^* and delivers compensation. In the long term, therefore, the two possibilities are again either termination in the absence of sufficiently informative shocks or maturity at U^* if such a transition arrives in time. This result resembles situations in which an employee on sick leave or in training receives little or no pay during their absence, but once they resume normal work, the same probationary structure applies as if they had been newly hired.

A pertinent limiting case arises when $\gamma_0 = 0$, such that recovery from a downturn is impossible. In that event a transition to the low state represents a permanent collapse of productivity. The optimal contract then resembles an invalidity scheme: at the moment of the accident the principal makes an immediate transfer, after which the agent receives a perpetual “pension,” corresponding to early retirement due to permanent incapacity. Incentive provision becomes irrelevant once productivity is lost forever.

The sample paths of the continuation value in the exhaustion-dominant regime, and the associated sample paths for the cumulative transfers are shown in Figure 1.3. The blue line represents the case in which no shock arrives before the probationary horizon expires; thus, the continuation value drifts down to zero and the agent is fired. The red dashed line shows the case in which a transition $H \rightarrow L$ is the first shock observed: the continuation value jumps immediately to U^* , a transfer is made, and the contract locks into the stationary phase forever. The green dotted-dashed line depicts a history with three successes during probation: each raises the continuation value above its previous level and prolongs the




Figure 1.3: Top panel: Sample paths of the agent's continuation value U_t in the exhaustion-dominant case, $\frac{\pi_1}{\pi_0} < \frac{\phi_1}{\phi_0}$, starting from $\theta_0 = H$ under three scenarios: no shocks (blue), early exhaustion with an immediate jump to U^* (red, dashed), and extended probation with spaced positive shocks followed by exponential exhaustion (green, dotted-dashed). Bottom panel: Corresponding cumulative transfers Y_t , with payments occurring only at success times and bonus sizes increasing with the prevailing continuation value.

probationary phase, until the contract is eventually terminated due to insufficient performance over time. No transfer is ever paid in this particular sample path.

To conclude, in the exhaustion-dominant regime the contract relies primarily on transitions, which are the more informative signals and the ones that eventually trigger payments, while successes play a secondary role by merely prolonging the probationary phase without paying on impact. The long-run outcome is either termination if no informative shocks occur in time, or maturity at U^* after a transition, where the relationship settles into the stationary first-best phase.

1.4.2. Success-Dominant Technology

I now turn to the opposite case, in which successes are the more informative signals of effort,

$$\frac{\phi_1}{\phi_0} < \frac{\pi_1}{\pi_0}.$$

In this regime the logic mirrors that of exhaustion-dominance: the contract relies on the more informative signal—here, successes—to provide incentives, whereas the less informative one—transitions—plays only a supporting role. Successes are tied to bonuses that eventually pay on impact, whereas transitions merely adjust continuation values across states to preserve promises and credibility through periods of low productivity. The formal characterization is given in the following proposition.

PROPOSITION 1.3. *Assume that $\frac{\phi_1}{\phi_0} < \frac{\pi_1}{\pi_0}$, $r \geq \pi_0 \left(1 + \frac{\phi_1 - \phi_0}{\pi_1 - \pi_0}\right)$, and*

$$\frac{\pi_1 - \pi_0}{c} v > 1 + \frac{(r + \gamma_0 + \phi_1)\pi_0}{(r + \gamma_0)(\pi_1 - \pi_0) + \pi_1\phi_0 - \pi_0\phi_1},$$

and let $U^ \equiv \frac{r + \gamma_0}{r} \frac{\pi_0 c}{(r + \gamma_0)(\pi_1 - \pi_0) + \pi_1\phi_0 - \pi_0\phi_1}$. There exists an optimal contract, which is given by the following:*

$$\begin{aligned} \chi^H(U_t) &= \frac{\gamma_0}{r + \gamma_0} (C_2 U^* + (1 - C_2) U_t) - U_t & p^H(U_t) &= 0, \\ \beta^H(U_t) &= \frac{c}{\pi_1 - \pi_0} - \frac{\phi_1 - \phi_0}{\pi_1 - \pi_0} \chi^H(U_t) & b^H(U_t) &= \beta^H(U_t) + U_t - U^*, \\ \chi^L(U_t) &= \begin{cases} \frac{r}{\gamma_0} U_t, & U_t \leq \frac{\gamma_0}{r + \gamma_0} U^*, \\ U^* - U_t, & \frac{\gamma_0}{r + \gamma_0} U^* < U_t \leq U^*, \end{cases} & p^L(U_t) &= 0, \\ w^L(U_t) &= r U^* \cdot \mathbb{1}\{U_t = U^*\} & \beta^L(U_t) &= b^L(U_t) = w^H(U_t) = 0, \end{aligned}$$

where $C_2 = \left(1 - \left(\frac{\pi_1\phi_0 - \pi_0\phi_1}{\phi_1(\pi_1 - \pi_0)}\right)^{\frac{1}{\alpha - 1}}\right) \in (0, 1)$, and $\alpha > 1$ is an implicit function of parameters $(r, \pi_0, \pi_1, \phi_0, \phi_1, \gamma_0)$.

Although the value functions in Figure 1.4 display the same qualitative features as in the exhaustion-dominant case—convergence towards the stationary level U^* and the absence of termination once that level is reached—the instruments that implement them now flip roles. Incentives are powered by successes: whenever a success arrives, the principal pays a bonus $b^H(U)$ and sets the sensitivity to successes to implement effort at minimal per-success cost. Transitions remain essential but do not pay on impact ($p^H \equiv 0$). Instead, χ^H and χ^L are chosen to stabilize the promise across states so that temporary spells of low productivity neither unravel incentives nor violate limited liability.

The structure is therefore the mirror image of the exhaustion-dominant regime. The more informative signal—successes—carries the incentive burden and eventually triggers payments, whereas the less informative signal—transitions—never pays on impact but shapes the path of continuation values. In




Figure 1.4: Value functions in the optimal contract under the success-dominant case: $\frac{\phi_1}{\phi_0} < \frac{\pi_1}{\pi_0}$.

particular, during a low-productivity spell the agent is not penalized: the promise is maintained until productivity returns. Moreover, a cycle $H \rightarrow L \rightarrow H$ acts as a positive, albeit less informative, signal of effort so that after returning to H , the continuation value is reset at a higher level that extends the length of the probatory phase.

The same concavity logic explains why transitions are not discarded altogether in this regime. Relying only on successes would make continuation values jump too sharply, undermining stability. Using transitions to adjust promises during low-productivity spells ensures a smoother trajectory of U , even though only successes ultimately trigger payments.

The dynamic implications mirror this logic. Starting in H with a promise $U \in (0, U^*)$, the continuation value drifts down deterministically in the absence of informative events, and termination occurs if no success arrives in time. A success during the probatory phase jumps the promise up to U^* and pays a bonus on impact; from that moment the relationship remains in the stationary phase with efficient effort in H , rest in L , and no further risk of termination. Transitions $H \rightarrow L$ along the way do not trigger payments; they are accommodated by χ so that the promised utility remains credible and, after a return to H , the promise sits at the level consistent with U^* and continued performance pay. When the state is L , the promise is simply stabilized with no transfers until a transition back to H occurs, at which point the same probationary logic applies. In the long run, therefore, the two possibilities are again either termination if no success arrives within the window, or maturity at U^* following a success, after which incentive compatibility is enforced at minimal cost and only promise-keeping remains relevant.

COROLLARY 1.2. *Assume that $\theta_0 = H$. The contract begins with a probationary phase. The agent is*

promised an initial continuation value strictly between zero and U^* ,

$$\hat{U}_0^H = U^* \left[1 - \left(\frac{U^*}{\alpha} \frac{r}{r + \gamma_0} \frac{r + \gamma_0 + \phi_1}{(\pi_1 - \pi_0)v - c} \right)^{\frac{1}{\alpha-1}} \right] \in (0, U^*),$$

If no success or transition $H \rightarrow L$ occurs before

$$T^\varnothing = -\frac{1}{\kappa} \log \left(1 - \frac{\hat{U}_0^H}{U^*} \right),$$

the continuation value drifts down smoothly until $U_{T^\varnothing} = 0$ and the contract is terminated.

If there is any transition $H \rightarrow L$ during the probatory phase, this phase is paused until the state returns to H again (keeping the agent's continuation value constant). After returning to the H state, the remaining probatory time is extended by an amount $\Delta T > 0$.

If instead a success occurs at some $t < T^\varnothing$, the continuation value jumps to U^* . From then on the contract never terminates. At that instant, the principal pays a bonus $b^H(U_t) > 0$, and the continuation value oscillates between two values:

- As long as the state is H , the continuation value is constant at U^* , and the principal pays a bonus $b^H(U^*) > \frac{c}{\pi_1 - \pi_0}$ with every success.
- As long as the state is L , the continuation value is constant at $\frac{\gamma_0}{r + \gamma_0} U^*$, and no transfers are paid.

In other words, if the relationship starts in the high state, the agent faces a probationary phase similar in spirit to the exhaustion-dominant case: unless sufficiently informative events arrive within a finite window, the contract is terminated. Here the phase ends as soon as the first informative success is observed: a success directly rewards effort, jumps the promise to U^* , and locks the relationship into the stationary phase with efficient effort in H and rest in L . If instead a transition $H \rightarrow L$ occurs during probation, the phase is paused with the promise held fixed; once productivity returns to H , the remaining probation is extended by a constant amount $\Delta T > 0$ (reflecting that an $H \rightarrow L \rightarrow H$ cycle is a weaker but still positive signal of effort). From then on, the first success pushes the promise to U^* , and the relationship becomes stationary. This logic resembles a trial period in which the employer requires some early signal of performance or resilience: a demonstrated success or evidence of recovery after an adverse shock is enough to lock the relationship into a secure, long-term phase. Once in the long-term phase, bonuses are paid in spells of high productivity, but they are higher than they would be absent any extenuation effects to compensate for the spells of low productivity in which there is no compensation.

If the relationship starts in the low state, the picture is simpler and mirrors the exhaustion-dominant case. The agent begins with a smaller promise, $\hat{U}_0^L = \frac{\gamma_0}{r + \gamma_0} \hat{U}_0^H$, receives no transfers while in L , and faces no risk of termination. When a transition $L \rightarrow H$ occurs, the promise jumps to $\hat{U}_0^H$, and from that moment, the path proceeds exactly as in Corollary 1.2: the phase is probationary until the first success arrives, which both pays a bonus on impact and moves the promise to U^* , after which the relationship remains in the stationary phase.

When $\gamma_0 = 0$, the exhaustion-dominant regime delivers an invalidity-style contract with perpetual payments after an accident. The success-dominant regime behaves very differently: if recovery is impossible, a transition $H \rightarrow L$ eliminates the arrival of successes forever, so the agent can never again produce the informative signals that sustain incentives. Since transitions are weaker and now permanent, the principal prefers immediate termination with no compensation if the first shock is an $H \rightarrow L$ event. Similarly, if the relationship starts in L under $\gamma_0 = 0$, it is never initiated, as no success can arrive at any time in the future.




Figure 1.5: Top panel: Sample paths of the agent's continuation value U_t in the success-dominant case, $\frac{\pi_1}{\pi_0} > \frac{\phi_1}{\phi_0}$, starting from $\theta_0 = H$ under three scenarios: no shocks (blue), early success with a temporary setback (red, dashed), and extended probation with eventual promotion (green, dotted-dashed). Bottom panel: Corresponding cumulative transfers Y_t when payments occur only at success times during productive spells, with larger bonuses associated with higher continuation values.

Figure 1.5 illustrates sample paths in the success-dominant regime. The blue line again represents the case with no shocks before the horizon, so the contract is terminated. The red dashed line corresponds to an early success: the continuation value jumps directly to U^* , a bonus is paid, and from then on the relationship is mature, with the continuation value oscillating between U^* when the state is H and $\frac{\gamma_0}{r+\gamma_0}U^* < U^*$ when the state is L . The green dotted-dashed line illustrates the case where a transition

$H \rightarrow L$ occurs before any success. In this phase the continuation value is frozen until the state returns to H , at which point it increases and the probationary phase is extended. A second spell in L has the same effect. Eventually a success arrives while in H , which both pays a bonus and pushes the continuation value to U^* , locking the contract into the stationary phase.

To summarize, in the success-dominant regime the more informative signal (successes) carries the incentive burden and eventually triggers payments, while the less informative signal (exhaustion) shapes the path of continuation values without paying on impact. The long-run outcome is either termination if no success arrives in time, or maturity at U^* following a success, after which incentive compatibility no longer binds and promise-keeping becomes the relevant constraint. This situation mirrors the exhaustion-dominant case with the roles of signals reversed.

1.4.3. Optimality of a^* under private effort

The efficient effort process a^* prescribes maximal effort in state H and minimal effort in state L until termination. While many moral hazard settings feature underprovision of effort—since reducing effort can reduce the agent’s rent—the principal here has no incentive to do so. To demonstrate this point, suppose that in state H the principal induces zero effort ($a_t = 0$) when the agent’s continuation value is U . The success intensity then falls to π_0 , and the transition rate to L falls to ϕ_0 . Ignoring the now-slack incentive constraint $(\pi_1 - \pi_0)\beta_t^H + (\phi_1 - \phi_0)\chi^H \leq c$, the principal optimally sets all the transfers and the sensitivity to successes to zero, $\beta^H(U) = b^H(U) = p^H(U) = w^H(U) = 0$, and chooses the continuation adjustment $\chi^H(U) = -\frac{r}{r+\gamma_0}U \leq 0$ so that her marginal value is equalized before and after a transition. The principal does not gain from such a deviation if

$$rV_H(U) \geq \pi_0 v + V'_H(U) \left[rU + \phi_0 \frac{r}{r+\gamma_0} U \right] + \phi_0 \left[V_L \left(\frac{\gamma_0}{r+\gamma_0} U \right) - V_H(U) \right], \quad (1.4.7)$$

which is precisely the optimality condition for high effort in the HJB equation (1.4.5). The concavity of V_H and the fact that $V_H(U) \geq V_L \left(\frac{\gamma_0}{r+\gamma_0} U \right)$ for all U ensure that this inequality holds for all $U \in [0, U^*]$ since both

$$V_H(U) - V_H(0) - UV'_H(U) \geq 0 \quad \text{and} \quad V_H(U) - V_L \left(\frac{\gamma_0}{r+\gamma_0} U \right) - \left(U - \frac{\gamma_0 U}{r+\gamma_0} \right) V'_H(U) \geq 0.$$

Hence, the principal strictly prefers to induce maximal effort in state H at all continuation values, confirming the optimality of the efficient effort process.

1.5. EXTENSIONS AND ROBUSTNESS

In this section I relax some of the modelling and analysis assumptions and discuss how the optimal contract changes when these assumptions are relaxed. First, I address how to modify the contract to remain incentive compatible when the productivity state (and thus transitions) is private information of the agent; thus, the principal must rely on the agent’s reports. Second, I allow the productivity of effort in the L state to be small but positive, and show that the efficient effort process and the optimal contract implementing it are the same as in the baseline if the productivity in the L state is small

enough. Third, I discuss the implications when the evolution of the productivity status is determined by the accumulation of past effort rather than by the current effort only.

1.5.1. Private productivity state

A key assumption of my model is that the productivity status is public and can be contracted on. This assumption fits better examples where the exhaustion events are verifiable accidents or injuries. However, it is not difficult to think of cases where exhaustion and recovery can be faked by the agent. For instance, psychological burnout is typically not easy for the employer to detect. Therefore, employee must convince the employer that he or she is unable to work. Because the principal has commitment power when designing the contract, the Revelation Principle states that the contract can incorporate reports of the agent about the productivity state once this state is his private information, and that we can look for a contract that induces truthful revelation of the state on the equilibrium path.

Assuming that the optimal contract still exhibits the structure of the probationary phase and either termination or promotion, it is not difficult to adapt the contract to make the effort profile a^* and truthful revelation of the state θ_t incentive compatible. Indeed, to induce truthful revelation of the state, the contract cannot reward or punish any report. Otherwise, the agent would overreport the transitions $H \rightarrow L$ if he can obtain some reward by doing so, as in the contracts described in Section 1.4. If, instead, the contract were to punish those reports to prevent opportunistic overreporting, the agent would be tempted to hide the transition to still benefit from eventual successes that happen with positive probability in any history. In other words, the sensitivities to the transitions (in this case, to the reports made by the agents) and the associated transfers must satisfy $\chi^H(U) = p^H(U) = \chi^L(U) = p^L(U) = 0$. All the load of incentives should then be put in the sensitivity of the agent's continuation value to successes when the state is H : $\beta^H = \frac{c}{\pi_1 - \pi_0}$. Paying success-contingent bonuses during these H -spells are the best reward that the principal can offer. To maintain the incentive to report truthfully, the principal can commit to pay a flow transfer $w^L(U) = rU \geq 0$ while the state is reported to be $\theta_t = L$.

To illustrate that such a contract satisfies incentive compatibility, consider first the case in which the agent's continuation value is U^* ,⁹ so that the mature, permanent phase has been achieved. If the agent's true state is $\theta_t = H$ and reports so, the agent expects to obtain bonuses $b^H(U^*) = \frac{c}{\pi_1 - \pi_0}$ with intensity π_1 and pays the effort cost c , until a transition to the state L is realized and reported, providing the agent continuation value U^* thereafter. Thus, the agent at time t is indeed delivered continuation value $U_t = U^*$. If instead the agent reports to be in state L , he would equally obtain continuation value U^* , so no misleading report is profitable. The same argument works for the case in which the true state is L .

Now consider that the contract is still in the probationary phase, at some time t in which the agent is promised a continuation value $U_t \in (0, U^*)$. On the equilibrium path, if the state is H , the continuation value drifts downwards until it hits the zero lower bound or until a success is realized. If the state is L , the probationary phase is paused, and the continuation value remains constant at U_t until a transition to H is realized. In contrast to the contract described in Proposition 1.3, the probationary phase is resumed from that moment on without any extension. If there were any extension associated with it,

⁹In this case, we have $U^* = \frac{\pi_0}{r} \frac{c}{\pi_1 - \pi_0}$.

the agent would like to repeatedly report transitions between H and L to extend for free the length of the probationary phase. If the agent with a true state H were to report to be in state L between time t and time $\tilde{\tau}$, the agent would obtain the flow transfer as long as the state does not switch:

$$\int_t^{\tilde{\tau}^-} e^{-(r+\phi_0)(s-t)} \left(\underbrace{rU_t}_{=w^L(U_t)} + \underbrace{\phi_0 U_t}_{\text{actual switch to } L} \right) ds + e^{-(r+\phi_0)(\tilde{\tau}-t)} U_t = U_t \quad \text{for all } \tilde{\tau}.$$

Thus, the agent has no incentive to falsely report being in state L . Instead, if the agent in true state L were to report being in state H , the probationary phase would not stop and the contract would behave as if the continuation value were drifting downwards: $\tilde{U}_s = U^* - (U^* - U_t)e^{r(s-t)}$ for $s \in (t, \tilde{\tau})$. The agent could then obtain a success with intensity π_0 , earning him a bonus $b^H(\tilde{U}_s)$ and a jump to the permanent phase, as long as there no actual switch to the H state. The expected payoff he would obtain is as follows:

$$\int_t^{\tilde{\tau}^-} e^{-(r+\pi_0+\gamma_0)(s-t)} \left[\pi_0 \left(\underbrace{\frac{c}{\pi_1 - \pi_0} + \tilde{U}_s - U^* + U^*}_{=b^H(\tilde{U}_s)} \right) + \gamma_0 \tilde{U}_s \right] ds + e^{-(r+\pi_0+\gamma_0)(\tilde{\tau}-t)} \tilde{U}_{\tilde{\tau}} = U_t \quad \text{for all } \tilde{\tau}.$$

Hence the agent has no incentive to misreport being in state H when the true state is L . Indeed, such a contract is actually the solution to the system of HJB equations where the incentive compatibility constraints are $\chi^H(U) = p^H(U) = \chi^L(U) = p^L(U) = 0$ and $w^L(U) = rU$. The value functions admit a closed-form solution:

$$\begin{aligned} V_H(U) &= V_H^{fb}(U) + \left(W(H) - \frac{\pi_0 v}{r} \right) \left(1 - \frac{U}{U^*} \right)^{1 + \frac{\pi_1}{r} + \frac{\phi_1}{r+\gamma_0}}, \\ V_L(U) &= V_L^{fb}(U) + \left(W(L) - \frac{\pi_0 v}{r} \right) \left(1 - \frac{U}{U^*} \right)^{1 + \frac{\pi_1}{r} + \frac{\phi_1}{r+\gamma_0}}. \end{aligned}$$

In this case, both value functions achieve the first-best payoff only at $U = U^*$ in either state. A representation of these value functions can be seen in Figure 1.6.

In Figure 1.7 we also display three sample paths of the continuation value at the optimal contract when starting at $\theta_0 = H$. The blue line represents the evolution when no shock occurs before the deadline of the probationary phase, so that the agent is fired. The red dashed line represents the case in which a success is the first realized shock, thus inducing the promotion of the agent to the permanent phase with constant continuation value U^* . The green dotted-dashed line represents the case in which several transitions are realized (and are reported by the agent). The continuation value remains constant during the spells in the L state, and drifts downwards when the state returns to H .

1.5.2. Positive productivity in the L state

In the baseline model I adopted the extreme assumption that effort does not affect output in state L . Indeed this situation can be seen as a rather restrictive case, since for instance mild illness may reduce




Figure 1.6: Value functions in the optimal contract when θ_t is private information.

but not completely annihilate productivity. For some $\delta \in (0, 1)$, consider the following:

$$p_t(a_t, L) = \pi_0 + (\pi_1 - \pi_0)\delta a_t,$$

and keep all of the other baseline model features unchanged. The lower δ is, the more severe the state L is for productivity. Setting $\delta = 0$ corresponds to the baseline model, and setting $\delta = 1$ corresponds to the case in which exhaustion has no effect on productivity. The efficient effort process with interior δ can therefore involve positive effort in the L state if the loss in productivity is not significant enough.

Assuming that $(\pi_1 - \pi_0)v > c$, inducing $a_t = 1$ when the state is $\theta_t = H$ is optimal for achieving aggregate welfare. Define a cutoff $\hat{\delta} \in (0, 1)$ as follows:

$$\hat{\delta} = \frac{c}{(\pi_1 - \pi_0)v} \frac{r + \gamma_1 + \phi_1}{r + \gamma_0 + \phi_1} + \frac{\gamma_0 - \gamma_1}{r + \gamma_0 + \phi_1}.$$

If the loss in productivity derived from exhaustion satisfies $\delta < \hat{\delta}$, the optimal effort process sets $a_t = 0$ if $\theta_t = L$. In that case, the optimal contracts derived in section 1.4 remain unchanged and optimal, since the incentive compatibility constraint in state L remains slack. As mentioned, inducing rest is easy and inexpensive for the principal. The extreme baseline assumption $\delta = 0$ is thus a useful simplification for the case of sufficiently small δ . This situation is more likely to occur when the recovery rate γ_0 is high and the discount rate r is small (these two increase $\hat{\delta}$) since the option of resting in state L and waiting for a quick transition back to H overcomes the short-run benefit of exerting effort when exhausted. Instead, high recovery rate γ_1 and high ratio $\frac{(\pi_1 - \pi_0)v}{c}$ make rest less desirable.

If instead the loss of productivity is such that $\delta > \hat{\delta}$, the efficient effort process sets maximal effort




Figure 1.7: Top panel: Sample paths of the agent's continuation value U_t starting from $U_0 < U^*$ and $\theta_0 = H$ under three scenarios: no shocks (blue), early success (red, dashed), and pauses/exhaustion spells (green, dotted-dashed). Bottom panel: Corresponding cumulative transfers Y_t , illustrating how bonuses and rest-period payments are delivered along each path.

when the state is L as well. Inducing that effort level when effort is unobserved might however be too costly for the principal. There are two reasons for this statement. First, transitions from the L state to the H state are, if anything, a bad signal about effort, since high effort decreases the likelihood of those transitions. Positive rewards linked to those events induce the agent to exert less effort. Second, although successes are positive signals about effort, the informativeness load in those signals is diminished compared with that of successes in state H : $\frac{\delta\pi_1 + (1-\delta)\pi_0}{\pi_0} < \frac{\pi_1}{\pi_0}$. A worse signal forces the principal to concede higher rewards for the same effort level, thus augmenting the limited liability rent that the agent obtains.

1.5.3. Accumulated fatigue

In shift work and safety-critical jobs, exhaustion usually reflects accumulated load and missed rest. Sustained work builds a stock of fatigue, while rest enable that stock to dissipate; the higher the accumulated fatigue is, the easier it is to become exhausted and the more difficult it is to bounce back.

To capture this situation, one can replace effort-driven transitions with a privately evolving fatigue stock A_t that increases with effort and decays with rest. Hence the hazards $H \rightarrow L$ and $L \rightarrow H$ depend on A_t . This modelling choice is close in spirit to that of [Sakamoto \(2014\)](#).

This change introduces path dependence into the observable regime. Under current effort transitions, θ_t is a two-state Markov chain with memoryless (exponential) spell lengths. With fatigue, an outside observer would detect duration dependence: the hazard of leaving H increases over the spell as fatigue accumulates, and the hazard of leaving L increases as fatigue dissipates. The efficient effort rule remains simple: work in H , rest in L . Thus, A_t drifts upwards in productive spells and downwards in unproductive spells, generating endogenous cycles.

This addition of accumulated fatigue introduces a new informational difficulty: because effort is hidden, the stock A_t is also hidden. If the agent deviates, he privately knows how fatigue has evolved and thus forecasts transitions more accurately than the principal does. The environment therefore combines hidden action with a hidden state (private information). Despite these technical difficulties, the study of optimal stationary contracts remains tractable, which can be found in the Online Appendix.

From an informativeness perspective, the relevant signal changes with the level of fatigue. If A_t were known (or correctly inferred) to be low, an observable transition from H to L is a rare event. It typically requires a long history of high effort to accumulate sufficient fatigue. In that case, a transition is a strong signal of past effort relative to success. In contrast, when A_t is high, exhaustion is likely in the short term even if the agent shirks, so transitions become weakly informative while successes—which load directly on current effort via $\pi_1 > \pi_0$ —correspond to the higher likelihood ratio. Hence, on the equilibrium path, as fatigue builds within an H -spell and then dissipates in L , the signal that should be rewarded more intensely shifts over time: transitions early in H , successes late in H .

1.6. CONCLUSION

In this paper I study dynamic contracts in environments where effort increases performance as well as the risk of exhaustion. When exhaustion is observed, it carries a double message: it is bad news for near-term performance but a sign that the worker has been trying hard. Optimal contracts thus reward such episodes to induce effort. If exhaustion is the more informative signal of effort, pay is delivered at the onset of, and during, those low-productivity periods. Anticipating this promise of downtime pay encourages the worker to exert high effort. If successes are more informative, pay takes the form of larger bonuses when good results arrive. Those bonuses are larger when exhaustion episodes are more likely or last longer—compensating for the loss of income during those periods.

Beyond exhaustion, one may consider several settings where the setup analyzed here can be applied. In software and product work, a push to add new features is followed by a stabilization phase—testing, integration, and fixes—that adds little new output by itself but is essential to keep creating value. In work that relies on high motivation and deep focus, periods of rest are necessary to restore both. In tenant farming, heavy cultivation within a season raises the harvest and is then followed by a period of fallow or lighter cultivation on the same plot. In manufacturing, running equipment at very high use lifts output on the day but brings forward cleaning and maintenance that interrupt production.

Recognising that high- and low-productivity phases are both part of the technology—and that effort shapes the timing of each—is essential for optimal incentive design.

BIBLIOGRAPHY

- I. Albulescu, L. Macsinga, A. Rusu, C. Sulea, M. Bodnaru, and B.T. Tulbure (2022). “Give me a break! A systematic review and meta-analysis on the efficacy of micro-breaks for increasing well-being and performance.” *PLOS ONE*, 17(8).
- S. Bavafa and I. Jónasson (2024). The Distributional Impact of Fatigue on Performance. *Management Science*, 70(5), 3319–3337.
- M. Baucells and L. Zhao (2019). It Is Time to Get Some Rest. *Management Science*, 65(4), 1717–1734.
- S. E. Bechtold, R. E. Janaro, and D. W. L. Sumners (1984). Maximization of Labor Productivity Through Optimal Rest-Break Schedules. *Management Science*, 30(12), 1442–1458.
- B. Bernard (2024). Continuous-Time Games with Imperfect and Abrupt Information. *Review of Economic Studies*, 91(4), 1988–2052.
- B. Bernard (2025). Continuous-Time Stochastic Games with Imperfect Public Monitoring. *Working paper*, University of Wisconsin–Madison.
- B. Biais, T. Mariotti, J.-C. Rochet, and S. Villeneuve (2010). Large Risks, Limited Liability, and Dynamic Moral Hazard. *Econometrica*, 78(1), 73–118.
- S. Board and M. Meyer-ter-Vehn (2022). A Reputational Theory of Firm Dynamics. *American Economic Journal: Microeconomics*, 14(2), 44–80.
- M. Cabral and Z. B. Cullen (2019). Estimating the Value of Public Insurance Using Complementary Private Insurance. *American Economic Journal: Economic Policy*, 11(2), 38–73.
- C. H. Colla, W. H. Dow, A. Dube, and V. Lovell (2014). Early Adoption of the *San Francisco* Paid Sick Leave Ordinance. *Journal of Policy Analysis and Management*, 33(4), 1026–1056.
- M. Collewet and J. Sauermann (2017). Working Hours and Productivity. *Labour Economics*, 47, 96–106.
- P. M. DeMarzo and Y. Sannikov (2017). Learning, Termination, and Payout Policy in Dynamic Incentive Contracts. *Review of Economic Studies*, 84(1), 182–236.
- D. Dragone (2009). I Am Getting Tired: Effort and Fatigue in Intertemporal Decision-Making. *Journal of Economic Psychology*, 30(4), 552–562.
- A. Fernandes and C. Phelan (2000). A Recursive Formulation for Repeated Agency with History Dependence. *Journal of Economic Theory*, 91(2), 223–247.
- M. Fischer, J. Geyer, and N. R. Ziebarth (2024). Subsidizing Disability Insurance: Usage, Beneficiaries, and (Lack of) Spillovers. NBER Working Paper No. 32198.

- J. R. Guardado and N. R. Ziebarth (2019). Worker Investments in Safety, Workplace Accidents, and Compensating Wage Differentials. *International Economic Review*, 60(1), 133–155.
- M. Halac and A. Prat (2016). Managerial Attention and Worker Performance. *Journal of Political Economy*, 124(3), 637–682.
- B. Holmström (1979). Moral Hazard and Observability. *The Bell Journal of Economics*, 10(1), 74–91.
- M. R. Ibanez and M. W. Toffel (2020). Managing Workload to Reduce Worker Injury. *Management Science*, 66(10), 4397–4420.
- D. S. Kc and C. Terwiesch (2009). The Impact of Workload on Service Time and Patient Safety: An Econometric Analysis of Hospital Operations. *Management Science*, 55(9), 1486–1498.
- K. Lavetti (2023). Compensating Wage Differentials in Labor Markets: Empirical Challenges and Applications. *Journal of Economic Perspectives*, 37(3), 189–212.
- G. Maideu-Morera (2024). Amenity-Biased Technical Change. *Working paper*
- D. Martimort and S. Straub (2016). How to Design Infrastructure Contracts in a Warming World: A Critical Appraisal of Public–Private Partnerships. *International Economic Review*, 57(1), 61–88.
- A. Mas (2025). Non-Wage Amenities. *Handbook of Labor Economics*, Volume 6, 373–446. Elsevier.
- D. Nocetti (2008). Industrial Fatigue Redux. *Economics Letters*, 99, 286–289.
- H. Sakamoto (2014). Dynamic Resource Management under the Risk of Regime Shifts. *Journal of Environmental Economics and Management*, 68(1), 1–19.
- M. Sáez-Martí (2025). Take a Break: A Model of Fatigue, Recovery, and the Economics of Remote Work. Cowles Foundation Discussion Paper No. 2268R1, Yale University.
- Y. Sannikov (2008). A Continuous-Time Version of the Principal–Agent Problem. *Review of Economic Studies*, 75(3), 957–984.
- Y. Sannikov (2014). Moral Hazard and Long-Run Incentives. *Working paper*.
- Seibold, A., Seitz, S. and Siegloch, S. (2025). Privatizing Disability Insurance. *Econometrica*, 93(5), 1697–1737.
- P. Tucker (2003). The Impact of Rest Breaks upon Accident Risk and Performance. *Work & Stress*, 17(2), 123–137.
- U.S. Bureau of Labor Statistics (2024). *Employee Benefits in the United States—March 2024*. Fact Sheet.
- M. Vander Weerdt, P. Stoddard-Dare, and L. DeRigne (2023). Is Paid Sick Leave Bad for Business? A Systematic Review. *American Journal of Industrial Medicine*, 66(6), 429–440.
- J. Y. Zhu (2013). Optimal Contracts with Shirking. *Review of Economic Studies*, 80(2), 812–839.

Appendices

1.A. MAIN PROOFS

Proof of Proposition 1.1. Let me first write the objective the expected payoff of the agent with such a contract:

$$\begin{aligned}
 U(a; \mathcal{C}) = & w_1 + (\pi_0 + (\pi_1 - \pi_0)a_1)b_1 - ca_1 \\
 & + (\pi_0 + (\pi_1 - \pi_0)a_1) \cdot (\phi_0 + (\phi_1 - \phi_0)a_1)w_2^{SL} \\
 & + (1 - \pi_0 - (\pi_1 - \pi_0)a_1) \cdot (\phi_0 + (\phi_1 - \phi_0)a_1)w_2^{FL} \\
 & + (\pi_0 + (\pi_1 - \pi_0)a_1) \cdot (1 - \phi_0 - (\phi_1 - \phi_0)a_1) \\
 & \quad \cdot [w_2^{SH} + (\pi_0 + (\pi_1 - \pi_0)a_2^{SH})b_2^{SH} - ca_2^{SH}] \\
 & + (1 - \pi_0 - (\pi_1 - \pi_0)a_1) \cdot (1 - \phi_0 - (\phi_1 - \phi_0)a_1) \\
 & \quad \cdot [w_2^{FH} + (\pi_0 + (\pi_1 - \pi_0)a_2^{FH})b_2^{FH} - ca_2^{FH}]
 \end{aligned}$$

Since there is limited liability, the agent can always shirk and obtain a nonnegative payoff, so that one can ignore the participation constraint of the agent. Consequently, it is optimal for the principal to set $w_1 = 0$, as it does not affect incentive to exert effort and is not required to meet the participation constraint.

Now consider the incentive compatibility constraints for inducing effort in period 2 when the agent is productive ($\theta_2 = H$). The payoff function is affine in both a_2^{SH} and a_2^{FH} , so the corresponding incentive compatibility constraints are:

$$a_2^{jH} = \begin{cases} 1 & \text{if } b_2^{jH} > \frac{c}{\pi_1 - \pi_0} \\ \in \{0, 1\} & \text{if } b_2^{jH} = \frac{c}{\pi_1 - \pi_0} \\ 0 & \text{if } b_2^{jH} < \frac{c}{\pi_1 - \pi_0} \end{cases} \quad j \in \{S, F\}$$

Notice that the bonus b_2^{jH} is the only incentive tool to induce a_2^{jH} , a it is enough to set it at $\frac{c}{\pi_1 - \pi_0}$ to induce $a_2^{jH} = 1$. Any excess bonus is helpful only to provide incentives for effort in period 1, but one can achieve the same incentive by increasing w_2^{jH} . Thus, without loss of generality I will consider only $b_2^{jH} = 0$ to induce $a_2^{jH} = 0$ and $b_2^{jH} = \frac{c}{\pi_1 - \pi_0}$ to induce $a_2^{jH} = 1$.

Now consider the incentive compatibility condition for effort in period 1. The agent optimally chooses

$a_1 = 1$ if and only if:

$$\begin{aligned}
& (\pi_1 - \pi_0)b_1 + [\pi_1\phi_1 - \phi_0\pi_0]w_2^{SL} + [(1 - \pi_1)\phi_1 - (1 - \pi_0)\phi_0]w_2^{FL} \\
& + [\pi_1(1 - \phi_1) - \pi_0(1 - \phi_0)] [w_2^{SH} + (\pi_0 + (\pi_1 - \pi_0)a_2^{SH})b_2^{SH} - ca_2^{SH}] \\
& + [(1 - \pi_1)(1 - \phi_1) - (1 - \pi_0)(1 - \phi_0)] [w_2^{FH} + (\pi_0 + (\pi_1 - \pi_0)a_2^{FH})b_2^{FH} - ca_2^{FH}] \geq c
\end{aligned} \tag{IC_1}$$

In order to induce $a_1 = 0$, the principal can set optimally $b_1 = w_2^{SL} = w_2^{FL} = w_2^{SH} = w_2^{FH} = 0$. However, the left-hand side of the above inequality can still be positive if the principal sets $b_2^{SH} = \frac{c}{\pi_1 - \pi_0}$ and $b_2^{FH} = 0$, and the condition $\pi_1(1 - \phi_1) > \pi_0(1 - \phi_0)$ is satisfied. Alas, even in that case the left-hand side does not exceed the right-hand side:

$$[\pi_1(1 - \phi_1) - \pi_0(1 - \phi_0)] \frac{c\pi_0}{\pi_1 - \pi_0} = \pi_0 c \left[1 - \frac{\phi_1\pi_1 - \phi_0\pi_0}{\pi_1 - \pi_0} \right] < c$$

Thus, one needs to find which is the most efficient tool amongst b_1 , w_2^{SL} , w_2^{FL} , w_2^{SH} and w_2^{FH} to induce a_1 . I build a Lagrangean to minimize the cost of the principal subject to (IC₁):

$$\begin{aligned}
\mathcal{L} = & -\pi_1 b_1 - \pi_1 \phi_1 w_2^{SL} - (1 - \pi_1) \phi_1 w_2^{FL} - \pi_1 (1 - \phi_1) w_2^{SH} - (1 - \pi_1) (1 - \phi_1) w_2^{FH} \\
& + \lambda \left[(\pi_1 - \pi_0) b_1 + [\pi_1 \phi_1 - \phi_0 \pi_0] w_2^{SL} + [(1 - \pi_1) \phi_1 - (1 - \pi_0) \phi_0] w_2^{FL} \right. \\
& \left. + [\pi_1 (1 - \phi_1) - \pi_0 (1 - \phi_0)] w_2^{SH} + [(1 - \pi_1) (1 - \phi_1) - (1 - \pi_0) (1 - \phi_0)] w_2^{FH} \right]
\end{aligned}$$

where I have omitted all the irrelevant terms in this optimisation problem for the sake of exposition. The transfer w_2^{SL} combines two positive signals for effort, and thus it is the most efficient one. To see this, notice that:

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial w_2^{SL}} &= -\pi_1 \phi_1 + \lambda [\pi_1 \phi_1 - \pi_0 \phi_0] = 0 \\
\iff \lambda &= \lambda^* \equiv \frac{\pi_1 \phi_1}{\pi_1 \phi_1 - \pi_0 \phi_0} > 0
\end{aligned}$$

And plugging such value for $\lambda = \lambda^*$ in the derivatives with respect to the remaining control variables, a negative value is obtained for any parametrisation, and hence it is optimal to set them equal to zero:

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial b_1} \Big|_{\lambda=\lambda^*} &= -\pi_1 + \frac{\pi_1 \phi_1}{\pi_1 \phi_1 - \pi_0 \phi_0} [\pi_1 - \pi_0] < 0 \\
\frac{\partial \mathcal{L}}{\partial w_2^{FL}} \Big|_{\lambda=\lambda^*} &= -\phi_1(1 - \pi_1) + \frac{\pi_1 \phi_1}{\pi_1 \phi_1 - \pi_0 \phi_0} [\phi_1(1 - \pi_1) - \phi_0(1 - \pi_0)] < 0 \\
\frac{\partial \mathcal{L}}{\partial w_2^{SH}} \Big|_{\lambda=\lambda^*} &= -\pi_1(1 - \phi_1) + \frac{\pi_1 \phi_1}{\pi_1 \phi_1 - \pi_0 \phi_0} [\pi_1(1 - \phi_1) - \pi_0(1 - \phi_0)] < 0 \\
\frac{\partial \mathcal{L}}{\partial w_2^{SL}} \Big|_{\lambda=\lambda^*} &= -(1 - \pi_1)(1 - \phi_1) + \frac{\pi_1 \phi_1}{\pi_1 \phi_1 - \pi_0 \phi_0} [(1 - \pi_1)(1 - \phi_1) - (1 - \pi_0)(1 - \phi_0)] < 0
\end{aligned}$$

So w_2^{SL} is the only positive transfer, and is set to meet condition (IC₁) with equality. There are four

cases depending on whether $a_2^{jH} = 0$ or $a_2^{jH} = 1$ is induced in period 2 for each $j \in \{S, F\}$:

	$b_2^{SH} = \frac{c}{\pi_1 - \pi_0}$	$b_2^{SH} = 0$
$b_2^{FH} = \frac{c}{\pi_1 - \pi_0}$	$w_2^{SL} = \frac{c}{\pi_1 - \pi_0} \frac{\pi_1 - \pi_0 [1 - (\phi_1 - \phi_0)]}{\pi_1 \phi_1 - \pi_0 \phi_0}$	$w_2^{SL} = \frac{c}{\pi_1 - \pi_0} \frac{\pi_1 - \pi_0 [1 - (\phi_1 - \phi_0 + \pi_1(1 - \phi_1) - \pi_0(1 - \phi_0))]}{\pi_1 \phi_1 - \pi_0 \phi_0}$
$b_2^{FH} = 0$	$w_2^{SL} = \frac{\pi_0 c}{\pi_1 - \pi_0} + \frac{(1 - \pi_0)c}{\pi_1 \phi_1 - \pi_0 \phi_0}$	$w_2^{SL} = \frac{c}{\pi_1 \phi_1 - \pi_0 \phi_0}$

It remains to show which effort profile the principal wants to induce given the corresponding transfer schemes. I need to consider eight possible profiles $(a_1, a_2^{SH}, a_2^{FH}) \in \{0, 1\}^3$. Denote by $V(a_1, a_2^{SH}, a_2^{FH})$ the principal's payoff from the optimal contract that induces profile $(a_1, a_2^{SH}, a_2^{FH})$.

Recall that I assumed that the optimal second-best profile absent any fatigue effect ($\phi_1 = \phi_0 = 0$) is $(1, 1, 1)$:

$$\frac{\pi_1 - \pi_0}{c} v \geq 1 + \frac{\pi_0}{\pi_1 - \pi_0} \frac{1}{1 - \pi_1} \quad (1.A.1)$$

In particular, equation (1.A.1) implies that inducing maximal effort in period 2 only is preferred to inducing some effort in period 2 only, and to induce no effort at all, so I can ignore the profiles $(0, 0, 0)$, $(0, 0, 1)$ and $(0, 1, 0)$.

Now, I prove that inducing effort in period 1 only is preferred to induced maximal effort in period 2 only:

$$\begin{aligned} V(1, 0, 0) > V(0, 1, 1) &\iff \pi_1 v - \pi_1 \phi_1 \frac{c}{\pi_1 \phi_1 - \pi_0 \phi_0} + \pi_0 v > \pi_0 v + \phi_0 \pi_0 v + (1 - \phi_0) \pi_1 \left(v - \frac{c}{\pi_1 - \pi_0} \right) \\ &\iff \frac{\pi_1 - \pi_0}{c} v > \frac{\pi_1}{\phi_0} \left(\frac{\phi_1}{\pi_1 \phi_1 - \pi_0 \phi_0} - \frac{1 - \phi_0}{\pi_1 - \pi_0} \right) \end{aligned}$$

The right-hand side is decreasing in ϕ_1 , so it can be bounded from above by:

$$\frac{\pi_1}{\phi_0} \left(\frac{\phi_1}{\pi_1 \phi_1 - \pi_0 \phi_0} - \frac{1 - \phi_0}{\pi_1 - \pi_0} \right) \leq \frac{\pi_1}{\phi_0} \left(\frac{\phi_0}{\pi_1 \phi_0 - \pi_0 \phi_0} - \frac{1 - \phi_0}{\pi_1 - \pi_0} \right) = 1 + \frac{\pi_0}{\pi_1 - \pi_0}$$

which is smaller than the right-hand side of equation (1.A.1). Thus, I only need to consider profiles such that $a_1 = 1$.

I first consider the profiles $(1, 1, 1)$, $(1, 1, 0)$ and $(1, 0, 0)$; and I then verify that the profile $(1, 0, 1)$ cannot improve with respect to the three simultaneously. Profile $(1, 1, 1)$ is preferred to $(1, 1, 0)$ if and only if:

$$V(1, 1, 1) \geq V(1, 1, 0) \iff \frac{\pi_1 - \pi_0}{c} v \geq 1 + \frac{\pi_0}{\pi_1 - \pi_0} \frac{1}{1 - \pi_1} \frac{\pi_1 \phi_1 (1 - \pi_0 - \phi_0) - \pi_0 \phi_0 (1 - \pi_1 - \phi_1)}{(1 - \phi_1)(\pi_1 \phi_1 - \pi_0 \phi_0)} \quad (1.A.2)$$

Profile $(1, 1, 1)$ is preferred to $(1, 0, 0)$ if and only if:

$$V(1, 1, 1) \geq V(1, 0, 0) \iff \frac{\pi_1 - \pi_0}{c} v \geq 1 + \frac{\pi_0}{\pi_1 - \pi_0} \frac{\pi_1 \phi_1 (1 - \phi_0) - \pi_0 \phi_0 (1 - \phi_1)}{(1 - \phi_1)(\pi_1 \phi_1 - \pi_0 \phi_0)} \quad (1.A.3)$$

The right-hand side of equation (1.A.2) can be shown to be larger than that of (1.A.3). Profile $(1, 1, 0)$

is preferred to $(1, 0, 0)$ if and only if:

$$V(1, 1, 0) \geq V(1, 0, 0) \iff \frac{\pi_1 - \pi_0}{c} v \geq 1 + \frac{\pi_0}{\pi_1 - \pi_0} \frac{(\phi_1 - \phi_0)\pi_0}{(1 - \phi_1)(\pi_1\phi_1 - \pi_0\phi_0)} \quad (1.A.4)$$

The right-hand side of equation (1.A.4) can be shown to be smaller than those of (1.A.2) and (1.A.3). Therefore, I can rank the profiles as follows: profile $(1, 1, 1)$ is preferred if (1.A.2) is satisfied; if (1.A.2) is not satisfied but (1.A.4) is, then $(1, 1, 0)$ is preferred; if none of them are satisfied, then $(1, 0, 0)$ is preferred by the principal.

Cutoff $\underline{\phi}$ is the value of ϕ_1 such that equation (1.A.2) holds with equality. Notice that this equation can be rearranged as a quadratic equation in ϕ_1 . In addition, the right-hand side converges to $+\infty$ as $\phi_1 \rightarrow 1$; and reduces to the right-hand side of equation (1.A.1) when $\phi_1 \rightarrow \phi_0$, so that there is a unique value $\underline{\phi} \in (\phi_0, 1)$ such that equation (1.A.2) holds with equality. The inequality is therefore satisfied when $\phi_1 \leq \underline{\phi}$.

Similarly, cutoff $\bar{\phi}$ is the value of ϕ_1 such that equation (1.A.4) holds with equality. Notice that this equation can also be rearranged as a quadratic equation in ϕ_1 . In addition, the right-hand side converges to $+\infty$ as $\phi_1 \rightarrow 1$; and reduces to $1 + \frac{\pi_0}{\pi_1 - \pi_0}$ when $\phi_1 \rightarrow \phi_0$, which is smaller than the right-hand side of equation (1.A.1), so that there is a unique value $\bar{\phi} \in (\phi_0, 1)$ such that equation (1.A.4) holds with equality. The inequality is therefore satisfied when $\phi_1 \leq \bar{\phi}$.

Finally, noticing that the right-hand side of equation (1.A.4) is strictly increasing in ϕ_1 on $(\phi_0, 1)$, and taking into account the ranking of equations (1.A.2) and (1.A.4), I conclude that $\underline{\phi} < \bar{\phi}$. Hence, the cutoffs satisfy $\phi_0 < \underline{\phi} < \bar{\phi} < 1$.

It remains to show that the profile $(1, 0, 1)$ does not improve with respect to $(1, 1, 1)$, and $(1, 0, 0)$ simultaneously (it can improve on each of them individually, though). If $(1, 0, 1)$ is preferred to $(1, 1, 1)$, then it must be that:

$$V(1, 0, 1) \geq V(1, 1, 1) \iff \frac{\pi_1 - \pi_0}{c} v \leq 1 + \frac{\pi_0}{\pi_1 - \pi_0} \frac{(\phi_1 - \phi_0)\pi_0}{(1 - \phi_1)(\pi_1\phi_1 - \pi_0\phi_0)} \quad (1.A.5)$$

And if $(1, 0, 1)$ is preferred to $(1, 0, 0)$, then it must be that:

$$V(1, 0, 1) \geq V(1, 0, 0) \iff \frac{\pi_1 - \pi_0}{c} v \geq 1 + \frac{\pi_0}{\pi_1 - \pi_0} \frac{1}{1 - \pi_1} \frac{\pi_1\phi_1(1 - \pi_0 - \phi_0) - \pi_0\phi_0(1 - \pi_1 - \phi_1)}{(1 - \phi_1)(\pi_1\phi_1 - \pi_0\phi_0)} \quad (1.A.6)$$

So that one must verify:

$$\begin{aligned} \frac{1}{1 - \pi_1} \frac{\pi_1\phi_1(1 - \pi_0 - \phi_0) - \pi_0\phi_0(1 - \pi_1 - \phi_1)}{(1 - \phi_1)(\pi_1\phi_1 - \pi_0\phi_0)} &\leq \frac{(\phi_1 - \phi_0)\pi_0}{(1 - \phi_1)(\pi_1\phi_1 - \pi_0\phi_0)} \\ &\iff \pi_1 \leq \pi_0 \end{aligned}$$

Which is a contradiction. This concludes the proof. $\square$

Proof of Lemma 1.1. For a given pair (Y, a) of a contract Y and an effort process a , recall the definition of the continuation payoff of the agent as the process $(U_t^a)_t$ that satisfies:

$$U_t^a = \mathbb{E} \left[\int_t^\infty e^{-r(s-t)} (dY_s - ca_s ds) \middle| \mathcal{F}_t \right].$$

Let me also define the expected lifetime utility of the agent conditional on the public information at t :

$$\mathcal{U}_t^a = \mathbb{E} \left[\int_0^\infty e^{-rs} (dY_s - ca_s ds) \middle| \mathcal{F}_t \right] \quad (1.A.7)$$

$$= \int_0^{t-} e^{-rs} (dY_s - ca_s ds) + e^{-rt} U_t^a \quad (1.A.8)$$

Notice that $\mathcal{U}_t^a$ is the expectation of a given random variable conditional on $\mathcal{F}_t$, and hence the process $(\mathcal{U}_t^a)_t$ is a $\mathcal{F}_t$ -martingale under $\mathbb{P}^{a,\theta}$. From the martingale representation theorem:

$$\mathcal{U}_t^a = \mathcal{U}_0^a + \int_0^t e^{-rs} (\beta_s^a dM_s^a + \chi_s^{a,H} dN_s^{a,H} + \chi_s^{a,L} dN_s^{a,L}) \quad (1.A.9)$$

for some $\mathcal{F}_t$ -predictable processes $\{(\beta_t^a)_t, (\chi_t^{a,H})_t, (\chi_t^{a,L})_t\}$, where the processes

$$\begin{aligned} M_t^a &= X_t - \int_0^t p_\theta(a_s) ds \\ &= X_t - \int_0^t [\pi_0 + (\pi_1 - \pi_0) a_s \mathbb{1}\{\theta_{s-} = H\}] ds \\ N_t^{a,H} &= J_t^H - \int_0^t [\phi_0 + (\phi_1 - \phi_0) a_s] \mathbb{1}\{\theta_{s-} = H\} ds \\ N_t^{a,L} &= J_t^L - \int_0^t [\gamma_0 + (\gamma_1 - \gamma_0) a_s] \mathbb{1}\{\theta_{s-} = L\} ds \end{aligned}$$

are $\mathcal{F}_t$ -martingales under $\mathbb{P}^{a,\theta}$. Combining the evolution $d\mathcal{U}_t^a$ from equations (1.A.8) and (1.A.9), one obtains the evolution of the continuation value:

$$d\mathcal{U}_t^a = r\mathcal{U}_t^a dt - dY_t + ca_t dt + \beta_t^a dM_t^a + \chi_t^{a,H} \mathbb{1}\{\theta_{t-} = H\} dN_t^{a,H} + \chi_t^{a,L} \mathbb{1}\{\theta_{t-} = L\} dN_t^{a,L}$$

Which is equivalent to the law of motion (1.3.2) when one plugs a^* . Notice that β_t^a measures the sensitivity of the continuation value to successes, and $\chi_t^{a,\theta}$ measures the sensitivity of the continuation payoff to switches in the productivity state θ .

The idea is to show, as usual, that the optimal effort profile a^* induces a $\mathcal{F}$ -martingale $(\mathcal{U}_t^{a^*})_t$, and that any deviation $\hat{a}_t$ induces a $\mathcal{F}$ -supermartingale $(\mathcal{U}_t^{\hat{a}})_t$.

Consider a profile $\hat{a}$ such that $(\hat{a}_s)_s$ is chosen until some time t , and then reverts to $a = (a_s)_s$. The

expected lifetime utility from such profile is:

$$\mathcal{U}_t^{\hat{a}} = \int_0^t e^{-rs} (dY_s - c\hat{a}_s ds) + e^{-rt} U_t^a$$

So that:

$$\begin{aligned} \mathcal{U}_t^{\hat{a}} &= \mathcal{U}_t^a + \int_0^t e^{-rs} c(a_s - \hat{a}_s) ds \\ &= \mathcal{U}_0^a + \int_0^t e^{-rs} (\beta_s^a dM_s^a + \chi_s^{a,H} \mathbb{1}_{\{\theta_{s-} = H\}} dN_s^{a,H} + \chi_s^{a,L} \mathbb{1}_{\{\theta_{s-} = L\}} dN_s^{a,L}) \\ &\quad + \int_0^t e^{-rs} c(a_s - \hat{a}_s) ds \\ &= \mathcal{U}_0^a + \int_0^t e^{-rs} (\beta_s^a dM_s^{\hat{a}} + \chi_s^{a,H} \mathbb{1}_{\{\theta_{s-} = H\}} dN_s^{\hat{a},H} + \chi_s^{a,L} \mathbb{1}_{\{\theta_{s-} = L\}} dN_s^{\hat{a},L}) \\ &\quad + \int_0^t e^{-rs} [\beta_s^a \mathbb{1}_{\{\theta_{s-} = H\}} (\pi_1 - \pi_0)(\hat{a}_s - a_s) + \chi_s^{a,H} \mathbb{1}_{\{\theta_{s-} = H\}} (\phi_1 - \phi_0)(\hat{a}_s - a_s)] ds \\ &\quad + \int_0^t e^{-rs} \chi_s^{a,L} \mathbb{1}_{\{\theta_{s-} = L\}} (\gamma_1 - \gamma_0)(\hat{a}_s - a_s) ds \\ &\quad + \int_0^t e^{-rs} c(a_s - \hat{a}_s) ds \end{aligned}$$

Recall that the processes (β, χ^H, χ^L) are predictable, so under $\mathbb{P}^{\hat{a}, \theta}$, the drift of $\mathcal{U}_t^{\hat{a}}$ has the same sign as:

$$([\beta_t^a (\pi_1 - \pi_0) + \chi_t^{a,H} (\phi_1 - \phi_0)] \mathbb{1}_{\{\theta_{t-} = H\}} + \chi_t^{a,L} (\gamma_1 - \gamma_0) \mathbb{1}_{\{\theta_{t-} = L\}} - c) (\hat{a}_t - a_t) dt$$

Maximal effort is induced in state $\theta_t = H$ if and only if:

$$\beta_t^a (\pi_1 - \pi_0) + \chi_t^{a,H} (\phi_1 - \phi_0) \geq c.$$

That is, the principal can provide incentives to the agent by rewarding or punishing either successes or observable transitions. Minimal effort is induced in state $\theta_t = L$ if and only if:

$$\chi_t^{a,L} (\gamma_1 - \gamma_0) \leq c.$$

In this case, successes cannot be used to provide incentives, since the agent has no influence on them. $\square$

Proof of Proposition 1.2. In this proof I show that closed form candidates for the value functions, together with the candidates to optimal control that are described, indeed solve the system of HJB equations (1.4.5)-(1.4.6). The usual verification step can be found and proved in Lemma 1.2.

Let $U^* = \frac{\phi_0 c}{(\phi_1 - \phi_0)r}$, and consider the following candidates for value functions defined on the domain

$[0, U^*]$:

$$V_H(U) = V_H^{fb}(U) - \frac{[(\pi_1 - \pi_0)v - c](r + \gamma_0)}{(r + \gamma_0 + \phi_1)r} \left[1 - \frac{U}{U^*}\right]^\alpha$$

$$V_L(U) = V_L^{fb}(U) - \frac{[(\pi_1 - \pi_0)v - c]\gamma_0}{(r + \gamma_0 + \phi_1)r} \left[1 - \frac{U}{U^*} \frac{r + \gamma_0}{\gamma_0}\right]^\alpha \mathbb{1} \left\{0 \leq U < \frac{\gamma_0}{r + \gamma_0} U^*\right\}$$

As I will verify later, these value functions are concave and single-peaked at interior values for U . I now check the optimality of the controls, for which I will need the derivatives of the value functions:

$$V_H'(U) = -1 + \frac{[(\pi_1 - \pi_0)v - c](r + \gamma_0)}{(r + \gamma_0 + \phi_1)r} \frac{\alpha}{U^*} \left[1 - \frac{U}{U^*}\right]^{\alpha-1} \geq -1$$

$$V_L'(U) = \begin{cases} -1 + \frac{(\pi_1 - \pi_0)v - c}{r + \gamma_0 + \phi_1} \frac{r + \gamma_0}{r} \frac{\alpha}{U^*} \left[1 - \frac{U}{U^*} \frac{r + \gamma_0}{\gamma_0}\right]^{\alpha-1} > -1 & \text{if } 0 \leq U < \frac{\gamma_0}{r + \gamma_0} U^* \\ -1 & \text{if } \frac{\gamma_0}{r + \gamma_0} U^* \leq U \leq U^* \end{cases}$$

Notice in particular that $V_H'(U^*) = V_L'(U^*) = -1$. Let me first look at the HJB equation for the state $\theta_t = H$. Denote by $\lambda(U) \geq 0$ the multiplier associated to the incentive compatibility constraint for each value of the state variable U . The only nonzero controls are χ^H , p^H and β^H , where $\chi^H(U) - p^H(U) = U^* - U$, $\beta^H(U) = C_1(U^* - U)$, and C_1 is a constant to be specified later. Taking derivatives with respect to the control variable χ^H , and evaluating at my candidates to optimal controls, it follows that:

$$(\phi_1 - \phi_0)\lambda(U) + \phi_1 \left[\underbrace{V_L'(U + \underbrace{\chi^H - p^H}_{=U^*})}_{=-1} - V_H'(U) \right] = 0, \quad (\chi^H)$$

which pins down the value of the multiplier:

$$\lambda(U) = \frac{\phi_1}{\phi_1 - \phi_0} [V_H'(U) + 1] \geq 0,$$

with strict inequality if $U < U^*$. I now perform the same step with the control variable β^H :

$$(\pi_1 - \pi_0)\lambda(U) + \pi_1 \left[V_H'(\underbrace{U + \beta^H - b^H}_{=C_1 U^* + (1 - C_1)U}) - V_H'(U) \right] = 0 \quad (\beta^H)$$

$$\iff \frac{\phi_1}{\phi_1 - \phi_0} [V_H'(U) + 1] = \frac{\pi_1}{\pi_1 - \pi_0} [V_H'(U) - V_H'(C_1 U^* + (1 - C_1)U)]$$

Given my candidates for value functions, the above equation is satisfied for all U if and only if:

$$C_1 = 1 - \left(\frac{\phi_1 \pi_0 - \phi_0 \pi_1}{\pi_1 (\phi_1 - \phi_0)} \right)^{\frac{1}{\alpha-1}} \in (0, 1)$$

Taking derivatives with respect to the remaining control variables and evaluating at the optimal controls,

it follows that:

$$\begin{aligned}
\phi_1 [-V'_L(U + \chi^H - p^H) - 1] &= \phi_1 [-V'_L(U^*) - 1] = 0 & (p^H) \\
\pi_1 [-V'_H(U + \beta^H - b^H) - 1] &= \pi_1 [-V'_H(C_1 U^* + (1 - C_1)U) - 1] \leq 0 & (b^H) \\
-1 - V'_H(U) &\leq 0 & (w^H)
\end{aligned}$$

So indeed my proposed candidate satisfies the KKT first order conditions for a constrained maximum. The Hessian at the candidate point is:

$$\mathcal{H}_{ess_H}(U) = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & \pi_1 V''_H(C_1 U^* + (1 - C_1)U) & 0 & -\pi_1 V''_H(C_1 U^* + (1 - C_1)U) & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & -\pi_1 V''_H(C_1 U^* + (1 - C_1)U) & 0 & \pi_1 V''_H(C_1 U^* + (1 - C_1)U) & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix},$$

with $V''_H(\cdot) \leq 0$, and strict inequality if $U < U^*$. To verify the second-order condition, I consider feasible directions that are nonnegative in the corner variables and orthogonal to the gradient. The only nonzero part of the Hessian lies in the β^H - b^H block. For directions of the form $(0, x, 0, y, 0)$, I compute:

$$d^\top \mathcal{H}_{ess_H} d = \pi_1 V''_H(C_1 U^* + (1 - C_1)U)(x - y)^2 \leq 0$$

This confirms that the restricted second-order condition holds.

I now look at the HJB equation for the state $\theta_t = L$. In this case, the incentive compatibility constraint is slack for my candidate $\chi^L(U) = \frac{r}{\gamma_0}U \geq 0$, so I will ignore the constraint. Taking derivatives with respect to the control variables and evaluating at the optimal controls, it follows that:

$$\begin{aligned} \gamma_0 [V'_H(U + \chi^L - p^L) - V'_L(U)] &= \begin{cases} \gamma_0 [V'_H(U^{\frac{r+\gamma_0}{\gamma_0}}) - V'_L(U)] & \text{if } U < \frac{\gamma_0}{r+\gamma_0} U^* \\ \gamma_0 [V'_H(U^*) - V'_L(U)] & \text{if } \frac{\gamma_0}{r+\gamma_0} U^* \leq U \leq U^* \end{cases} \\ &= \begin{cases} 0 & \text{if } U < \frac{\gamma_0}{r+\gamma_0} U^* \\ 0 & \text{if } \frac{\gamma_0}{r+\gamma_0} U^* \leq U \leq U^* \end{cases} \end{aligned} \quad (\chi^L)$$

$$\pi_0 [V'_L(U + \beta^L - b^L) - V'_L(U)] = \pi_0 [V'_L(U) - V'_L(U)] = 0 \quad (\beta^L)$$

$$\begin{aligned} \gamma_0 [-V'_H(U + \chi^L - p^L) - 1] &= \begin{cases} \gamma_0 [-V'_H(U^{\frac{r+\gamma_0}{\gamma_0}}) - 1] & \text{if } U < \frac{\gamma_0}{r+\gamma_0} U^* \\ \gamma_0 [-V'_H(U^*) - 1] & \text{if } \frac{\gamma_0}{r+\gamma_0} U^* \leq U \leq U^* \end{cases} \\ &= \begin{cases} < 0 & \text{if } U < \frac{\gamma_0}{r+\gamma_0} U^* \\ 0 & \text{if } \frac{\gamma_0}{r+\gamma_0} U^* \leq U \leq U^* \end{cases} \end{aligned} \quad (p^L)$$

$$\pi_0 [-V'_L(U + \beta^L - b^L) - 1] = \pi_0 [-V'_L(U) - 1] \begin{cases} < 0 & \text{if } U < \frac{\gamma_0}{r+\gamma_0} U^* \\ = 0 & \text{if } \frac{\gamma_0}{r+\gamma_0} U^* \leq U \leq U^* \end{cases} \quad (b^L)$$

$$-1 - V'_L(U) \begin{cases} < 0 & \text{if } U < \frac{\gamma_0}{r+\gamma_0} U^* \\ = 0 & \text{if } \frac{\gamma_0}{r+\gamma_0} U^* \leq U \leq U^* \end{cases} \quad (w^L)$$

So indeed my proposed candidate satisfies the KKT first order conditions for a constrained maximum. The Hessian at the candidate point is:

$$\mathcal{H}_{ess_L}(U) = \begin{pmatrix} \gamma_0 V''_H(U + \chi^L(U)) & 0 & -\gamma_0 V''_H(U + \chi^L(U)) & 0 & 0 \\ 0 & \pi_0 V''_L(U) & 0 & -\pi_0 V''_L(U) & 0 \\ -\gamma_0 V''_H(U + \chi^L(U)) & 0 & \gamma_0 V''_H(U + \chi^L(U)) & 0 & 0 \\ 0 & -\pi_0 V''_L(U) & 0 & \pi_0 V''_L(U) & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix},$$

with $V''_L(\cdot) \leq 0$, and strict inequality if $U < U^*$. To verify the second-order condition, I consider feasible directions that are nonnegative in the corner variables and orthogonal to the gradient. For directions of the form $d = (x, y, z, s, 0)$, I compute:

$$d^\top \mathcal{H}_{ess_L} d = \pi_0 V''_L(U)(y - s)^2 + \gamma_0 V''_H(U + \chi^L(U))(x - z)^2 \leq 0$$

This confirms that the restricted second-order condition holds.

It remains to show that the proposed value functions indeed solve the HJB equations. Both equations are satisfied for all values of U if and only if α solves:

$$\left(r + \pi_1 + \phi_1 - \pi_1 \left(\frac{\phi_1 \pi_0 - \pi_1 \phi_0}{\pi_1(\phi_1 - \phi_0)} \right)^{\frac{\alpha}{\alpha-1}} \right) - \alpha \left(r + \frac{\phi_1 \pi_0 - \pi_1 \phi_0}{\phi_1 - \phi_0} \left(1 - \left(\frac{\phi_1 \pi_0 - \pi_1 \phi_0}{\pi_1(\phi_1 - \phi_0)} \right)^{\frac{1}{\alpha-1}} \right) \right) = 0$$

The left-hand side of the above equation is strictly decreasing in α in the exhaustion-dominant case ($\phi_1 \pi_0 > \pi_1 \phi_0$), with a strictly positive limit when $\alpha \rightarrow 1^+$, and with a strictly negative limit for α large

enough, which means that there is a unique solution with $\alpha > 1$.

Finally, I verify the role of my assumptions in guaranteeing that my contract is well defined. First, limited liability imposes that no transfer is negative, so it must be verified that $p^H(0) \geq 0$:

$$\begin{aligned} p^H(0) &= \chi^H(0) - U^* \\ &= \frac{c}{\phi_1 - \phi_0} - U^* - \beta^H(0) \frac{\pi_1 - \pi_0}{\phi_1 - \phi_0} \\ &= \frac{c}{\phi_1 - \phi_0} - U^* - C_1 U^* \frac{\pi_1 - \pi_0}{\phi_1 - \phi_0} \\ &= \frac{c}{\phi_1 - \phi_0} \left(1 - \frac{\phi_0}{r} - \frac{\pi_1 - \pi_0}{\phi_1 - \phi_0} \frac{\phi_0}{r} C_1 \right) \end{aligned}$$

Thus, nonnegativity is satisfied if the discount factor is large enough:

$$r \geq \phi_0 \left(1 + \frac{\pi_1 - \pi_0}{\phi_1 - \phi_0} C_1 \right)$$

In particular, since $C_1 < 1$, it is enough to assume that $r \geq \phi_0 \left(1 + \frac{\pi_1 - \pi_0}{\phi_1 - \phi_0} \right)$. Next, to ensure that the optimal initialising continuation value is positive, I verify that the slope of the principal value function is strictly positive at $U = 0$:

$$\begin{aligned} V'_H(0) = V'_L(0) &= -1 + \frac{\alpha}{U^*} \frac{[(\pi_1 - \pi_0)v - c](r + \gamma_0)}{(r + \gamma_0 + \phi_1)r} > 0 \\ \iff \frac{(\pi_1 - \pi_0)v - c}{r + \gamma_0 + \phi_1} \frac{r + \gamma_0}{r} &> \frac{c\phi_0}{r(\phi_1 - \phi_0)} \frac{1}{\alpha} \\ \iff \frac{\pi_1 - \pi_0}{c} v > 1 + \frac{\phi_0}{\phi_1 - \phi_0} \frac{r + \gamma_0 + \phi_1}{r + \gamma_0} \frac{1}{\alpha} \end{aligned}$$

In particular, since $\alpha > 1$, it is enough to assume $\frac{\pi_1 - \pi_0}{c} v \geq 1 + \frac{\phi_0}{\phi_1 - \phi_0} \frac{r + \gamma_0 + \phi_1}{r + \gamma_0}$.

□

LEMMA 1.2 (Verification). *Let $V : \mathbb{R}_+ \times \{H, L\} \rightarrow \mathbb{R}$ be a continuous function that is piecewise C^1 in its first argument and concave. Assume V satisfies the system of HJB equations (1.4.5)-(1.4.6) on the domain $[0, U^*]$, and is extended linearly with slope -1 for $U > U^*$. Then any contract Y^* that attains the HJB system is optimal among admissible $\mathcal{F}_t$ -adapted contracts implementing a^* .*

Proof of Lemma 1.2. Consider any admissible $\mathcal{F}_t$ -adapted contract Y that implements the effort process a^* . By the Lebesgue Decomposition Theorem, we decompose Y into an absolutely continuous part (w_t), a pure jump part tied to successes (b_t), pure jumps tied to transitions (p_t^H, p_t^L), other pure jumps (ΔY_t), and a singular continuous part (Y_t^{sc}):

$$dY_t = w_t dt + b_t dX_t + p_t^H dJ_t^{H \rightarrow L} + p_t^L dJ_t^{L \rightarrow H} + \Delta Y_t + dY_t^{sc}.$$

Let $\mathcal{V}_t = e^{-rt} V(U_t, \theta_t) + \int_0^t e^{-rs} (v dX_s - dY_s)$ be the principal's expected total payoff process. By the

Meyer-Itô formula for semimartingales, we analyze the dynamic evolution of $\mathcal{V}_t$. Consider the case where the current state is high, $\theta_{t-} = H$. We obtain:

$$\begin{aligned}
e^{rt} d\mathcal{V}_t = & \left(\pi_1(v - b_t) - w_t - \phi_1 p_t^H + \frac{\partial V}{\partial U}(U_{t-}, H)[rU_{t-} + c - \pi_1 \beta_t - \phi_1 \chi_t^H - w_t] - rV(U_{t-}, H) \right. \\
& + \pi_1[V(U_{t-} + \beta_t - b_t, H) - V(U_{t-}, H)] + \phi_1[V(U_{t-} + \chi_t^H - p_t^H, L) - V(U_{t-}, H)] \left. \right) dt \\
& + [v - b_t + V(U_{t-} + \beta_t - b_t, H) - V(U_{t-}, H)] dM_t \\
& + [-p_t^H + V(U_{t-} + \chi_t^H - p_t^H, L) - V(U_{t-}, H)] dN_t^{H \rightarrow L} \\
& - \left[1 + \frac{\partial V}{\partial U}(U_{t-}, H) \right] dY_t^{sc} \\
& + [V(U_{t-} - \Delta Y_t, H) - V(U_{t-}, H) - \Delta Y_t],
\end{aligned}$$

where $dM_t = dX_t - \pi_1 dt$ and $dN_t^{H \rightarrow L} = dJ_t^{H \rightarrow L} - \phi_1 dt$. The first term (multiplied by dt) is less or equal to zero since V solves the HJB system. The second and third terms are zero in expectation, since both M_t and $N_t^{H \rightarrow L}$ are martingales, the contract is admissible,¹⁰ and the shape of V solving the HJB system: bounded from above and linearly decreasing in U for $U > U^*$.¹¹ The fourth term is no greater than zero since $\frac{\partial}{\partial U} V(U, H) \geq -1$ for all U . The fifth term is not greater than zero since V is concave in U and again $\frac{\partial}{\partial U} V(U, H) \geq -1$. All the inequalities become equalities at the contract Y^* that solves the HJB system. An analogous derivation holds for the case $\theta_{t-} = L$, where the HJB equation for the low state ensures the drift is non-positive.

Thus, $\mathcal{V}_t$ is a supermartingale. For any $T > 0$:

$$V(U_0, \theta_0) = \mathcal{V}_0 \geq \mathbb{E}[\mathcal{V}_T] = \mathbb{E} \left[\int_0^T e^{-rs} (vdX_s - dY_s) \right] + \mathbb{E} [e^{-rT} V(U_T, \theta_T)].$$

Taking the limit as $T \rightarrow \infty$:

$$V(U_0, \theta_0) \geq \mathbb{E} \left[\int_0^\infty e^{-rs} (vdX_s - dY_s) \right] + \liminf_{T \rightarrow \infty} \mathbb{E} [e^{-rT} V(U_T, \theta_T)].$$

¹⁰Indeed, the admissibility condition states that $\mathbb{E}[\int_0^\infty e^{-rt} dY_t] < \infty$, which first implies that b_t is integrable. Moreover, it also implies that the transversality condition $\liminf_{t \rightarrow \infty} \mathbb{E}[e^{-rT} U_T] = 0$ is satisfied. To see this, notice that:

$$\mathbb{E}[e^{-rT} U_T] = \mathbb{E} \left[\int_T^\infty e^{-rt} (dY_t - ca_t dt) \right].$$

By limited liability, we can find the following bounds:

$$0 \leq \mathbb{E}[e^{-rT} U_T] \leq \mathbb{E} \left[\int_T^\infty e^{-rt} dY_t \right].$$

The desired result follows by taking limit as $T \rightarrow \infty$, since both upper and lower bounds must vanish.

¹¹For $U > U^*$, the value function satisfies $V(U, H) = V(U^*, H) + U^* - U$, which means that $V(U, H) \geq K - U$ for some constant K . Thus, $e^{-rT} V(U_T, H) \geq e^{-rT} (K - U_T)$. Taking expectation and limits one obtains:

$$\lim_{T \rightarrow \infty} \mathbb{E}[e^{-rT} V(U_T, H)] \geq \lim_{T \rightarrow \infty} (e^{-rT} K - \mathbb{E}[e^{-rT} U_T]) = 0.$$

The admissibility condition (finite expected cost) implies the transversality condition on the promised utility: $\lim_{T \rightarrow \infty} \mathbb{E}[e^{-rT}U_T] = 0$. Since $V(U, \theta)$ behaves linearly as $-U$ for large U , this implies $\lim_{T \rightarrow \infty} \mathbb{E}[e^{-rT}V(U_T, \theta_T)] = 0$. Therefore, $V(U_0, \theta_0)$ is an upper bound on the principal's payoff. By a similar argument, one obtains that the contract Y^* that solves the system of HJB equations achieves the supremum:

$$V(U_0, \theta_0) = \mathcal{V}_0 = \mathbb{E}[\mathcal{V}_T] = \lim_{T \rightarrow \infty} \mathbb{E}[\mathcal{V}_T] = \mathbb{E} \left[\int_0^\infty e^{-rs} (v dX_s - dY_s^*) \right].$$

□

Proof of Corollary 1.1. The claim follows from observing that the value functions $V_H(U)$ and $V_L(U)$ are concave in U , and from taking first-order conditions with respect to U . The assumptions in Proposition 1.2 ensure that the optimal value $\hat{U}_0^H$ is interior, and hence so is $\hat{U}_0^L = \frac{\gamma_0}{r+\gamma_0} \hat{U}_0^H$.

The stopping time $T^\varnothing$ is then calculated as the time it takes for the continuation value to move from $\hat{U}_0^H$ at time 0 to $U_T = 0$ in the absence of any shock:

$$dU_t = \kappa(U_t - U^*)dt, \quad U_{T^\varnothing} = 0, \quad U_0 = \hat{U}_0^H \quad \Rightarrow \quad \hat{U}_0^H = U^*(1 - e^{-\kappa T^\varnothing}).$$

Finally, I show how to obtain the extensions ΔT after each success during the probatory phase. At time t , if the continuation value is U_t , the remaining until termination if no other shock arrives is given by:

$$U_t = U^*(1 - e^{\kappa(t-T^\varnothing)}). \quad (1.A.10)$$

If instead a success is realized at time t , the continuation value jumps from U_t to

$$U_t + \beta^H(U_t) - b^H(U_t) = C_1 U^* + (1 - C_1)U_t > U_t$$

And the new remaining time is now given by:

$$C_1 U^* + (1 - C_1)U_t = U^*(1 - e^{\kappa(t-T-\Delta T)}). \quad (1.A.11)$$

Combining (1.A.10) and (1.A.11), it follows that:

$$\Delta T = -\frac{\log(1 - C_1)}{\kappa} > 0$$

which does not depend on calendar time t nor in the continuation value U_t .

□

Proof of Proposition 1.3. In this proof I show that closed form candidates for value functions, together with the candidates to optimal control that are described, indeed solve the system of HJB equations (1.4.5)-(1.4.6). The usual verification step can be found and proved in Lemma 1.2.

Let $U^* = \frac{r+\gamma_0}{r} \frac{\pi_0 c}{(r+\gamma_0)(\pi_1-\pi_0)+\pi_1\phi_0-\pi_0\phi_1} > 0$, and consider the following candidates for value functions defined on the domain $[0, U^*]$:

$$\begin{aligned} V_H(U) &= V_H^{fb}(U) - \frac{[(\pi_1 - \pi_0)v - c](r + \gamma_0)}{(r + \gamma_0 + \phi_1)r} \left[1 - \frac{U}{U^*}\right]^\alpha \\ V_L(U) &= V_L^{fb}(U) - \frac{[(\pi_1 - \pi_0)v - c]\gamma_0}{(r + \gamma_0 + \phi_1)r} \left[1 - \frac{U}{U^*} \frac{r + \gamma_0}{\gamma_0}\right]^\alpha \mathbb{1} \left\{0 \leq U < \frac{\gamma_0}{r + \gamma_0} U^*\right\} \end{aligned}$$

As I will verify later, these value functions are concave and single-peaked at interior values for U . I now check the optimality of the controls, for which I will need the derivatives of the value functions:

$$\begin{aligned} V_H'(U) &= -1 + \frac{[(\pi_1 - \pi_0)v - c](r + \gamma_0)}{(r + \gamma_0 + \phi_1)r} \frac{\alpha}{U^*} \left[1 - \frac{U}{U^*}\right]^{\alpha-1} \geq -1 \\ V_L'(U) &= \begin{cases} -1 + \frac{(\pi_1 - \pi_0)v - c}{r + \gamma_0 + \phi_1} \frac{r + \gamma_0}{r} \frac{\alpha}{U^*} \left[1 - \frac{U}{U^*} \frac{r + \gamma_0}{\gamma_0}\right]^{\alpha-1} > -1 & \text{if } 0 \leq U < \frac{\gamma_0}{r + \gamma_0} U^* \\ -1 & \text{if } \frac{\gamma_0}{r + \gamma_0} U^* \leq U \leq U^* \end{cases} \end{aligned}$$

Notice in particular that $V_H'(U^*) = V_L'(U^*) = -1$. Let me first look at the HJB equation for the state $\theta_t = H$. Denote by $\lambda(U) \geq 0$ the multiplier associated to the incentive compatibility constraint for each value of the state variable U . The only nonzero controls are β^H , b^H and χ^H , where $\beta^H(U) - b^H(U) = U^* - U$, $\chi^H(U) = \frac{\gamma_0}{r + \gamma_0} (C_2 U^* + (1 - C_2)U) - U$, and C_2 is a constant to be specified later. Taking derivatives with respect to the control variable β^H , and evaluating at my candidates to optimal controls, it follows that:

$$(\pi_1 - \pi_0)\lambda(U) + \pi_1 \underbrace{\left[V_H'(U + \underbrace{\beta^H - b^H}_{=U^*}) - V_H'(U) \right]}_{=-1} = 0, \quad (\beta^H)$$

which pins down the value of the multiplier:

$$\lambda(U) = \frac{\pi_1}{\pi_1 - \pi_0} [V_H'(U) + 1] \geq 0,$$

with strict inequality if $U < U^*$. I now perform the same step with the control variable χ^H :

$$(\phi_1 - \phi_0)\lambda(U) + \phi_1 \left[V_L' \left(U + \underbrace{\chi^H - p^H}_{=\frac{\gamma_0}{r + \gamma_0} (C_2 U^* + (1 - C_2)U)} \right) - V_H'(U) \right] = 0 \quad (\chi^H)$$

$$\iff \frac{\pi_1}{\pi_1 - \pi_0} [V_H'(U) + 1] = \frac{\phi_1}{\phi_1 - \phi_0} \left[V_H'(U) - V_L' \left(\frac{\gamma_0}{r + \gamma_0} (C_2 U^* + (1 - C_2)U) \right) \right] \quad (1.A.12)$$

Given my candidates for value functions, the above equation is satisfied for all U if and only if:

$$C_2 = 1 - \left(\frac{\pi_1 \phi_0 - \pi_0 \phi_1}{\phi_1 (\pi_1 - \pi_0)} \right)^{\frac{1}{\alpha-1}} \in (0, 1) \quad (1.A.13)$$

Taking derivatives with respect to the remaining control variables and evaluating at the optimal controls, it follows that:

$$\pi_1 [-V'_H(U + \beta^H - b^H) - 1] = \pi_1 [-V'_H(U^*) - 1] = 0 \quad (b^H)$$

$$\phi_1 [-V'_L(U + \chi^H - p^H) - 1] = \phi_1 \left[-V'_L \left(\frac{\gamma_0}{r + \gamma_0} (C_2 U^* + (1 - C_2)U) \right) - 1 \right] \leq 0 \quad (p^H)$$

$$-1 - V'_H(U) \leq 0 \quad (w^H)$$

So indeed my proposed candidate satisfies the KKT first order conditions for a constrained maximum. The Hessian at the candidate point is:

$$\mathcal{H}ess_H(U) = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & \phi_1 V''_L(\frac{\gamma_0}{r+\gamma_0}(C_2 U^* + (1 - C_2)U)) & 0 & -\phi_1 V''_L(\frac{\gamma_0}{r+\gamma_0}(C_2 U^* + (1 - C_2)U)) & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & -\phi_1 V''_L(\frac{\gamma_0}{r+\gamma_0}(C_2 U^* + (1 - C_2)U)) & 0 & \phi_1 V''_L(\frac{\gamma_0}{r+\gamma_0}(C_2 U^* + (1 - C_2)U)) & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix},$$

with $V''_L(\cdot) \leq 0$, and strict inequality if $U < U^*$. To verify the second-order condition, I consider feasible directions that are nonnegative in the corner variables and orthogonal to the gradient. The only nonzero part of the Hessian lies in the $\chi^H - p^H$ block. For directions of the form $d = (0, x, 0, y, 0)$, I compute:

$$d^\top \mathcal{H}ess_H d = \phi_1 V''_L(\frac{\gamma_0}{r+\gamma_0}(C_2 U^* + (1 - C_2)U))(x - y)^2 \leq 0$$

This confirms that the restricted second-order condition holds.

I now look at the HJB equation for the state $\theta = L$. In this case, the incentive compatibility constraint is slack for my candidate $\chi^L(U) = \frac{r}{\gamma_0} U \geq 0$, so I will ignore the constraint. Taking derivatives with respect to the control variables and evaluating at the optimal controls, it follows that:

$$\begin{aligned}
\pi_0 [V'_L(U + \beta^L - b^L) - V'_L(U)] &= \pi_0 [V'_L(U) - V'_L(U)] = 0 & (\beta^L) \\
\gamma_0 [V'_H(U + \chi^L - p^L) - V'_L(U)] &= \begin{cases} \gamma_0 [V'_H(U^{\frac{r+\gamma_0}{\gamma_0}}) - V'_L(U)] & \text{if } U < \frac{\gamma_0}{r+\gamma_0} U^* \\ \gamma_0 [V'_H(U^*) - V'_L(U)] & \text{if } \frac{\gamma_0}{r+\gamma_0} U^* \leq U \leq U^* \end{cases} & (\chi^L) \\
&= \begin{cases} 0 & \text{if } U < \frac{\gamma_0}{r+\gamma_0} U^* \\ 0 & \text{if } \frac{\gamma_0}{r+\gamma_0} U^* \leq U \leq U^* \end{cases} \\
\pi_0 [-V'_L(U + \beta^L - b^L) - 1] &= \pi_0 [-V'_L(U) - 1] \begin{cases} < 0 & \text{if } U < \frac{\gamma_0}{r+\gamma_0} U^* \\ = 0 & \text{if } \frac{\gamma_0}{r+\gamma_0} U^* \leq U \leq U^* \end{cases} & (b^L) \\
\gamma_0 [-V'_H(U + \chi^L - p^L) - 1] &= \begin{cases} \gamma_0 [-V'_H(U^{\frac{r+\gamma_0}{\gamma_0}}) - 1] & \text{if } U < \frac{\gamma_0}{r+\gamma_0} U^* \\ \gamma_0 [-V'_H(U^*) - 1] & \text{if } \frac{\gamma_0}{r+\gamma_0} U^* \leq U \leq U^* \end{cases} & (p^L) \\
&= \begin{cases} < 0 & \text{if } U < \frac{\gamma_0}{r+\gamma_0} U^* \\ 0 & \text{if } \frac{\gamma_0}{r+\gamma_0} U^* \leq U \leq U^* \end{cases} \\
-1 - V'_L(U) &\begin{cases} < 0 & \text{if } U < \frac{\gamma_0}{r+\gamma_0} U^* \\ = 0 & \text{if } \frac{\gamma_0}{r+\gamma_0} U^* \leq U \leq U^* \end{cases} & (w^L)
\end{aligned}$$

So indeed my proposed candidate satisfies the KKT first order conditions for a constrained maximum. The Hessian at the candidate point is:

$$\mathcal{Hess}_L(U) = \begin{pmatrix} \gamma_0 V''_H(U + \chi^L(U)) & 0 & -\gamma_0 V''_H(U + \chi^L(U)) & 0 & 0 \\ 0 & \pi_0 V''_L(U) & 0 & -\pi_0 V''_L(U) & 0 \\ -\gamma_0 V''_H(U + \chi^L(U)) & 0 & \gamma_0 V''_H(U + \chi^L(U)) & 0 & 0 \\ 0 & -\pi_0 V''_L(U) & 0 & \pi_0 V''_L(U) & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix},$$

with $V''_L(\cdot) \leq 0$, and strict inequality if $U < U^*$. To verify the second-order condition, I consider feasible directions that are nonnegative in the corner variables and orthogonal to the gradient. For directions of the form $d = (x, y, z, s, 0)$, I compute:

$$d^\top \mathcal{Hess}_L d = \pi_0 V''_L(U)(y - s)^2 + \gamma_0 V''_H(U + \chi^L(U))(x - z)^2 \leq 0$$

This confirms that the restricted second-order condition holds.

It remains to show that the proposed value functions indeed solve the HJB equations. Both equations are satisfied for all values of U if and only if α and C_2 solve:

$$C_2 = 1 + \frac{r + \gamma_0}{\gamma_0} \left(1 - \frac{(\pi_1 - \pi_0)(\pi_1 + \phi_1 + r(1 - \alpha))}{\alpha(\pi_1 \phi_0 - \pi_0 \phi_1)} \right) \quad (1.A.14)$$

Combining (1.A.13) and (1.A.14), a single equation for α is obtained:

$$\frac{r + \gamma_0}{\gamma_0} \left(1 - \frac{(\pi_1 - \pi_0)(\pi_1 + \phi_1 + r(1 - \alpha))}{\alpha(\pi_1\phi_0 - \pi_0\phi_1)} \right) + \left(\frac{\pi_1\phi_0 - \pi_0\phi_1}{\phi_1(\pi_1 - \pi_0)} \right)^{\frac{1}{\alpha-1}} = 0$$

The left-hand side of the above equation is strictly increasing in α in the success-dominant case ($\pi_1\phi_0 > \phi_1\pi_0$), with a strictly negative limit when $\alpha \rightarrow 1^+$, and a strictly positive limit for when $\alpha \rightarrow +\infty$, which means that there is a unique solution with $\alpha > 1$.

Finally, I verify the role of my assumptions in guaranteeing that my contract is well defined. First, limited liability imposes that no transfer is negative, so I need to verify that $b^H(0) \geq 0$:

$$\begin{aligned} b^H(0) &= \beta^H(0) - U^* \\ &= \frac{c}{\pi_1 - \pi_0} - U^* - \chi^H(0) \frac{\phi_1 - \phi_0}{\pi_1 - \pi_0} \\ &= \frac{c}{\pi_1 - \pi_0} - U^* \left(1 + C_2 \frac{\phi_1 - \phi_0}{\pi_1 - \pi_0} \frac{\gamma_0}{r + \gamma_0} \right) \\ &= \frac{c}{\pi_1 - \pi_0} - \frac{r + \gamma_0}{r} \frac{\pi_0 c}{(r + \gamma_0)(\pi_1 - \pi_0) + \pi_1\phi_0 - \pi_0\phi_1} \left(1 + C_2 \frac{\phi_1 - \phi_0}{\pi_1 - \pi_0} \frac{\gamma_0}{r + \gamma_0} \right) \end{aligned}$$

Thus, nonnegativity is satisfied if the discount factor is large enough:

$$r \geq \pi_0 \frac{(r + \gamma_0)(\pi_1 - \pi_0)}{(r + \gamma_0)(\pi_1 - \pi_0) + \pi_1\phi_0 - \pi_0\phi_1} \left(1 + C_2 \frac{\phi_1 - \phi_0}{\pi_1 - \pi_0} \frac{\gamma_0}{r + \gamma_0} \right)$$

In particular, since $C_2 < 1$, it is enough to assume that $r \geq \pi_0 \left(1 + \frac{\phi_1 - \phi_0}{\pi_1 - \pi_0} \right)$. Next, to ensure that the optimal initialising continuation value is positive, I verify that the slope of the principal value function is strictly positive at $U = 0$:

$$\begin{aligned} V'_H(0) &= V'_L(0) = -1 + \frac{\alpha}{U^*} \frac{[(\pi_1 - \pi_0)v - c](r + \gamma_0)}{(r + \gamma_0 + \phi_1)r} > 0 \\ \iff \frac{(\pi_1 - \pi_0)v - c}{r + \gamma_0 + \phi_1} \frac{r + \gamma_0}{r} &> \frac{r + \gamma_0}{r} \frac{\pi_0 c}{(r + \gamma_0)(\pi_1 - \pi_0) + \pi_1\phi_0 - \pi_0\phi_1} \frac{1}{\alpha} \\ \iff \frac{\pi_1 - \pi_0}{c} v &> 1 + \frac{(r + \gamma_0 + \phi_1)\pi_0}{(r + \gamma_0)(\pi_1 - \pi_0) + \pi_1\phi_0 - \pi_0\phi_1} \frac{1}{\alpha} \end{aligned}$$

In particular, since $\alpha > 1$, it is enough to assume $\frac{\pi_1 - \pi_0}{c} v \geq 1 + \frac{(r + \gamma_0 + \phi_1)\pi_0}{(r + \gamma_0)(\pi_1 - \pi_0) + \pi_1\phi_0 - \pi_0\phi_1}$.

□

Proof of Corollary 1.2. The claim follows from observing that the value functions $V_H(U)$ and $V_L(U)$ are concave in U , and from taking first-order conditions with respect to U . The assumptions in Proposition 1.3 ensure that the optimal value $\hat{U}_0^H$ is interior, and hence so is $\hat{U}_0^L = \frac{\gamma_0}{r + \gamma_0} \hat{U}_0^H$.

The stopping time $T^\emptyset$ is then calculated as the time it takes for the continuation value to move from

$\hat{U}_0^H$ at time 0 to $U_T = 0$ in the absence of any shock:

$$dU_t = \tilde{\kappa}(U_t - U^*)dt, \quad U_{T^0} = 0, \quad U_0 = \hat{U}_0^H \quad \Rightarrow \quad \hat{U}_0^H = U^*(1 - e^{-\tilde{\kappa}T^0}),$$

where

$$\tilde{\kappa} = r + \frac{\pi_1\phi_0 - \pi_0\phi_1}{\pi_1 - \pi_0} \left(1 - \frac{\gamma_0}{r + \gamma_0} \left(\frac{\pi_1\phi_0 - \pi_0\phi_1}{\phi_1(\pi_1 - \pi_0)} \right)^{\frac{1}{\alpha-1}} \right) > 0.$$

Finally, I show how to obtain the extensions ΔT after each transition during the probatory phase. At time t , if the continuation value is t , the remaining until termination if no other shock arrives is given by:

$$U_t = U^*(1 - e^{\tilde{\kappa}(t-T)}). \quad (1.A.15)$$

Suppose instead a transition from H to L is realized at time t , the continuation value jumps from U_t to

$$U_t + \chi^H(U_t) - p^H(U_t) = \frac{\gamma_0}{r + \gamma_0} (C_2 U^* + (1 - C_2) U_t),$$

and the probatory phase is stopped until there is a new transition from L to H . At such moment, the continuation value jumps to

$$\frac{\gamma_0}{r + \gamma_0} (C_2 U^* + (1 - C_2) U_t) + \chi^L \left(\frac{\gamma_0}{r + \gamma_0} (C_2 U^* + (1 - C_2) U_t) \right) = C_2 U^* + (1 - C_2) U_t > U_t.$$

The new remaining time is now given by:

$$C_2 U^* + (1 - C_2) U_t = U^*(1 - e^{\tilde{\kappa}(t-T-\Delta T)}). \quad (1.A.16)$$

Combining (1.A.15) and (1.A.16), it follows that:

$$\Delta T = -\frac{\log(1 - C_2)}{\tilde{\kappa}} > 0$$

which does not depend on calendar time t nor in the continuation value U_t .

□

1.B. FIRST-BEST EFFORT PROFILE

I derive here the first-best effort profile as shown in Section 1.3.1. Recall that the problem is:

$$\begin{aligned} rW(H) &= \sup_{a_H \in [0,1]} \left\{ v(\pi_0 + (\pi_1 - \pi_0)a_H) - ca_H + [\phi_0 + (\phi_1 - \phi_0)a_H] (W(L) - W(H)) \right\}, \\ rW(L) &= \sup_{a_L \in [0,1]} \left\{ v\pi_0 - ca_L + [\gamma_0 + (\gamma_1 - \gamma_0)a_L] (W(H) - W(L)) \right\}. \end{aligned} \quad (\mathcal{P}^{fb})$$

PROPOSITION 1.1. *Assume $(\pi_1 - \pi_0)v > c$. Then the optimal effort process a^* is bang-bang: full effort in the high state and no effort in the low state,*

$$a_t^* = \begin{cases} 1 & \text{if } \theta_t = H, \\ 0 & \text{if } \theta_t = L. \end{cases}$$

The corresponding value functions are

$$\begin{aligned} W(H) &= \frac{r + \gamma_0}{r + \gamma_0 + \phi_1} \frac{\pi_1 v - c}{r} + \frac{\phi_1}{r + \gamma_0 + \phi_1} \frac{\pi_0 v}{r} > 0, \\ W(L) &= \frac{\gamma_0}{r + \gamma_0 + \phi_1} \frac{\pi_1 v - c}{r} + \frac{r + \phi_1}{r + \gamma_0 + \phi_1} \frac{\pi_0 v}{r} \in (0, W(H)). \end{aligned}$$

Proof of Proposition 1.1. It is straightforward to verify that the HJB equations are satisfied when plugging $W(H)$, $W(L)$ and a^* . I just need to show that a^* indeed prescribes an optimal control. The objective functions are linear in effort, and thus $a_t = 1$ is optimal when $\theta_t = H$ if and only if:

$$\begin{aligned} &(\pi_1 - \pi_0)v - c + (\phi_1 - \phi_0)(W(L) - W(H)) \geq 0 \\ \iff &[(\pi_1 - \pi_0)v - c] \left(1 - \frac{\phi_1 - \phi_0}{r + \gamma_0 + \phi_1} \right) \geq 0 \\ \iff &[(\pi_1 - \pi_0)v - c] \frac{r + \gamma_0 + \phi_0}{r + \gamma_0 + \phi_1} \geq 0 \end{aligned}$$

And thus it is optimal if and only if $(\pi_1 - \pi_0)v \geq c$. Finally, $a_t = 0$ is optimal when $\theta_t = L$ if and only if:

$$\begin{aligned} &0 - c + (\gamma_1 - \gamma_0)(W(H) - W(L)) \leq 0 \\ \iff &-c + [(\pi_1 - \pi_0)v - c] \frac{\gamma_1 - \gamma_0}{r + \gamma_0 + \phi_1} \leq 0 \\ \iff &(\gamma_1 - \gamma_0)(\pi_1 - \pi_0)v - (r + \gamma_1 + \phi_1)c \leq 0 \end{aligned}$$

Since effort slows down the recovery rate, $\gamma_1 < \gamma_0$, and thus the marginal benefit of effort is always negative, so I have a corner solution with $a_t = 0$. $\square$

Chapter 2

On the (Relative) Merits of Money Burning for Optimal Contracting

Javier González-Morín and David Martimort

ABSTRACT

We consider the following screening model of procurement. An agent (the seller) has private information on his cost parameter. A principal (the buyer) learns an *ex post* signal on this parameter. The signal is private information to the principal and proper incentives to reveal this signal must be designed. In related contexts, money burning, i.e., the *ex post* destruction of some of the gains from trade, has shown to be useful to provide such incentives. We demonstrate that money burning allows the principal to implement the first-best output with zero information rent for the agent; although it is never optimal to do so since output distortions are less costly. More generally, money burning is rarely optimal, and only used as a tool of last resort if output distortions are no longer feasible. In particular, when output must be chosen before the non-verifiable signal realizes, money burning becomes more attractive.

2.1. INTRODUCTION

We consider the following screening model of procurement. The principal (the buyer, *she*) contracts for the delivery of good or service. The agent in charge (the seller, *he*) has private information on his cost parameter. In this adverse selection context, the principal should offer a menu of contracts to maximise her surplus and let the agent select the right option within the menu according to his type. However, achieving successful self-selection requires that the principal leaves some information rent to the agent. To reduce this rent that she views as costly, the principal must distort quantities downwards. The so called *rent/efficiency* trade-off as exemplified in [Baron and Myerson \(1982\)](#) is at play.

Of course, a more direct access to information would facilitate the principal's task in extracting more information rent from her agent and could thus be useful to better regulate the agent. As a result, lesser output distortions might be implemented; and the aforementioned rent/efficiency trade-off would be tilted towards efficiency. Benchmarking of the seller's performances with those of related sellers on the market place, or even more extreme forms of yardstick information where individual compensations are linked to peers' performances are all information management techniques that are often used in practice to this end and that fall under this umbrella.

On the theory side, the benefit of an additional ex post public signal, which is correlated with the agent's type and can be used by a principal to fully extract the agent's information rent is well documented, at least since the work of [Riordan and Sappington \(1988\)](#).¹ When the principal has access to an ex post public signal on the agent's cost parameter, the optimal contract can be designed so as to fully extract the agent's information rent and implement the first-best.² Contingent payments that depend on how the agent's report is confirmed or not by the principal's signal suffice to do the trick.

Much less is known when the principal ex post signal is non-verifiable, or private information. Indeed, there are many instances where the principal may obtain information that cannot be verified by a Court of Law. For instance, there could be a leakage of information, such as insider knowledge being revealed by an employee of the seller's unit, that makes the principal learn about the agent. In such scenario, the principal cannot write down an explicit contract contingent on the realized signal. The principal may behave opportunistically and report whatever signal would minimize her payment to the agent. The principal thus faces the challenge of designing ex ante a contract that prevents her own opportunism ex post. In addition to the more standard incentive compatibility constraints on the agent, incentive constraints on the principal's side must also be imposed to induce her to truthful report whatever signal she has learned. In our context, simple sell-out contracts that leave the principal's payoff unchanged across signal realizations and make the agent residual claimant for the gains from trade solve the credibility problem.

Two scenarios could then envisioned. In the first scenario, there is no waste of money and whatever is paid by the principal for the agent's services ends up being pocketed by the agent. This requirement is akin to assuming that trade is balanced, or that any putative waste of resource would be renegotiated away on efficiency grounds. Solving simultaneously with the sole use of a single payment scheme the principal's and the agent's incentive problems may be difficult in those circumstances.³ Breaking the budget could thus be attractive; the second scenario that is the focus of this paper. Money burning, since it creates a wedge between what the principal pays for the agent's services and what the latter receives, precisely performs this task. In practice, money burning can be viewed as a metaphor for the costs of conflicts arising from contradictory information, costly and useless actions, or artificial costs implemented to prevent voluntary resignations or contractual terminations.

¹More recent works on this front are [Kessler et al. \(2005\)](#), [Krähmer \(2020\)](#) and [Kattwinkel and Knoepfle \(2023\)](#).

²The logic is also familiar from the work of [Crémer and McLean \(1985\)](#), [Crémer and McLean \(1988\)](#), [McAfee and Reny \(1992\)](#), [Lopomo et al. \(2022\)](#) and [Rahman \(2012a\)](#) in multi-agent environments.

³This tension is reminiscent of the moral hazard in teams problem à la [Holmström \(1982\)](#) or to the difficulty of satisfying requirement of budget balance in mechanism design environment as in [Laffont and Maskin \(1980\)](#).

A priori, one could think that this extra degree of freedom in contracting might allow to implement cheaper incentives and thereby improve the terms of the rent/efficiency trade-off. Money burning gives an extra possibility for keeping the principal's payoff constant across realizations of her signal, satisfying thereby her incentive constraints, without changing the agent's payments and thus perturbing his own incentives. As a result, money burning could allow to reduce output distortions and lower informational rents. To illustrate, we demonstrate that money burning certainly allows to implement first-best output while it permits full rent extraction.

Although attractive and intuitive, this line of reasoning stressing the beneficial role of money burning is somewhat incomplete. Money burning may just not be optimal, even when allowed. The waste of surplus might be too large. The intuition is the following. Remember that the principal's incentives hold provided that her net payoff is constant over possible realizations of her signal. Without money burning, this constancy can be achieved by varying at the same time payments and outputs. A higher payment to the agent in case the observed signal confirms the agent's report should then be compensated by a greater output so as to leave the principal's net surplus unchanged. Under broad circumstances, the value of contracting is positive at a given pair of signal-cost realizations and it is always better to use such variations in output and payment rather than just burn money. Under weak conditions (technically, the principal's problem should be regular, i.e., local incentive compatibility constraints are sufficient for global incentive compatibility and the agent's participation constraint binds for the worst realization of his cost parameter), it is never optimal to burn money. Adjusting output requirements and payments to the principal's signal suffice. Fine tuning production and payment is always preferable as long as surplus has a positive social value *even when taking incentive constraints* into account.

Of course, adjusting output itself might not be always feasible. It is especially so when production is a commitment that requires investments or dedicated assets that cannot be acquired or sold at wish along the production process. In those circumstances, output can certainly be contingent on the agent's report early on but might not be on the principal's own claim *ex post*, once she learns her signal. Together with contingent payments, money burning becomes a complementary tool to write a contract contingent on the signal. We show in fact that money burning puts the agent under countervailing incentives (Lewis and Sappington, 1989a,b). Exaggerating his cost allows the agent to produce a lower quantity targeted to a less efficient type at a lower cost; saving thereby an information rent. Yet, exaggerating cost makes it more likely that the principal's signal strongly conflicts with this claim; which in turn increases the probability of money burning. Since money burning destroys surplus and the agent ends up being residual claimant for that surplus, increasing that probability reduces the benefits of exaggerating costs. Money burning then ends up being useful on two grounds. First, it might allow to extract more rent from the agent while keeping the same output distortions as absent any signal. This scenario arises when the incentive problem remains regular. The sole role of money burning is that it reduces the payment for the agent's services but the standard Baron-Myerson output distortion prevails. Second, money burning cannot only extract more rent but also reduces output distortions. This case arises when countervailing incentives are strong enough to keep a whole subset of the agent's types indifferent between participating or not. To prove the existence of such milder distortions, we use techniques from the principal-agent literature with type-dependent participation

constraints, especially the work of [Jullien \(2000\)](#) and [Martimort and Stole \(2022\)](#).

LITERATURE REVIEW. Our paper borrows its framework from [Riordan and Sappington \(1988\)](#). Those authors have also developed a simple procurement model with the additional ingredient that the principal receives an ex post signal correlated with the agent's private information. There, the signal is verifiable and contractible, allowing for first-best implementation and full rent extraction under fairly weak conditions on the conditional distribution of the signal and the agent's cost function. The intuition is that correlation between the ex post signal and the agent's type can be used to cross-check his report with the observed signal and punish (at least stochastically) any discrepancy.⁴ This paper extends this basic environment on two fronts. First, the ex post signal is non-verifiable; which *a priori* prevents first-best implementation. Second, we elaborate on the potential benefits of money burning under those circumstances.

On this front, [Strausz \(2006\)](#) also considered a pure adverse selection procurement model with binary signals and types. This author shows how the non-verifiability of the signal prevents the principal from implementing the first-best output with full rent extraction but can still be useful to improve the rent-efficiency trade-off even when budget balance is maintained. [Dequiedt and Martimort \(2015\)](#) also considered an adverse selection setup in which two agents have correlated types which are drawn from a continuous distribution.⁵ The principal must run bilateral contracts with each of them, so that the report made by one of the agents cannot be publicly used for the allocation regarding the other agent. However, because the principal may still use the information that she gathers, the contract allows for principal's discretion in the last step. In both [Strausz \(2006\)](#) and [Dequiedt and Martimort \(2015\)](#), the benefits of using the signal, though non-verifiable, is that the agent's virtual cost (in the Myersonian parlance) can now be tailored to its precise realization; which lead to implement output contingent on this signal. Neither [Strausz \(2006\)](#) nor [Dequiedt and Martimort \(2015\)](#) consider the possibility of burning money. Because the present paper provides conditions under which money burning is actually useless, we also offer in passing a rationale for the approach taken by this earlier literature.

Subjective evaluations have also been an important and lively topic for managerial economics over the recent years. More specifically, departing from the benchmark à la [Holmström \(1979\)](#) where principals may rely on objective measures of performances to solve a moral hazard problem, several authors have investigated the design of compensation schemes when only non-verifiable measures of performances are available. Here, the work of [MacLeod \(2003\)](#) stands as a leading reference.⁶ This author studies a moral hazard environment in which agent's performance is private information to the principal. Preventing the principal's opportunism might thus call for a fixed payment that dampens any incentives on the agent's side. Money burning is not only sufficient but also necessary to provide incentives to both parties. Since money burning is costly, it is not used unless in the worst state where the agent's performance is the lowest. There is a form of wage compression but it comes also with reduced effort in comparison with

⁴This implementation of the first-best bears some resemblance with the results of [Crémer and McLean \(1985\)](#), [Crémer and McLean \(1988\)](#) and [McAfee and Reny \(1992\)](#) in multi-agent environments. The difference is that in those environments, there is no ex post signal and cross-checking arises by comparing agents' reports.

⁵See also [Celik et al. \(2021\)](#).

⁶Relevant references on this front also include [Baker et al. \(1994\)](#), [Rajan and Reichelstein \(2006\)](#), [Rahman \(2012b\)](#), [Zábojník \(2014\)](#) and [Cheng \(2021\)](#) among others.

the scenario of an objective performance measure. In contrast, a screening environment has (generally) more contracting variables available. The principal’s incentive constraints can also be adjusted by modifying trading volume; which might make money burning just suboptimal. This result bears thus some resemblance with [Fuchs \(2007\)](#) who studied a repeated moral hazard environment. His main finding is that the principal can rely on the threat of terminating the contract with the agent instead of wasting payments. Termination destroys the remaining surplus from the relationship, even if such surplus is positive. However, the principal must rely on such threat to provide incentives to the agent.⁷ Along the same vein, [Lang \(2023\)](#) demonstrates that stochastic compensations may be more attractive than money burning because a risk-averse agent may incur some disutility from that randomness which has a similar impact to money burning.

Finally, [Khalil et al. \(2015\)](#) study a mixed model with both adverse selection and moral hazard elements. The principal can there also learn a non-verifiable signal and money burning is allowed. Those authors demonstrate that the date at which the non-verifiable signal is learned matters a great deal. If the effort decision is taken after the signal is realised, money burning is *a priori* optimal but rescaling the project in response to the signal may be a more attractive strategy. This result is reminiscent of the sub-optimality found for money burning in our pure adverse selection contexts. Yet, we also demonstrate the attractiveness of money burning when production is committed earlier on.

ORGANIZATION OF THE PAPER. Section 2.2 presents our model. Section 2.3 presents several benchmarks: the complete information scenario, the case where no signal is available (which boils down to [Baron and Myerson \(1982\)](#)’s well known model), the case where an ex post signal is available but it can be contracted upon. Section 2.4 describes the set of incentive feasible allocations when the ex post signal is non-verifiable. Section 2.5 characterizes the optimal contract when output can be made contingent on the principal’s report on her signal. Section 2.5.1 shows that money burning can be used to implement the first best outcome but Section 2.5.2 argues that it is not optimal when the principal’s problem is regular. Section 2.6 characterizes the optimal contract when output cannot be made contingent on the principal’s report on her signal. Money burning may now be attractive in case of strong conflicts between the principal’s and the agent’s assessments of costs. We again characterize the optimal contract under those circumstances and show the role of money burning as vehicle for countervailing incentives that might facilitate rent extraction but also ease production. Proofs are relegated to 2.A.

2.2. THE MODEL

TECHNOLOGY AND PREFERENCES. We consider the following simple model of procurement. A principal (thereafter *she*) wants to purchase a good or service from an agent (*he*). The principal derives utility $v(q, t)$ from purchasing a non-negative quantity q at price y :

$$v(q, y) = S(q) - y$$

where the gross surplus function S is twice continuously differentiable, increasing and strictly concave ($S' > 0 > S''$), and satisfies the Inada conditions $\lim_{q \rightarrow 0} S'(q) = +\infty$, $\lim_{q \rightarrow \infty} S'(q) = 0$ and $S(0) = 0$.

⁷A similar result can also be found in [Levin \(2003\)](#).

Let $\mathcal{Q} \subseteq \mathbb{R}_+$ be the set of possible quantities and $\mathcal{T} \subseteq \mathbb{R}$ is the set of possible payments.

The agent has a quasi-linear utility function:

$$u(q, t; \theta) = t - \theta q$$

where θ is his marginal cost of production and t stands for the agent's reward, which *a priori* should be lower than the principal's payment when money burning is allowed.

INFORMATION STRUCTURE. The agent has private information on the cost parameter θ . This parameter is drawn from a distribution according to a cumulative distribution function F , atomless and continuously differentiable, admitting a density function f and with support $\Theta = [\underline{\theta}, \bar{\theta}]$ with $0 < \underline{\theta} < \bar{\theta}$. As it is standard in most of the screening literature,⁸ we shall impose the following familiar assumption:

ASSUMPTION 2.1. MONOTONE HAZARD RATE PROPERTY:

$$\frac{F(\theta)}{f(\theta)} \quad \text{non-decreasing.}$$

This standard procurement model is augmented by adding the possibility for the principal to learn an informative signal σ on the agent's type. The set Σ of possible signal realizations of this signal is a compact interval of the real line, with its lowest and its maximal elements respectively denoted by $\underline{\sigma}$ and $\bar{\sigma}$. Let $\Delta(\Sigma)$ be the set of probability distributions on Σ and let $g(\sigma|\theta) \in \Delta(\Sigma)$ be the distribution of signals conditional on θ . This conditional distribution is assumed to have full support so that no signal may reveal a given type with certainty.

Following an approach already taken in [MacLeod \(2003\)](#), this signal σ stands for a subjective evaluation of the agent's performance. It is private information to the principal and thus non-verifiable by any third party. In a moral hazard context with subjective evaluation, [MacLeod \(2003\)](#) argues that a money burning is a useful tool to provide incentives to both the principal and the agent. Accordingly, we follow this author and allow the principal's payment, y , and the agent's rewards t to differ, i.e., $y = t + b$, with $b \in \mathcal{B} = [0, B]$ with B being an upper bound on the amount of money that can be burned.

ASSUMPTIONS. Let denote by $g_\theta(\sigma|\theta)$, the derivative with respect to the cost parameter θ of the conditional density of signal. For future reference, the conditional cumulative distribution will be denoted as $G(\sigma|\theta) = \int_{\underline{\sigma}}^{\sigma} g(\tilde{\sigma}|\theta) d\tilde{\sigma}$. For technical reasons, we shall assume that g is twice continuously differentiable in θ and that the score $\frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)}$ is absolutely continuous in θ . Furthermore, we shall also impose the following conditions.

ASSUMPTION 2.2. MONOTONE LIKELIHOOD RATIO PROPERTY:

$$\frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)} \quad \text{non-decreasing in } \sigma, \quad \forall \theta \in \Theta.$$

⁸[Bagnoli and Bergstrom \(2005\)](#).

Assumption 2.2 is a familiar requirement that ensures that higher values of the signal σ are more likely to arise when the agent's true type increases. It is well known that this condition also implies first-order stochastic dominance, $G_\theta(\sigma|\theta) \leq 0$, i.e., as the agent's type increases it becomes less likely that signals become good news. Next condition focuses on scenarios where there is some form of decreasing returns in this respect.

ASSUMPTION 2.3. CONVEXITY OF THE CONDITIONAL CUMULATIVE DISTRIBUTION FUNCTION:

$$G_{\theta\theta}(\sigma|\theta) \geq 0, \quad \forall \sigma \in \Sigma, \forall \theta \in \Theta$$

Convexity of the conditional distribution of signal is a more technical assumption needed at several places below to ensure that global incentive compatibility constraints always hold.

EXAMPLE 2.1. As a running example, we take θ to be uniform on $[\underline{\theta}, \bar{\theta}]$:

$$F(\theta) = \frac{\theta - \underline{\theta}}{\bar{\theta} - \underline{\theta}} \quad \text{and} \quad f(\theta) = \frac{1}{\bar{\theta} - \underline{\theta}} \quad \text{for } \theta \in [\underline{\theta}, \bar{\theta}],$$

and the signal distribution conditional on θ is:

$$g(\sigma|\theta) = (1 - \alpha) + \alpha(2\sigma F(\theta) + 2(1 - \sigma)[1 - F(\theta)]) \quad \text{for } \sigma \in [0, 1],$$

where $\alpha \in [0, 1]$ is a measure of the informativeness of the signal: if $\alpha = 0$, the signal is uniformly distributed on $[0, 1]$; if $\alpha > 0$, the signal σ is positively correlated with θ . In particular, Assumptions 2.1, 2.2 and 2.3 are satisfied.

Moreover, we specify $S(q) = 2q^{1/2}$, so that Inada conditions are satisfied.

TERMINOLOGY. Since the relevant incentive problem is to deter over-reporting of the cost parameter θ , we define news relative to the agent's report. A high realization of σ will be called *good news*, as it is more consistent with a high reported cost and therefore makes over-reporting less likely. A low realization of σ will be called *bad news*, as it is more indicative of a low cost parameter and therefore more suggestive of over-reporting. Thus, "good" and "bad" refer to consistency with the report, not to whether the signal is favorable or unfavorable about the true type.

MECHANISMS. A direct revelation mechanism (*DRM*) in our environment is an allocation rule $\{(q(\hat{\theta}, \hat{\sigma}), t(\hat{\theta}, \hat{\sigma}), b(\hat{\theta}, \hat{\sigma}))\}_{\hat{\theta} \in \Theta, \hat{\sigma} \in \Sigma}$, that stipulates a quantity $q(\hat{\theta}, \hat{\sigma})$, a payment $t(\hat{\theta}, \hat{\sigma})$ to the agent and an amount of money to be burnt $b(\hat{\theta}, \hat{\sigma})$ as a function of the agent's report $\hat{\theta} \in \Theta$ on her type and the principal's report on his signal $\hat{\sigma} \in \Sigma$.

TIMING. The game unfolds as follows.

1. The principal proposes a *DRM*, namely $\{(q(\hat{\theta}, \hat{\sigma}), t(\hat{\theta}, \hat{\sigma}), b(\hat{\theta}, \hat{\sigma}))\}_{\hat{\theta} \in \Theta, \hat{\sigma} \in \Sigma}$.

2. The agent learns θ .
3. The agent accepts or rejects the mechanism. Following acceptance, the agent sends a report $\hat{\theta} \in \Theta$ to the mechanism. Otherwise, the principal and the agent both get their reservation values that are normalized at zero.
4. The principal observes the report $\hat{\theta}$, learns the signal σ , and then reports $\hat{\sigma} \in \Sigma$ to the mechanism.^{9,10}
5. The allocation $(q(\hat{\theta}, \hat{\sigma}), t(\hat{\theta}, \hat{\sigma}), b(\hat{\theta}, \hat{\sigma}))$ is implemented. Payoffs are realized.

A *DRM* as described above, with the sequential timing of communication stands as a metaphor for other indirect mechanisms that would appear as more natural in practice but may be more complex to analyze without the apparatus of direct communication. To illustrate, the production process may be a dynamic one, with earlier production stages producing a signal on the agent's underlying cost that is privately observed by the principal. The agent's report in the direct communication game corresponds to the choice of a whole menu of production plans before the informative signal realizes and is learned by the principal in the indirect version. Such a plan stipulates a final output, a payment received the principal and the amount of money burnt as a function of this output. A menu is a family of such plans indexed by the signal realization. Finally, the principal picks an item within a menu of such contingent plans according to the signal she has learned. To reflect the sequential turn of actions in the indirect version of the game, communication in the direct version is also sequential. As proved in a related context by [Dequiedt and Martimort \(2015\)](#), the Revelation Principle still applies in our procurement context provided incentive constraints are properly defined for both the principal and the agent.¹¹

In the first part of the paper, we will consider an environment where such output plans, contingent on the principal's report on his signal are feasible. In practice, it corresponds to scenarios where production is an ongoing process that can easily be modified. Output adjustments may be useful to better screen the agent's cost parameter. In Section 2.6, we consider an alternative timing where the signal is learned after output has been committed. This scenario is meant to capture environments where the agent must incur some (unmodeled) investment (productive assets) that determines the scale of production and, in the short run, no additional investment or no sale of existing assets is feasible.

2.3. BENCHMARKS

COMPLETE INFORMATION. When θ is common knowledge, there is no point using a subjective evaluation and money burning is useless. A simple forcing contract can be used to induce the first-best output and extracts all surplus from the agent. The first-best outcome entails:

$$S'(q^{fb}(\theta)) = \theta, \quad U^{fb}(\theta) = 0, \quad b^{fb}(\theta) = 0, \quad \forall \theta \in \Theta. \quad (2.3.1)$$

⁹The exact timing has no impact. It could be that the principal learns her signal before observing the agent's report with no change in the optimal contract.

¹⁰Our timing is such that the principal learns her private information after having designed the mechanism; in contrast with [Cella \(2008\)](#) who analyzed an informed principal's problem with correlated information.

¹¹See 2.B for details.

For future reference, let $W^{fb}(\theta) = S(q^{fb}(\theta)) - \theta q^{fb}(\theta)$ be the first-best surplus and let $\mathcal{W}^{fb} = \mathbb{E}_\theta(W^{fb}(\theta))$ its expected value.

ASYMMETRIC INFORMATION WITH NO SIGNAL. This scenario boils down to the standard procurement model à la [Baron and Myerson \(1982\)](#). In particular, there is no need for money burning while the rest of the solution is well known and highlights the main features of the rent-efficiency trade-off of screening models:

$$S'(q^{bm}(\theta)) = \theta + \frac{F(\theta)}{f(\theta)}, \quad U^{bm}(\theta) = \int_\theta^{\bar{\theta}} q^{bm}(x) dx \geq 0, \quad b^{bm}(\theta) = 0, \quad \forall \theta \in \Theta. \quad (2.3.2)$$

There is a downward distortion of output for all types except at the top and all types except the least-efficient one receives a positive information rent. Assumption 2.1 ensures that q^{bm} so defined is non-increasing as requested by incentive compatibility while the assumptions made on the surplus function ensure that q^{bm} is always positive so that all types always produce.

ASYMMETRIC INFORMATION WITH A VERIFIABLE SIGNAL. Suppose now that the signal σ is verifiable and thus can be contracted upon. Clearly, money burning is useless in this context. Money burning wastes resource and it has no benefits whatsoever in relaxing the agent's incentive constraints; payments contingent on σ suffices.

From [Riordan and Sappington \(1988\)](#), and following the logic of [Cr  mer and McLean \(1988\)](#) and [McAfee and Reny \(1992\)](#)'s more general results, it is well-known that the first-best outcome can be implemented with no surplus left to the agent under weak assumptions on the distribution of the verifiable signal. To see how, consider the following payment schedule

$$t(\hat{\theta}, \sigma) = S(q^{fb}(\hat{\theta})) - W^{fb}(\hat{\theta}) + z(\hat{\theta}, \sigma) \quad \forall \hat{\theta} \in \Theta, \forall \sigma \in \Sigma. \quad (2.3.3)$$

Those payments make the agent residual claimant for the choice of the output $q^{fb}(\hat{\theta})$ but request him to pay a fee $W^{fb}(\hat{\theta})$ independent of σ that is designed to extract the agent's surplus. This fee is augmented by an extra payment $z(\hat{\theta}, \sigma)$ contingent on the realized signal. This payment corrects any stochastic discrepancy that may arise between the agent's report $\hat{\theta}$ and the realized ex post signal.

PROPOSITION 2.1. *Suppose that σ is verifiable and that Assumption 2.3 holds. There exists an optimal contract that implements the first-best outcome (2.3.1). In such optimal contract, money burning is useless and the transfers described in (2.3.3) are such that:*

$$z(\hat{\theta}, \sigma) = q^{fb}(\hat{\theta}) \frac{g_\theta(\sigma|\hat{\theta})}{g(\sigma|\hat{\theta})} \frac{1}{\mathbb{E}_\sigma \left(\left(\frac{g_\theta(\sigma|\hat{\theta})}{g(\sigma|\hat{\theta})} \right)^2 \middle| \hat{\theta} \right)} \quad (2.3.4)$$

The first-best implementation when σ is verifiable relies on the possibility for the principal to offer lotteries $z(\hat{\theta}, \sigma)$ that tie the agent's ex post payment to the realization of the signal; punishing the agent when the ex post signal conflicts too much with his report and rewarding otherwise. In particular, such lotteries can be reinterpreted as the principal estimating the type of the agent by maximum likelihood

methods, and comparing this estimate with the agent's report. The score $s(\sigma|\theta) \equiv \frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)}$ indicates the sign of the gradient of the log-likelihood with respect to θ . If the agent reports $\hat{\theta}$ and the realized signal is such that $s(\sigma|\hat{\theta}) < 0$, then the principal would find that the likelihood would be higher with a lower report of the type. Since the main incentive problem is to prevent over-reporting the type, the principal must punish such combination of signal and report, and thus $z(\hat{\theta}, \sigma) < 0$. On the contrary, the principal rewards the agent with $z(\hat{\theta}, \sigma) > 0$ when $s(\sigma|\hat{\theta}) > 0$, and there is neither a punishment nor a reward when the reported type of the agent matches the maximum likelihood estimate based on σ . The size of rewards and punishments is scaled by the inverse of the Fisher information $\mathcal{I}(\theta) \equiv \mathbb{E}_\sigma \left(\left(\frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)} \right)^2 \middle| \theta \right)$, which is the variance of the score. If the signal contains little information about θ , the principal must use extreme lotteries to prevent misreporting. Thus, the higher this information, the smaller the scale of the transfers for any realization of the signal.




Figure 2.1: Transfers that implement the first-best allocation of output and zero rent when σ is verifiable. For $\sigma = 0.5$, the score of the signal is zero, which corresponds to a case with neither good nor bad news: the agent is compensated for the total cost of production $\theta q^{fb}(\theta)$. For $\sigma = 0$, the principal punishes the agent with a negative transfer, in order to disincentivize over-reporting. For $\sigma = 1$, the principal rewards the agent with higher transfers.

2.4. INCENTIVE-FEASIBLE ALLOCATIONS

When deciding whether to accept the mechanism or not, the agent must form an expectation about the possible reports of the principal. Expecting a truthful report of the principal's subjective signal, the agent's payoff must be computed taking expectation over σ conditional on the realised type θ . Accordingly, we define the agent's *interim rent* when his type is θ as

$$U(\theta) = \mathbb{E}_\sigma (t(\theta, \sigma) - \theta q(\theta, \sigma) \mid \theta).$$

The agent's participation constraint can thus be written as

$$U(\theta) \geq 0, \quad \forall \theta \in \Theta. \quad (2.4.1)$$

Being given that the principal will report truthfully her subjective signal, the agent's incentive-compatibility constraints write as

$$U(\theta) = \max_{\hat{\theta} \in \Theta} \mathbb{E}_{\sigma} \left(t(\hat{\theta}, \sigma) - \theta q(\hat{\theta}, \sigma) | \theta \right), \quad \forall \theta \in \Theta. \quad (2.4.2)$$

Additional incentive compatibility constraints must be incorporated to indeed induce the principal to truthfully report her signal. Importantly, note that these constraints are written ex post, once the agent has reported (truthfully in equilibrium) her type θ . For any such report, the principal must prefer not to manipulate her subjective evaluation of the agent; a condition written as

$$S(q(\theta, \sigma)) - t(\theta, \sigma) - b(\theta, \sigma) \geq S(q(\theta, \hat{\sigma})) - t(\theta, \hat{\sigma}) - b(\theta, \hat{\sigma}), \quad \forall \theta \in \Theta, \forall (\sigma, \hat{\sigma}) \in \Sigma^2. \quad (2.4.3)$$

Because σ is not payoff-relevant, the role of σ and $\hat{\sigma}$ can be reversed. In fact, the principal should be made indifferent between any of her available reports and (2.4.3) amounts to the existence of a function $H(\theta)$ that correspond to the constant value of the principal's net payoff over all realizations of her signal:

$$S(q(\theta, \sigma)) - t(\theta, \sigma) - b(\theta, \sigma) = H(\theta) \quad \forall \theta \in \Theta, \forall \sigma \in \Sigma. \quad (2.4.4)$$

The principal is thus willing to report truthfully as she is indifferent between any of her available reports. The contract takes the form of a sell-out contract. Anticipating such truthful reporting, the agent is right to evaluate his own incentives to report the truth and participate with the true expectation on signal conditional on his type.

The principal's incentive constraint (2.4.4) deserves further comments. In a standard moral hazard à la MacLeod (2003), the principal can only use transfers to incentivize her agent. She can reduce the wage being paid by reporting a particular realization of the subjective signal on the agent's performance and she would thus always choose that very report. But then the payment schedule can no longer depend on performance and incentives to exert effort can no longer be provided. Money burning is thus necessary to simultaneously induce truthful reporting by the principal and incentivize effort by the agent. In our adverse selection setup, money burning is no longer necessary. It could still be possible to satisfy (2.4.4) without burning money (i.e., setting $b(\theta, \sigma) \equiv 0$). In particular, different combinations of outputs and transfers $(q(\theta, \sigma), t(\theta, \sigma))$ could be implemented for different realizations of σ the principal's payoff is kept constant. Hence, money burning is not necessary in a screening environment.

Using the expression of the transfers $t(\theta, \sigma)$ coming from (2.4.4), the agent's incentive compatibility constraints can be written in terms of interim payoffs as

$$U(\theta) = \max_{\hat{\theta} \in \Theta} -H(\hat{\theta}) + \mathbb{E}_{\sigma} \left(S(q(\hat{\theta}, \sigma)) - \theta q(\hat{\theta}, \sigma) - b(\hat{\theta}, \sigma) | \theta \right), \quad \forall \theta \in \Theta. \quad (2.4.5)$$

This expression makes it clear that the agent's is made residual claimant for the maximization of the overall surplus $S(q) - \theta q - b$, choosing among production plans $q(\hat{\theta}, \sigma)$ contingent on the ex post signal σ that each corresponds to a payment back to the principal worth $H(\hat{\theta})$.

THE PRINCIPAL'S PROBLEM. Expressing also the maximand in terms of the agent's interim rent, the principal's problem is to choose an allocation $(q(\theta, \sigma), b(\theta, \sigma), U(\theta))_{(\theta, \sigma) \in \Theta \times \Sigma}$ that solves

$$\begin{aligned} (\mathcal{P}) : \max_{(q, b, U)} \quad & \mathbb{E}_{\theta, \sigma} (S(q(\theta, \sigma)) - \theta q(\theta, \sigma) - b(\theta, \sigma) - U(\theta)) \\ \text{subject to} \quad & (2.4.1), (2.4.5) \text{ and} \\ & b(\theta, \sigma) \in [0, B], \quad \forall (\theta, \sigma) \in \Theta \times \Sigma. \end{aligned} \quad (2.4.6)$$

This expression of the principal's maximization problem showcases an extension of the classic rent-efficiency trade-off. The principal would like to maximise the surplus that is created and reduce the rent that is left to the agent. In sharp contrast with the standard procurement model, part of the surplus can here be burnt. Yet, money burning can also be used to relax the agent's incentive compatibility constraint (2.4.5).

NECESSARY CONDITIONS FOR INCENTIVE COMPATIBILITY. A standard approach in mechanism design consists in studying a relaxed maximization problem where the whole set of incentive compatibility constraints (2.4.5) (whose description already encompasses the characterization of the principal's incentive compatibility constraints (2.4.4) is replaced by a necessary Envelope Condition. Next Lemma expresses this condition.

LEMMA 2.1. *An allocation $(q(\theta, \sigma), b(\theta, \sigma), U(\theta))$ satisfying (2.4.5) is necessarily such that U is absolutely continuous and thus a.e. differentiable with a derivative when it exists given by*

$$\dot{U}(\theta) = \mathbb{E}_{\sigma} \left(-q(\theta, \sigma) + (S(q(\theta, \sigma)) - \theta q(\theta, \sigma) - b(\theta, \sigma)) \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \middle| \theta \right). \quad (2.4.7)$$

2.5. WHEN THE SIGNAL IS LEARNED BEFORE PRODUCTION

2.5.1. On the (Possible) Value of Money Burning

The full surplus-extraction result of Proposition 2.1 that holds even in the absence of any money burning crucially depends on the fact that the signal σ was supposed to be verifiable. Next Proposition shows that with non-verifiability of the signal the first-best can no longer be implemented unless money is burned.

PROPOSITION 2.2. *Suppose that σ is non-verifiable. The first-best allocation $(q^{fb}(\theta), U^{fb}(\theta))$*

1. *cannot be implemented without money burning;*

2. can be implemented with money burning when Assumption 2.3 holds and B is large enough. It is so by using money burning following bad news, i.e.,

$$b(\theta, \sigma) = \begin{cases} B & \text{if } \sigma \in [\underline{\sigma}, \sigma^{fb}(\theta)), \\ 0 & \text{if } \sigma \in (\sigma^{fb}(\theta), \bar{\sigma}], \end{cases} \quad (2.5.1)$$

where $\sigma^{fb}(\theta)$ solves

$$-q^{fb}(\theta) = BG_{\theta}(\sigma^{fb}(\theta)|\theta). \quad (2.5.2)$$

The intuition for Item 1. is straightforward and can easily be grasped from the expression of the transfers in (2.3.3). Lotteries are no longer possible when the principal can manipulate her report on the realized signal and money burning is not feasible. The principal would just report the signal with the lowest such transfer $\min_{\sigma \in \Sigma} z(\hat{\theta}, \sigma)$; which voids lotteries of their value.

Item 2. shows that, if enough money can be burned, full extraction of the agent's surplus is possible and the first-best output can be implemented provided that the principal burns the largest possible amount B whenever the signal she observes is sufficiently bad news relative to the agent's report. To understand this result, observe that the agent wants to overstate his cost parameter θ and claim he has a cost $\theta + d\theta$ to produce a lower quantity $q^{fb}(\theta + d\theta)$ at a lower marginal cost and thereby save $q^{fb}(\theta + d\theta)d\theta \approx q^{fb}(\theta)d\theta$. Of course, this effect would still be present in the absence of any ex post signal. With an ex post signal that can be manipulated by the principal, (2.4.4) constraint requires that the principal gets a payoff independent of the signal. In other words, the agent ends up being residual claimant not only for the whole gains from trade but also for the money burned. To create countervailing incentives¹² and thus better extract the agent's surplus, more money should thus be burned as the agent overstates his cost parameter.

2.5.2. On the Irrelevance of Money Burning

WHEN THE INCENTIVE PROBLEM IS REGULAR. In a first pass and to simplify the analysis, we first focus on a *relaxed* incentive problem ($\mathcal{P}^r$) where the requirement of incentive compatibility (2.4.5) reduces to the necessary condition (2.4.7) while the participation constraint (2.4.1) is now only imposed for the least efficient type:

$$U(\bar{\theta}) \geq 0. \quad (2.5.3)$$

DEFINITION 2.1. ($\mathcal{P}$) is said to be regular when its solutions also solve ($\mathcal{P}^r$).

A key assumption to ensure regularity is the following generalization of Assumption 2.1.

ASSUMPTION 2.4. GENERALIZED MONOTONE HAZARD RATE PROPERTY:

$$\frac{\frac{F(\theta)}{f(\theta)}}{1 + \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta)}{f(\theta)}} \begin{cases} \text{non-decreasing in } \theta, & \forall \sigma \in \Sigma, \\ \text{non-increasing in } \sigma, & \forall \theta \in \Theta, \end{cases} \quad (2.5.4)$$

¹²Lewis and Sappington (1989a,b).




Figure 2.2: In blue, the threshold for the signal σ conditional on θ that determines money burning. In purple, dashed, the probability of money burning conditional on the reported type θ .

with

$$1 + \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta)}{f(\theta)} > 0, \quad \forall (\theta, \sigma) \in \Theta \times \Sigma. \quad (2.5.5)$$

The first of those monotonicity conditions ensures that the second-best solution $q^{sb}(\theta, \sigma)$ remains non-increasing with type while the second monotonicity condition that follows from Assumption 2.2 shall ensure that the agent will produce more when signals are good news. Example 2.1 satisfies this property if $\alpha \leq \frac{1}{3}$, i.e., the signal is not too informative about the cost θ . The full characterization of the second-best solution is provided below.

PROPOSITION 2.3. *Suppose that Assumption 2.4 holds. The optimal contract solution to $(\mathcal{P}^r)$ entails the following features.*

1. The second-best output $q^{sb}(\theta, \sigma)$ entails a generalized Baron-Myerson distortion:

$$S'(q^{sb}(\theta, \sigma)) = \theta + \frac{\frac{F(\theta)}{f(\theta)}}{1 + \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta)}{f(\theta)}}, \quad \forall (\theta, \sigma) \in \Theta \times \Sigma. \quad (2.5.6)$$

$q^{sb}(\theta, \sigma)$ is positive, non-increasing (resp. non-decreasing) in θ (resp. in σ).

2. Money burning is useless:

$$b^{sb}(\theta, \sigma) = 0, \quad \forall (\theta, \sigma) \in \Theta \times \Sigma. \quad (2.5.7)$$

3. The agent's information rent is given by:

$$U^{sb}(\theta) = \int_{\theta}^{\bar{\theta}} \mathbb{E}_{\sigma} \left(q^{sb}(\tilde{\theta}, \sigma) - (S(q^{sb}(\tilde{\theta}, \sigma)) - \tilde{\theta} q^{sb}(\tilde{\theta}, \sigma)) \frac{g_{\theta}(\sigma|\tilde{\theta})}{g(\sigma|\tilde{\theta})} \middle| \tilde{\theta} \right) d\tilde{\theta}. \quad (2.5.8)$$

4. The incentive problem is regular provided that $\Delta\theta = \bar{\theta} - \underline{\theta}$ is small enough and

$$1 + \frac{\frac{G_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \left(\frac{F(\theta)}{f(\theta)} \right)^2}{1 + \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta)}{f(\theta)}} \frac{d}{d\sigma} \left(\frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \right) \geq 0, \quad \forall (\theta, \sigma) \in \Theta \times \Sigma. \quad (2.5.9)$$

Output Distortions. To understand (2.5.6), it is first useful to think about the case where the signal σ would be uninformative (i.e., $g_{\theta} \equiv 0$). In that polar scenario, the signal has no value and the optimal output exhibits the standard Baron-Myerson distortion expressed in (2.3.2). With an informative signal, the principal instead would like to link the agent's payment to what she learns from the signal if the latter was verifiable and the first-best output could be implemented as shown in Proposition 2.1. When this signal is non-verifiable, the principal's incentive compatibility constraints (2.4.4) limit how flexible payments can be made to fulfil this task. Output distortions, contingent on the signal's realization, are always needed for any realization of the signal while full rent extraction is no longer possible as shown in (2.5.8). Because the agent's rent is costly for the principal, downward distortions of output below the first-best level are needed to reduce this rent. The logic is familiar from most screening models. Yet, details differ. The comparison of (2.5.6) with the standard Baron-Myerson formula (2.3.2) showcases the nature of those output distortions. At a rough level, the logic of screening models à la [Baron and Myerson \(1982\)](#) still holds; the optimal output is such that the marginal benefit of production is equal to what can be viewed as (marginal) *generalized virtual cost* of the agent.¹³ Indeed, Assumption 2.2 implies that a signal σ such that $\frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} > 0$ (resp. < 0) is *good news* (resp. *bad news*) and it makes the principal think that the agent has not exaggerated his cost parameter. Output distortions are less (resp. more) attractive under those conditions.

No Money Burning. Remarkably, output distortions are not only necessary but are also sufficient to solve the bilateral incentive problem between the principal and her agent. When (2.5.5) holds, the principal's optimization problem remains strictly concave in output. Production remains valuable while money burning is necessarily costly. Only output is used to take into account signal realizations and no money needs to be burned. For a regular problem, imposing budget balance between the principal and the agent is enough to reach the informationally constrained optimum.

The Value of the Signal. Using the fact that no money is burned, we may rewrite (2.4.7) as

$$\dot{U}^{sb}(\theta) = \mathbb{E}_{\sigma} \left(-q^{sb}(\theta, \sigma) + (S(q^{sb}(\theta, \sigma)) - \theta q^{sb}(\theta, \sigma)) \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \middle| \theta \right).$$

¹³The expression was coined by [Dequiedt and Martimort \(2015\)](#).




Figure 2.3: Output schemes. In blue, the first-best level of output. For $\sigma = 0.5$ (purple), the score $\frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)}$ of the signal is zero, which corresponds to a case with neither good nor bad news: output hence equals the Baron-Myerson amount. For $\sigma = 0$ (green, dashed), the principal exacerbates the downwards distortions, in order to disincentivize over-reporting. For $\sigma = 1$ (orange, dotted), the principal softens the distortions in output.

Integrating by parts the second term in this expectation yields

$$\dot{U}^{sb}(\theta) = \mathbb{E}_\sigma \left(-q^{sb}(\theta, \sigma) | \theta \right) - \mathbb{E}_\sigma \left((S'(q^{sb}(\theta, \sigma)) - \theta) q_\sigma^{sb}(\theta, \sigma) \frac{G_\theta(\sigma|\theta)}{g(\sigma|\theta)} | \theta \right). \quad (2.5.10)$$

For a non-informative signal, $g_\theta = G_\theta \equiv 0$, and the second term above disappears. Otherwise, the fact that q^{sb} is distorted downwards (i.e., $S'(q^{sb}(\theta, \sigma)) \geq \theta$) together with $q_\sigma^{sb} \geq 0$ and $G_\theta \leq 0$ (which are both implied by Assumption 2.2) implies that this second term on the right-hand side of (2.5.10) is always non-negative. This sign captures the fact that, overall, the principal can design countervailing incentives for the agent once the signal is available, even if it remains non-verifiable.

Average Distortions. Different realizations of the signal σ induce different distortions on output, as shown in (2.5.6). When the realized signal is good news, the optimal contract implements smaller distortions; while larger distortions are introduced with bad news. However, thanks to strict convexity of the distortions in the score $\frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)}$, Jensen's inequality allows us to derive:

$$\mathbb{E}_\sigma \left(\frac{\frac{F(\theta)}{f(\theta)}}{1 + \frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta)}{f(\theta)}} \middle| \theta \right) \geq \frac{\frac{F(\theta)}{f(\theta)}}{1 + \mathbb{E}_\sigma \left(\frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)} \middle| \theta \right) \frac{F(\theta)}{f(\theta)}} = \frac{F(\theta)}{f(\theta)}, \quad (2.5.11)$$

with strict inequality provided that $g_\theta \neq 0$. In words, for each type θ , the expected distortion induced in the optimal contract is larger when the principal has a private signal available.

“VERY BAD NEWS”. As a corollary of those findings, money might only be burned when $(\mathcal{P})$ is regular if (2.5.5) fails for some combinations of signal and type. This possibility arises when σ is *very bad news* on the agent’s cost, i.e., $g_\theta(\sigma|\theta)$ is sufficiently negative at such a combination. It characterizes a scenario of strong conflict between the principal’s assessment of the agent’s performance and the latter’s report on his type. In Example 2.1, this occurs when $\alpha > \frac{1}{3}$.

To explore the consequences of this possibility for contract design, we first consider the case of a regular problem. Next Assumption, which is a strengthening of Assumption 2.1 gives us a clear configuration of the optimal contract under those circumstances.

ASSUMPTION 2.5. MIXED MONOTONICITY.

$$\frac{F(\theta)}{f(\theta)} + \frac{g(\sigma|\theta)}{g_\theta(\sigma|\theta)} \quad \text{non-decreasing in } \theta, \quad \forall \sigma \in \Sigma.$$

PROPOSITION 2.4. Suppose that Assumption 2.5 holds and that there exists $\theta^* \in (\underline{\theta}, \bar{\theta}]$ such that

$$1 + \frac{g_\theta(\underline{\sigma}|\theta^*)}{g(\underline{\sigma}|\theta^*)} \frac{F(\theta^*)}{f(\theta^*)} = 0. \quad (2.5.12)$$

Then, there exists a threshold value $\sigma^{sb}(\theta) \in (\underline{\sigma}, \bar{\sigma}]$ ¹⁴ defined as

$$1 + \frac{g_\theta(\sigma^{sb}(\theta)|\theta)}{g(\sigma^{sb}(\theta)|\theta)} \frac{F(\theta)}{f(\theta)} = 0, \quad \forall \theta \in [\theta^*, \bar{\theta}]. \quad (2.5.13)$$

such that solution to the relaxed problem $(\mathcal{P}^r)$ is characterized as follows.

1. The optimal output entails:


$$\begin{cases} q^{sb}(\theta, \sigma) = 0 & \text{if } \theta \in [\theta^*, \bar{\theta}] \text{ and } \sigma \in [\underline{\sigma}, \sigma^{sb}(\theta)], \\ S'(q^{sb}(\theta, \sigma)) = \theta + \frac{\frac{F(\theta)}{f(\theta)}}{1 + \frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta)}{f(\theta)}}, & \text{otherwise.} \end{cases} \quad (2.5.14)$$

2. Money is burned following very bad news only:


$$b^{sb}(\theta, \sigma) = \begin{cases} B & \text{if } \theta \in [\theta^*, \bar{\theta}] \text{ and } \sigma \in [\underline{\sigma}, \sigma^{sb}(\theta)], \\ 0 & \text{otherwise.} \end{cases} \quad (2.5.15)$$

Proposition 2.4 demonstrates that burning money is now useful in the case of a strong conflict between the principal’s assessment and the agent’s report; the agent claiming a cost parameter which is high enough while the principal’s evaluation is much more favorable. We already know from Proposition 2.3 that, following *bad news*, the principal wants to reduce the agent’s output. The novelty is that, when Assumption 2.5 and Condition (2.5.12) hold altogether, the principal prefers not to ask for any

¹⁴With the convention that $\sigma^{sb}(\theta) = \underline{\sigma}$ for $\theta \in [\underline{\theta}, \theta^*)$.



(a) Output schemes. In blue, the first-best level of output. For $\sigma = 0.5$ (purple), the score $\frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)}$ of the signal is zero, which corresponds to a case with neither good nor bad news: output hence equals the Baron-Myerson amount. For $\sigma = 0$ (green, dashed), the principal exacerbates the downwards distortions, in order to disincentivize over-reporting. Distortions can be large enough to bring output to zero for the upper tail of types $\theta > \theta^*$. For $\sigma = 1$ (orange, dotted), the principal softens the distortions in output.



(b) In blue, the threshold for the signal σ conditional on θ that determines money burning. In purple, dashed, the probability of money burning conditional on the reported type θ .

Figure 2.4: Signal-dependent distortions and money burning thresholds.

production so as there are no gains from trade. At the same time, even more surplus can be destroyed when money is also burned. Money burning acts as a credible threat to extract more of the agent's surplus when reports are strongly conflicting. Instead, when the principal's signal does not conflict much with the agent's report, production remains attractive. The principal finds profitable to reduce output but retains from burning money.

Nonregularity. The availability of the signal σ may cause the contracting problem to be non-regular. Indeed [Strausz \(2006\)](#) shows in his binary model without money burning that the participation constraint may be binding for both the inefficient and the efficient type. In our continuous type space scenario, dealing with this problem requires the use of techniques borrowed from [Jullien \(2000\)](#) and [Martimort and Stole \(2022\)](#). In [2.C](#) we extend Propositions [2.3](#) and [2.4](#) to consider such a possibility.

2.6. WHEN THE SIGNAL IS LEARNED AFTER PRODUCTION

In order to further illustrate the use of money burning, we consider an alternative scenario where the signal σ is realised after output q has already been chosen by the agent. Under those circumstances, production can no longer be made contingent on the principal's report about the signal. Yet, payments can be delayed to account for this possibility. We will thus simplify notations accordingly by suppressing the dependence of output on signal and we shall write $q(\hat{\theta})$. The principal's incentive compatibility constraints ([2.4.4](#)) show that there must exist a function H' such that the principal overall payment, including money burned, must be kept constant across all realizations of the signal as:

$$t(\theta, \sigma) + b(\theta, \sigma) = H'(\theta) \quad \forall \theta \in \Theta, \forall \sigma \in \Sigma. \quad (2.6.1)$$

If the principal were to use different payments for the agent's services as the signal he receives varies, those payments should be compensated by equal variations in the amount of money burned. To illustrate, suppose that the principal wants to punish the agent in case she gets a favorable signal by lowering the payment $t(\theta, \sigma_1)$ when σ_1 is small enough while, at the same time, she would like to reward the agent when the signal is less favorable, i.e., setting $t(\theta, \sigma_2) > t(\theta, \sigma_1)$ when $\sigma_2 > \sigma_1$. Since output can no longer vary with the signal, the only way of making such signal-dependent rewards incentive compatible for the principal is to also increase money burning following σ_1 precisely by an amount $t(\theta, \sigma_2) - t(\theta, \sigma_1)$. In other words, money burning is given its best chance when the principal learns signal after production because it is only with money burning that the principal can credibly commit to incorporate her subjective evaluation in the agent's compensation.

Adapting our findings in Lemma [2.1](#), the agent's rent profile must now satisfy a simpler Envelope Condition

$$\dot{U}(\theta) = -q(\theta) - \mathbb{E}_\sigma \left(b(\theta, \sigma) \frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)} | \theta \right). \quad (2.6.2)$$

This condition is already illuminative of the key forces at play. By eventually choosing to burn money for all signal realizations σ such that $g_\theta(\sigma|\theta) < 0$, the principal creates countervailing incentives for his agent. On the one hand, the latter wants again to overstate his cost parameter θ and claim he has a cost $\theta + d\theta$ to save $q(\theta + d\theta)d\theta \approx q(\theta)d\theta$ in doing so. On the other hand, the likelihood of a favorable

signal (i.e., σ small enough) is greater at type θ than at type $\theta + d\theta$. Burning money for sufficiently favorable signals thus induces a significant cost on the exaggerating agent which limits those incentives to exaggerate costs.

We are now equipped for a characterization of the optimal contract. We shall distinguish two scenarios depending on how strong those countervailing incentives are.

2.6.1. Wage Compression

In the first scenario, the maximal amount of money that can be burned is not too large. The countervailing effect discussed above is of limited magnitude even if the principal uses money burning with some probability, doing so does not perturb too much the standard lessons of the Baron-Myerson scenario. The incentive problem $(\mathcal{P})$ remains regular, with the agent's participation constraint being again reduced to (2.5.3) and the Envelope Condition (2.6.2) being sufficient for (2.4.5).

PROPOSITION 2.5. *Suppose that Assumption 2.5 holds and again define θ^* as in (2.5.12). Then, the solution to the relaxed problem $(\mathcal{P}^r)$ entails the following properties.*

1. *The optimal output is always the Baron-Myerson outcome:*

$$q^{sb}(\theta) = q^{bm}(\theta), \quad \forall \theta \in \Theta. \quad (2.6.3)$$

2. *Money is burned following very bad news only:*

$$b^{sb}(\theta, \sigma) = \begin{cases} B & \text{if } \theta \in [\theta^*, \bar{\theta}] \text{ and } \sigma \in [\underline{\sigma}, \sigma^{sb}(\theta)], \\ 0 & \text{otherwise.} \end{cases} \quad (2.6.4)$$

$(\mathcal{P})$ is regular provided that

$$B \in \left[0, \min_{\theta \in [\theta^*, \bar{\theta}]} \left\{ \frac{-q^{bm}(\theta)}{G_{\theta}(\sigma^{sb}(\theta)|\theta)} \right\} \right]. \quad (2.6.5)$$

Proposition 2.5 delineates circumstances where the sole role of money burning is to reduce payments; a form of wage compression. It has no consequences whatsoever for the volume of trade and output remains at its Baron-Myerson level.

This result is reminiscent of MacLeod (2003). This author demonstrates that, under pure moral hazard, the optimal wage exhibits some form of compression. The principal commits to a fixed payment, independent of her own subjective evaluation of the agent, for all outcomes except the worst one and burn some money in that state. The incentive power of such a scheme is reduced in comparison with the ideal scenario where an objective performance evaluation would be available. As a result, the agent's effort and its liability rent are both reduced. Proposition 2.5 shows also that the agent's information rent is somewhat reduced in comparison in our pure screening scenario. Indeed, using the fact that

(2.5.3) is binding at the optimum for a regular problem and the Envelope Condition (2.6.2), a simple integration gives us the expression of the agent's information rent as

$$U^{sb}(\theta) = \int_{\theta}^{\bar{\theta}} (q^{bm}(\tilde{\theta}) + BG_{\theta}(\sigma^{sb}(\tilde{\theta})|\tilde{\theta})) d\tilde{\theta} \leq U^{bm}(\theta), \quad \forall \theta \in \Theta; \quad (2.6.6)$$

where the inequality is strict for all $\theta < \bar{\theta}$ since money burning occurs with positive probability for θ close enough of $\bar{\theta}$. From there, we may rewrite the agent's overall payment as

$$S(q^{bm}(\theta)) - BG(\sigma^{sb}(\theta)|\theta) - H^{sb}(\theta) = U^{sb}(\theta) + \theta q^{bm}(\theta) \quad (2.6.7)$$

$$\leq U^{bm}(\theta) + \theta q^{bm}(\theta) = S(q^{bm}(\theta)) - H^{bm}(\theta), \quad \forall \theta \in \Theta$$

where the right-hand side stands for the payment received in the absence of any signal and has been rewritten to stress, that even in that scenario, one possible implementation of the Baron-Myerson outcome is to let the agent enjoy all gross surplus of the transaction $S(q^{bm}(\theta))$ against a fee $H^{bm}(\theta)$. Condition (2.6.7) then shows that the agent is asked to disgorge more of the overall surplus of the transaction when a signal is available.

A remarkable finding of Proposition 2.5 is that the volume of trade remains unchanged, in contrast with MacLeod (2003)'s findings under pure moral hazard. The reason is simple. In our hidden information environment, the principal can use both quantity and payments for screening purposes. When production is chosen before any signal can be learned, money burning can help to save on payments but, as long as countervailing incentives are not too strong, the reduction of payments available when the signal can be used does not interact with output requirements. Instead, in a pure moral hazard environment, payments have always a direct impact on effort.

2.6.2. Output Expansion

Consider now a scenario where countervailing incentives are more pronounced. That is, condition (2.6.5) fails to hold: $B > \min_{\theta \in [\theta^*, \bar{\theta}]} \left\{ \frac{-q^{bm}(\theta)}{G_{\theta}(\sigma^{sb}(\theta)|\theta)} \right\}$. We consider a partially relaxed problem ($\mathcal{P}^{r'}$) where the participation constraint writes as (2.4.1), a pure state constraint, while incentive compatibility constraints (2.4.5) are still reduced to the Envelope Condition (2.6.2). The solution to this problem is now characterized.

PROPOSITION 2.6. *Suppose that Assumption 2.5 holds and again define θ^* as in (2.5.12). Then, the solution to the relaxed problem ($\mathcal{P}^{r'}$) entails the following properties.*

1. *There exists a probability measure μ whose support is $\Omega = \{\theta : U^{r'}(\theta) = 0\}$ such that*

$$S'(q^{r'}(\theta)) = \theta + \frac{F(\theta) - M(\theta)}{f(\theta)}, \quad \forall \theta \in \Theta. \quad (2.6.8)$$

where $M(\theta) = \int_{[\underline{\theta}, \theta]} \mu(d\tilde{\theta})$. In particular, we have

$$q^{r'}(\theta) \geq q^{bm}(\theta), \quad \forall \theta \in \Theta. \quad (2.6.9)$$

2. Money is burned following very bad news only

$$b^{r'}(\theta, \sigma) = \begin{cases} B & \text{if } \theta \in [\theta^*, \bar{\theta}] \text{ and } \sigma \in [\underline{\sigma}, \sigma^{r'}(\theta)], \\ 0 & \text{otherwise.} \end{cases} \quad (2.6.10)$$

where $\sigma^{r'}(\theta)$ is now defined as

$$1 + \frac{g_\theta(\sigma^{r'}(\theta)|\theta)}{g(\sigma^{r'}(\theta)|\theta)} \frac{F(\theta) - M(\theta)}{f(\theta)} = 0, \quad \forall \theta \in [\theta^*, \bar{\theta}]. \quad (2.6.11)$$

In particular, we have

$$\sigma^{r'}(\theta) \leq \sigma^{sb}(\theta) \text{ if } q^{r'}(\theta) \leq q^{fb}(\theta) \quad \forall \theta \in \Theta. \quad (2.6.12)$$

Moreover, it must be that

$$\min\{\Omega\} > \theta^* \quad (2.6.13)$$

and

$$q^{r'}(\theta) = -BG_\theta(\sigma^{r'}(\theta)|\theta), \quad \forall \theta \in \overset{\circ}{\Omega}. \quad (2.6.14)$$

The incentive constraints (2.4.5) hold provided that $q^{r'}$ is non-increasing and $\sigma^{r'}$ non-decreasing.

Proposition 2.6 accounts for the possibility that countervailing incentives are so strong that the agent's participation constraint binds on an interval with non-empty interior.¹⁵ The solution is then characterized by a probability measure over this set. On this set, money burning is optimally used to extract all the agent's information rent. Some distortions in output are still present, but these are lower compared to the Baron-Myerson solution. At the same time, money burning is less useful and limited in comparison with the scenario described in Proposition 2.5, since the rents of some mass of agent's types are already zero.

EXAMPLE 2.1. (continued) When condition (2.6.5) fails to hold, the type space $[\underline{\theta}, \bar{\theta}]$ is partitioned into three subsets under the optimal contract. On $[\underline{\theta}, \theta^*]$, output equals the Baron-Myerson level, no money is burned, and rents are positive. On (θ^*, θ') , output equals the Baron-Myerson level, some money is burned with $\sigma^{r'}(\theta) = \sigma^{sb}(\theta)$ but rents are still strictly positive. Finally, on $\Omega = [\theta', \bar{\theta}]$, output is constant at $q^{bm}(\theta')$, money is burnt with $\sigma^{r'}(\theta) \leq \sigma^{sb}(\theta)$, and rents are zero. The type $\theta' \in (\theta^*, \bar{\theta})$ is the smallest θ such that


$$-q^{bm}(\theta) - BG_\theta(\sigma^{sb}(\theta)|\theta) \geq 0,$$

which represents the value at which the rent would hit zero with the allocation described in Proposition 2.5 when (2.6.5) fails. As a consequence, there is output expansion on $\overset{\circ}{\Omega}$.


By combining (2.6.8) (2.6.11) and (2.6.14), we can find the cutoff $\sigma^{r'}(\theta)$ on $\overset{\circ}{\Omega}$:

$$S'(-BG_\theta(\sigma^{r'}(\theta)|\theta)) = \theta - \frac{g(\sigma^{r'}(\theta)|\theta)}{g_\theta(\sigma^{r'}(\theta)|\theta)}, \quad \forall \theta \in \overset{\circ}{\Omega}.$$


¹⁵If condition (2.6.5) holds, then the measure is such that $M(\theta) = 0$ for all $\theta \in \Theta$, meaning that the measure puts all the mass on the atom $\bar{\theta}$, which is the regular case.



(a) Output schemes. In blue (dotted), first-best output. In orange, Baron-Myerson scheme. In green (dashed), scheme with money burning.



(b) Rents. In blue (dashed), rents obtained by the agent in Baron-Myerson. In orange, rents with signals and money burning as defined in Proposition 2.6. There is a positive measure of types for which $U(\theta) = 0$.



(c) Money burning thresholds. In blue (dashed), threshold $\sigma^{sb}(\theta)$ as defined in 2.5.13. In orange, threshold $\sigma^{r'}(\theta)$ as defined in Proposition 2.6.

Figure 2.5: Optimal contract of Proposition 2.6.

Since $g(\sigma|\theta)$ is affine in θ in this example, g_θ and G_θ are constant in θ . It follows that the cutoff $\sigma^{r'}(\theta)$ is constant in θ on $\overset{\circ}{\Omega}$. From (2.6.14), it follows that $q^{r'}(\theta)$ is also constant on $\overset{\circ}{\Omega}$.

The measure is such that:

$$M(\theta) = \begin{cases} 0 & \text{if } \theta \leq \theta', \\ F(\theta) + f(\theta) \frac{g(\sigma^{r'}(\theta)|\theta)}{g_\theta(\sigma^{r'}(\theta)|\theta)} \in (0, 1) & \text{if } \theta' \leq \theta \leq \bar{\theta}, \end{cases}$$

In particular, the measure puts positive mass on the atom $\bar{\theta}$. A representation of the optimal contract can be found in Figure 2.5.

BIBLIOGRAPHY

- Bagnoli, M. and Bergstrom, T. (2005). Log-Concave Probability and Its Applications. *Economic Theory*, 26, 445–469.
- Baker, G., Gibbons, R., and Murphy, K. J. (1994). Subjective Performance Measures in Optimal Incentive Contracts. *The Quarterly Journal of Economics*, 109(4), 1125–1156.
- Baron, D. P. and Myerson, R. B. (1982). Regulating a Monopoly with Unknown Costs. *Econometrica*, 50(4), 911–930.
- Celik, G., Shin, D., and Strausz, R. (2021). Public Good Overprovision by a Manipulative Provider. *The RAND Journal of Economics*, 52(2), 314–333.
- Cella, M. (2008). Informed Principal with Correlation. *Games and Economic Behavior*, 64(2), 433–456.
- Cheng, C. (2021). Moral Hazard in Teams with Subjective Evaluations. *The RAND Journal of Economics*, 52(1), 22–48.
- Cr  mer, J. and McLean, R. P. (1985). Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent. *Econometrica*, 53(2), 345–361.
- Cr  mer, R. and McLean, R. P. (1988). Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions. *Econometrica*, 56(6), 1247–1257.
- Demougin, D. M. and Garvie, D. A. (1991). Contractual Design with Correlated Information under Limited Liability. *The RAND Journal of Economics*, 22(3), 477–489.
- Dequiedt, V. and Martimort, D. (2015). Vertical Contracting with Informational Opportunism. *American Economic Review*, 105, 2141–2182.
- Fuchs, W. (2007). Contracted with Repeated Moral Hazard and Private Evaluations. *American Economic Review*, 97, 1432–1448.
- Gary-Bobo, R. and Spiegel, Y. (2006). Optimal State-Contingent Regulation under Limited Liability. *The RAND Journal of Economics*, 37(2), 431–448.

- Holmström, B. (1979). Moral Hazard and Observability. *The Bell Journal of Economics*, 10(1), 74–91.
- Holmström, B. (1982). Moral Hazard in Teams. *The Bell Journal of Economics*, 13(2), 324–340.
- Jullien, B. (2000). Participation Constraints in Adverse Selection Models. *Journal of Economic Theory*, 93, 1–47.
- Kattwinkel, D. and Knoepfle, J. (2023). Costless Information and Costly Verification: A Case for Transparency. *Journal of Political Economy*, 131(2), 504–548.
- Kessler, A. S., Lülkesmann, C., and Schmitz, P. W. (2005). Endogenous Punishments in Agency with Verifiable Ex Post Information. *International Economic Review*, 46(4), 1207–1231.
- Khalil, F., Lawarrée, J., and Scott, T. J. (2015). Private Monitoring, Collusion and the Timing of Information. *The RAND Journal of Economics*, 46(4), 872–890.
- Krähmer, D. (2020). Information Disclosure and Full Surplus Extraction in Mechanism Design. *Journal of Economic Theory*, 187, 105020.
- Laffont, J. J. and Maskin, E. (1980). A Differential Approach to Dominant Strategy Mechanisms. *Econometrica*, 48(6), 1507–1520.
- Lang, M. (2023). Stochastic Contracts and Subjective Evaluations. *The RAND Journal of Economics*, 54(1), 104–134.
- Levin, J. (2003). Relational Incentive Contracts. *American Economic Review*, 93(3), 835–857.
- Lewis, T. and Sappington, D. (1989a). Countervailing Incentives in Agency Problems. *Journal of Economic Theory*, 49, 294–313.
- Lewis, T. and Sappington, D. (1989b). Inflexible Rules in Incentive Problems. *The American Economic Review*, 79, 69–84.
- Lopomo, G., Rigotti, L., and Shannon, C. (2022). Detectability, Duality, and Surplus Extraction. *Journal of Economic Theory*, 204, 105465.
- MacLeod, W. B. (2003). Optimal Contracting with Subjective Evaluation. *American Economic Review*, 93, 216–240.
- Martimort, D. and Stole, L. (2022). Participation Constraints in Discontinuous Adverse Selection Models. *Theoretical Economics*, 17, 1145–1181.
- McAfee, R. P. and Reny, P. (1992). Correlated Information and Mechanism Design. *Econometrica*, 60(2), 395–421.
- Milgrom, P. and Segal, I. (2002). Envelope Theorems for Arbitrary Choice Sets. *Econometrica*, 70, 583–601.
- Rahman, D. (2012a). Surplus Extraction on Arbitrary Type Spaces. Unpublished manuscript.

- Rahman, D. (2012b). But Who Will Monitor the Monitor? *American Economic Review*, 102(6), 2767–2797.
- Rajan, M. V. and Reichelstein, S. (2006). Subjective Performance Indicators and Discretionary Bonus Pools. *Journal of Accounting Research*, 44(3), 585–618.
- Riordan, M. and Sappington, D. (1988). Optimal Contracts with Public Ex Post Information. *Journal of Economic Theory*, 45, 189–199.
- Strausz, R. (2006). Interim Information in Long-Term Contracts. *Journal of Economics and Management Science*, 15, 1041–1067.
- Zábojník, J. (2014). Subjective Evaluations with Performance Feedback. *The RAND Journal of Economics*, 45(2), 341–369.

Appendices

2.A. MAIN PROOFS

PROOF OF PROPOSITION 2.1. Consider the payments as in (2.3.3) with, in addition,

$$z(\hat{\theta}, \sigma) = k(\hat{\theta}) \frac{g_{\theta}(\sigma|\hat{\theta})}{g(\sigma|\hat{\theta})} \quad \forall \hat{\theta} \in \Theta, \forall \sigma \in \Sigma \quad (2.A.1)$$

with a convenient function k to be found below.

Incentive compatibility constraints (2.4.2) are still valid when σ is verifiable. Those constraints thus rewrite as

$$U(\theta) = \max_{\hat{\theta} \in \Theta} S(q^{fb}(\hat{\theta})) - \theta q^{fb}(\hat{\theta}) - W^{fb}(\hat{\theta}) + k(\hat{\theta}) \mathbb{E}_{\sigma} \left(\frac{g_{\theta}(\sigma|\hat{\theta})}{g(\sigma|\hat{\theta})} | \theta \right), \quad \forall \theta \in \Theta. \quad (2.A.2)$$

By incentive compatibility plus the fact that $\mathbb{E}_{\sigma} \left(\frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} | \theta \right) = 0$, we check that there is full extraction of the agent's surplus when the agent tells the truth on his type, namely

$$U(\theta) = S(q^{fb}(\theta)) - \theta q^{fb}(\theta) - W^{fb}(\theta) = 0, \quad \forall \theta \in \Theta. \quad (2.A.3)$$

Observe that the maximand in (2.A.2) is absolutely continuous in θ when $\frac{g_{\theta}(\sigma|\hat{\theta})}{g(\sigma|\hat{\theta})}$ is so as we assumed. We can thus apply the Envelope Theorem in Milgrom and Segal (2002, Theorem 2), to show that U is absolutely continuous with a derivative (which is actually defined everywhere in the case under scrutiny) which is now worth

$$\dot{U}(\theta) = -q^{fb}(\theta) + k(\theta) \mathbb{E}_{\sigma} \left(\left(\frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \right)^2 | \theta \right), \quad \forall \theta \in \Theta. \quad (2.A.4)$$

Identifying this envelope condition with the derivative of (2.A.3), we must thus have

$$k(\theta) = \frac{q^{fb}(\theta)}{\mathbb{E}_{\sigma} \left(\left(\frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \right)^2 | \theta \right)}, \quad \forall \theta \in \Theta. \quad (2.A.5)$$

Taking together (2.A.2) and (2.A.3), global incentive compatibility holds when

$$(\theta - \hat{\theta})q^{fb}(\hat{\theta}) \geq k(\hat{\theta})\mathbb{E}_\sigma \left(\frac{g_\theta(\sigma|\hat{\theta})}{g(\sigma|\hat{\theta})} | \theta \right), \quad \forall (\theta, \hat{\theta}) \in \Theta^2$$

or, using (2.A.4),

$$(\theta - \hat{\theta})\mathbb{E}_\sigma \left(\left(\frac{g_\theta(\sigma|\hat{\theta})}{g(\sigma|\hat{\theta})} \right)^2 | \hat{\theta} \right) \geq \mathbb{E}_\sigma \left(\frac{g_\theta(\sigma|\hat{\theta})}{g(\sigma|\hat{\theta})} | \theta \right), \quad \forall (\theta, \hat{\theta}) \in \Theta^2. \quad (2.A.6)$$

Let us now define

$$\psi(\theta, \hat{\theta}) = \mathbb{E}_\sigma \left(\frac{g_\theta(\sigma|\hat{\theta})}{g(\sigma|\hat{\theta})} | \theta \right).$$

We may compute

$$\psi_\theta(\hat{\theta}, \hat{\theta}) = \mathbb{E}_\sigma \left(\left(\frac{g_\theta(\sigma|\hat{\theta})}{g(\sigma|\hat{\theta})} \right)^2 | \hat{\theta} \right), \quad \psi_{\theta\theta}(\theta, \hat{\theta}) = \mathbb{E}_\sigma \left(\frac{g_\theta(\sigma|\hat{\theta})g_{\theta\theta}(\sigma|\theta)}{g(\sigma|\hat{\theta})g(\sigma|\theta)} | \theta \right).$$

Integrating by parts yields

$$\psi_{\theta\theta}(\theta, \hat{\theta}) = - \int_{\underline{\sigma}}^{\bar{\sigma}} \frac{d}{d\sigma} \left(\frac{g_\theta(\sigma|\hat{\theta})}{g(\sigma|\hat{\theta})} \right) G_{\theta\theta}(\sigma|\theta) d\sigma.$$

When Assumptions 2.2 and 2.3 both hold, $\psi(\theta, \hat{\theta})$ is concave with respect to θ and thus

$$\psi(\theta, \hat{\theta}) \leq \psi(\hat{\theta}, \hat{\theta}) + \psi_\theta(\hat{\theta}, \hat{\theta})(\theta - \hat{\theta})$$

which, taking into account $\psi(\hat{\theta}, \hat{\theta}) = 0$, amounts to (2.A.6). $\square$

PROOF OF LEMMA 2.1. Observe that the maximand in (2.A.2) is absolutely continuous in θ for any $\hat{\theta}$ when g is twice continuously differentiable in θ as assumed. We can thus apply the Envelope Theorem in Milgrom and Segal (2002, Theorem 2), to show that U is absolutely continuous with a derivative, when it exists, which is defined as in (2.4.7). $\square$

PROOF OF PROPOSITION 2.2. We prove each item in turn.

1. Suppose that the first-best output profile $q^{fb}(\theta)$ is implementable without money burning. Then, using (2.4.5), there would exist a rent profile U and a function $H(\theta)$ such that

$$U(\theta) = \max_{\hat{\theta} \in \Theta} S(q^{fb}(\hat{\theta})) - \theta q^{fb}(\hat{\theta}) - H(\hat{\theta}), \quad \forall \theta \in \Theta. \quad (2.A.7)$$

We can again apply Lemma 2.1 to show that U is absolutely continuous with a derivative (almost everywhere) worth

$$\dot{U}(\theta) = -q^{fb}(\theta) \quad \forall \theta \in \Theta.$$

With full surplus extraction, we should have $U(\theta) = U^{fb}(\theta) = 0$ and thus $q^{fb}(\theta) = 0$ a.e.; a contradiction with the fact that output remains positive in (2.3.1).

2. Suppose that the first-best output allocation $(q^{fb}(\theta), U^{fb}(\theta))$ is implementable with money burning. Then, using (2.4.5), there would exist a function $H(\theta)$ such that

$$U^{fb}(\theta) = 0 = \max_{\hat{\theta} \in \Theta} S(q^{fb}(\hat{\theta})) - \theta q^{fb}(\hat{\theta}) - H(\hat{\theta}) - \mathbb{E}_\sigma \left(b(\hat{\theta}, \sigma) | \theta \right), \quad \forall \theta \in \Theta. \quad (2.A.8)$$

We can again apply the Envelope Theorem in Milgrom and Segal (2002, Theorem 2) to get

$$\dot{U}^{fb}(\theta) = 0 = -q^{fb}(\theta) - \mathbb{E}_\sigma \left(b(\theta, \sigma) \frac{g_\theta(\sigma | \theta)}{g(\sigma | \theta)} | \theta \right), \quad \forall \theta \in \Theta.$$

With full surplus extraction, we should thus have

$$-q^{fb}(\theta) - \mathbb{E}_\sigma \left(b(\theta, \sigma) \frac{g_\theta(\sigma | \theta)}{g(\sigma | \theta)} | \theta \right) = 0, \quad \forall \theta \in \Theta. \quad (2.A.9)$$

From (2.A.8), global incentive compatibility holds when

$$0 \geq S(q^{fb}(\hat{\theta})) - \theta q^{fb}(\hat{\theta}) - H(\hat{\theta}) - \mathbb{E}_\sigma \left(b(\hat{\theta}, \sigma) | \theta \right), \quad \forall (\theta, \hat{\theta}) \in \Theta^2$$

or,

$$(\theta - \hat{\theta})q^{fb}(\hat{\theta}) \geq \mathbb{E}_\sigma \left(b(\hat{\theta}, \sigma) | \hat{\theta} \right) - \mathbb{E}_\sigma \left(b(\hat{\theta}, \sigma) | \theta \right), \quad \forall (\theta, \hat{\theta}) \in \Theta^2. \quad (2.A.10)$$

Consider the principal's optimization problem when willing to implement the first-best allocation $(q^{fb}(\theta), U^{fb}(\theta) = 0)$. It consists in pointwise maximizing with respect to $b(\theta, \sigma) \in [0, B]$ the following objective

$$\mathbb{E}_\theta \left(S(q^{fb}(\theta)) - \theta q^{fb}(\hat{\theta}) - \mathbb{E}_\sigma \left(b(\hat{\theta}, \sigma) | \theta \right) \right)$$

subject to (2.A.9) only. Provided that (2.A.10) holds for the solution to this relaxed problem, we shall have find the optimal mechanism that extracts all surplus from the agent whatever his type and implements the first-best output.

Let denote by $\lambda(\theta)$ the Lagrange multiplier for (2.A.9). Because the corresponding Lagrangean is linear in $b(\theta, \sigma)$, the solution is bang-bang with

$$b(\theta, \sigma) \begin{cases} = B & \text{if } g(\sigma | \theta)f(\theta) + \lambda(\theta)g_\theta(\sigma | \theta) < 0, \\ \in [0, B] & \text{if } g(\sigma | \theta)f(\theta) + \lambda(\theta)g_\theta(\sigma | \theta) = 0, \\ = 0 & \text{if } g(\sigma | \theta)f(\theta) + \lambda(\theta)g_\theta(\sigma | \theta) > 0. \end{cases} \quad (2.A.11)$$

First, notice that $\lambda(\theta) = 0$ cannot be; because, otherwise, $b(\theta, \sigma) = 0$ for all $\sigma \in \Sigma$ and (2.A.9)

cannot hold. Second, suppose now that $\lambda(\theta) < 0$. Then, Assumption 2.2 would imply that $b(\theta, \sigma) = B$ for $\sigma \in [\sigma^{fb}(\theta), \bar{\sigma}]$ where $\sigma^{fb}(\theta)$ is defined as

$$\frac{g_\theta(\sigma^{fb}(\theta)|\theta)}{g(\sigma^{fb}(\theta)|\theta)} = -\frac{f(\theta)}{\lambda(\theta)}. \quad (2.A.12)$$

Moreover, $g_\theta(\sigma|\theta) > 0$ for all $\sigma \in (\sigma^{fb}(\theta), \bar{\sigma}]$. But then, the left-hand side of (2.A.9) is necessarily negative; a contradiction. Hence, $\lambda(\theta) > 0$. Then, Assumption 2.2 implies that $b(\theta, \sigma) = B$ for $\sigma \in [\underline{\sigma}, \sigma^{fb}(\theta)]$. Now, observe that we have $g_\theta(\sigma|\theta) < 0$ for all $\sigma \in [\underline{\sigma}, \sigma^{fb}(\theta)]$. Finally, $\sigma^{fb}(\theta)$ is defined as (2.5.2). In fact, Assumption 2.2 implies that $\int_{\underline{\sigma}}^{\sigma} g_\theta(\tilde{\sigma}|\theta) d\tilde{\sigma} = G_\theta(\sigma|\theta)$ is quasi-convex in σ and thus minimum at $\sigma_0(\theta)$ such that $g_\theta(\sigma_0(\theta)|\theta) = 0$. Condition (2.5.2) has thus a solution $\sigma^{fb}(\theta) \in [\underline{\sigma}, \sigma_0(\theta)]$ when B is large enough, namely

$$B \geq \max_{\theta \in \Theta} \left\{ \frac{-q^{fb}(\theta)}{G_\theta(\sigma_0(\theta)|\theta)} \right\}.$$

Let us check that (2.A.10) holds for the solution so obtained. First, notice that H is defined as

$$H(\theta) = S(q^{fb}(\theta)) - \theta q^{fb}(\theta) - BG_\theta(\sigma^{fb}(\theta)|\theta), \quad \forall \theta \in \Theta. \quad (2.A.13)$$

Second, global incentive compatibility amounts to checking that

$$(\theta - \hat{\theta})q^{fb}(\hat{\theta}) \geq B \left(\int_{\underline{\sigma}}^{\sigma^{fb}(\hat{\theta})} (g(\sigma|\hat{\theta}) - g(\sigma|\theta)) d\sigma \right), \quad \forall (\theta, \hat{\theta}) \in \Theta^2. \quad (2.A.14)$$

Using

$$-q^{fb}(\theta) = B \int_{\underline{\sigma}}^{\sigma^{fb}(\theta)} g_\theta(\sigma|\theta) d\sigma = BG_\theta(\sigma^{fb}(\theta)|\theta),$$

and inserting into (2.A.14) allows us to rewrite this condition as

$$(\hat{\theta} - \theta)G_\theta(\sigma^{fb}(\hat{\theta})|\hat{\theta}) \geq G(\sigma^{fb}(\hat{\theta})|\hat{\theta}) - G(\sigma^{fb}(\hat{\theta})|\theta), \quad \forall (\theta, \hat{\theta}) \in \Theta^2. \quad (2.A.15)$$

This condition holds under Assumption 2.3.

□

PROOF OF PROPOSITION 2.3. When the incentive problem is regular, we can recover from the binding participation constraint for the worst type $\bar{\theta}$ (2.5.3) and the Envelope Condition (2.4.7), the expression of the rent profile $U(\theta)$ in terms of the output allocation as (2.5.8). From there, the fee $H(\theta)$ is obtained as

$$U(\theta) = \mathbb{E}_\sigma (S(q(\theta, \sigma)) - \theta q(\theta, \sigma) - b(\theta, \sigma)|\theta) - H(\theta). \quad (2.A.16)$$

Inserting (2.5.8) into the maximand ($\mathcal{P}^r$) and integrating by parts yields that $q^{sb}(\theta, \sigma)$ and $b(\theta, \sigma)$

should now be pointwise-optimizers for

$$\mathbb{E}_{\theta, \sigma} \left(\left(1 + \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta)}{f(\theta)} \right) (S(q(\theta, \sigma)) - \theta q(\theta, \sigma) - b(\theta, \sigma)) - \frac{F(\theta)}{f(\theta)} q(\theta, \sigma) \right). \quad (2.A.17)$$

From there, the following consequences are immediately drawn.

1. Observe that Condition (2.5.5) implies that the coefficient of the surplus $S(q(\theta, \sigma)) - \theta q(\theta, \sigma)$ in this maximand remains positive. Hence, the optimization problem remains strictly concave. Pointwise optimization of this objective then yields the expression of the second-best output in (2.5.6). Inada conditions on S , as assumed, imply that $q^{sb}(\theta, \sigma)$ remains positive.
2. Because Condition (2.5.5) holds and the objective is linear in $b(\theta, \sigma)$, (2.5.7) immediately follows.

MONOTONICITY. Condition (2.5.4) implies that $q^{sb}(\theta, \sigma)$ is non-increasing in θ . That $q^{sb}(\theta, \sigma)$ is non-decreasing in σ follows from Assumption 2.2.

REGULARITY. Regularity follows from checking that the necessary condition (2.4.7) suffices for (2.4.5) while the participation constraint (2.4.1) follows from (2.5.3). We check each item in turn.

1. We rewrite the incentive compatibility conditions (2.4.5) using (2.5.7) as

$$U^{sb}(\theta) \geq \mathbb{E}_{\sigma} \left(S(q^{sb}(\hat{\theta}, \sigma)) - \theta q^{sb}(\hat{\theta}, \sigma) | \theta \right) - H^{sb}(\hat{\theta})$$

or

$$U^{sb}(\theta) - U^{sb}(\hat{\theta}) \geq \mathbb{E}_{\sigma} \left(S(q^{sb}(\hat{\theta}, \sigma)) - \theta q^{sb}(\hat{\theta}, \sigma) | \theta \right) - \mathbb{E}_{\sigma} \left(S(q^{sb}(\hat{\theta}, \sigma)) - \hat{\theta} q^{sb}(\hat{\theta}, \sigma) | \hat{\theta} \right), \forall (\theta, \hat{\theta}) \in \Theta^2.$$

Using the necessary condition (2.4.7), we rewrite the left-hand side above and get

$$\begin{aligned} & \int_{\hat{\theta}}^{\theta} \mathbb{E}_{\sigma} \left(-q^{sb}(\tilde{\theta}, \sigma) + (S(q^{sb}(\tilde{\theta}, \sigma)) - \tilde{\theta} q^{sb}(\tilde{\theta}, \sigma)) \frac{g_{\theta}(\sigma|\tilde{\theta})}{g(\sigma|\tilde{\theta})} | \tilde{\theta} \right) d\tilde{\theta} \\ & \geq \mathbb{E}_{\sigma} \left(S(q^{sb}(\hat{\theta}, \sigma)) - \theta q^{sb}(\hat{\theta}, \sigma) | \theta \right) - \mathbb{E}_{\sigma} \left((S(q^{sb}(\hat{\theta}, \sigma)) - \hat{\theta} q^{sb}(\hat{\theta}, \sigma)) | \hat{\theta} \right), \quad \forall (\theta, \hat{\theta}) \in \Theta^2. \end{aligned} \quad (2.A.18)$$

Define now ψ as

$$\psi(\theta, \hat{\theta}) = \mathbb{E}_{\sigma} \left(S(q^{sb}(\hat{\theta}, \sigma)) - \theta q^{sb}(\hat{\theta}, \sigma) | \theta \right). \quad (2.A.19)$$

We rewrite (2.A.18) as

$$\int_{\hat{\theta}}^{\theta} \psi_{\theta}(\tilde{\theta}, \hat{\theta}) d\tilde{\theta} \geq \psi(\hat{\theta}, \hat{\theta}) - \psi(\theta, \hat{\theta}) \quad \forall (\theta, \hat{\theta}) \in \Theta^2. \quad (2.A.20)$$

Condition (2.A.20) amounts to

$$\int_{\hat{\theta}}^{\theta} [\psi_{\theta}(\tilde{\theta}, \hat{\theta}) - \psi_{\theta}(\hat{\theta}, \hat{\theta})] d\tilde{\theta} \geq 0$$

$$\int_{\hat{\theta}}^{\theta} \int_{\hat{\theta}}^{\tilde{\theta}} \psi_{\theta\hat{\theta}}(\tilde{\theta}, x) dx d\tilde{\theta} \geq 0$$

It immediately follows that (2.A.18) is satisfied when

$$\psi_{\theta\hat{\theta}}(\theta, \hat{\theta}) \geq 0 \quad \forall (\theta, \hat{\theta}) \in \Theta^2. \quad (2.A.21)$$

Using (2.A.19), we compute

$$\psi_{\theta\hat{\theta}}(\theta, \hat{\theta}) = \int_{\underline{\sigma}}^{\bar{\sigma}} q_{\hat{\theta}}^{sb}(\hat{\theta}, \sigma) \left(-1 + (S'(q^{sb}(\hat{\theta}, \sigma)) - \theta) \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \right) g(\sigma|\theta) d\sigma. \quad (2.A.22)$$

Using (2.5.6), observe that

$$-1 + (S'(q^{sb}(\theta, \sigma)) - \theta) \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} = -\frac{1}{1 + \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta)}{f(\theta)}} < 0$$

Taken this with the monotonicity condition $q_{\hat{\theta}}^{sb} < 0$, we get $\psi_{\theta\hat{\theta}}(\theta, \theta) > 0$ and so (2.A.21) holds when $\Delta = \bar{\theta} - \underline{\theta}$ is small enough.

2. We now demonstrate that $U^{sb}(\theta)$ is decreasing and thus the conjecture that (2.4.1) can be replaced by (2.5.3) is correct. To this end, we compute

$$\begin{aligned} \mathbb{E}_{\sigma} \left((S(q^{sb}(\theta, \sigma)) - \theta q^{sb}(\theta, \sigma)) \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} | \theta \right) &= \int_{\underline{\sigma}}^{\bar{\sigma}} (S(q^{sb}(\theta, \sigma)) - \theta q^{sb}(\theta, \sigma)) g_{\theta}(\sigma|\theta) d\sigma \\ &= [S(q^{sb}(\theta, \sigma)) - \theta q^{sb}(\theta, \sigma) G_{\theta}(\sigma|\theta)]_{\underline{\sigma}}^{\bar{\sigma}} - \int_{\underline{\sigma}}^{\bar{\sigma}} (S'(q^{sb}(\theta, \sigma)) - \theta) q_{\sigma}^{sb}(\theta, \sigma) G_{\theta}(\sigma|\theta) d\sigma. \end{aligned}$$

Using $G_{\theta}(\underline{\sigma}|\theta) = G_{\theta}(\bar{\sigma}|\theta) = 0$ and (2.5.6), we rewrite the right-hand side as

$$- \int_{\underline{\sigma}}^{\bar{\sigma}} \frac{\frac{F(\theta)}{f(\theta)}}{1 + \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta)}{f(\theta)}} q_{\sigma}^{sb}(\theta, \sigma) G_{\theta}(\sigma|\theta) d\sigma.$$

The right-hand side of (2.4.7) thus becomes

$$- \int_{\underline{\sigma}}^{\bar{\sigma}} q^{sb}(\theta, \sigma) g(\sigma|\theta) \left(\frac{1}{1 + \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta)}{f(\theta)}} + \frac{\frac{G_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \left(\frac{F(\theta)}{f(\theta)} \right)^2}{\left(1 + \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta)}{f(\theta)} \right)^2} \frac{d}{d\sigma} \left(\frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} \right) \right) d\sigma. \quad (2.A.23)$$

Hence, U^{sb} is indeed certainly decreasing when Condition (2.5.9) holds.

□

PROOF OF PROPOSITION 2.4. Remember from the Proof of Proposition 2.3 that, when looking for a solution to $(\mathcal{P}^r)$, the principal must end up maximizing pointwise the expression (2.A.17). Suppose

that (2.5.5) fails at some pair (θ, σ) . From Assumption 2.2, this condition also fails for all $\sigma' \leq \sigma$. In other words,

$$1 + \frac{g_\theta(\sigma'|\theta)}{g(\sigma'|\theta)} \frac{F(\theta)}{f(\theta)} < 0, \quad \forall \sigma' \in [\underline{\sigma}, \sigma^{sb}(\theta)] \quad (2.A.24)$$

where $\sigma^{sb}(\theta) > \underline{\sigma}$ is uniquely defined (again from Assumption 2.2) by (2.5.13). when

$$1 + \frac{g_\theta(\underline{\sigma}|\theta)}{g(\underline{\sigma}|\theta)} \frac{F(\theta)}{f(\theta)} < 0. \quad (2.A.25)$$

Moreover, this latter condition itself requires $\theta \geq \theta^*$ for some $\theta^* > \underline{\theta}$ that solves (2.5.12), again from Assumption 2.5.

Pointwise optimization for (2.A.17) with respect to $b(\theta, \sigma)$ then leads to (2.5.15). The optimality condition (2.5.14) is also straightforward. $\square$

PROOF OF PROPOSITION 2.5. We first consider the relaxed problem $(\mathcal{P}^r)$. Following previous steps, the principal must end up maximizing pointwise the following expression (where we take into account the fact that output is no longer contingent on signal):

$$\mathbb{E}_\theta \left(S(q(\theta)) - \left(\theta + \frac{F(\theta)}{f(\theta)} \right) q(\theta) - \mathbb{E}_\sigma \left(b(\theta, \sigma) \left(1 + \frac{F(\theta)}{f(\theta)} \frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)} \right) | \theta \right) \right). \quad (2.A.26)$$

It immediately follows that the optimal output is given by (2.6.3) while, thanks to the bang-bang nature of the problem, the optimal level of money burned is characterized in (2.6.4).

REGULARITY. Given the structure of the contract laid out in (2.6.3) and (2.6.4), we may rewrite the incentive compatibility conditions (2.4.5) as

$$\begin{aligned} U^{sb}(\theta) &= S(q^{bm}(\theta)) - \theta q^{bm}(\theta) - BG(\sigma^{sb}(\theta)|\theta) - H^{bm}(\theta) \\ &\geq S(q^{bm}(\hat{\theta})) - \theta q^{bm}(\hat{\theta}) - BG(\sigma^{sb}(\hat{\theta})|\theta) - H^{bm}(\hat{\theta}) \end{aligned}$$

or

$$U^{sb}(\theta) - U^{sb}(\hat{\theta}) \geq (\hat{\theta} - \theta) q^{bm}(\hat{\theta}) - B(G(\sigma^{sb}(\hat{\theta})|\theta) - G(\sigma^{sb}(\hat{\theta})|\hat{\theta})), \quad \forall (\theta, \hat{\theta}) \in \Theta^2.$$

Using the necessary condition (2.6.2), we rewrite the right-hand side above and get

$$\int_{\theta}^{\hat{\theta}} (q^{bm}(\tilde{\theta}) + BG_\theta(\sigma^{sb}(\tilde{\theta})|\tilde{\theta})) d\tilde{\theta} \geq (\hat{\theta} - \theta) q^{bm}(\hat{\theta}) - B(G(\sigma^{sb}(\hat{\theta})|\theta) - G(\sigma^{sb}(\hat{\theta})|\hat{\theta})), \quad \forall (\theta, \hat{\theta}) \in \Theta^2. \quad (2.A.27)$$

Note that q^{bm} is non-increasing and thus

$$\int_{\theta}^{\hat{\theta}} q^{bm}(\tilde{\theta}) d\tilde{\theta} \geq (\hat{\theta} - \theta) q^{bm}(\hat{\theta}), \quad \forall (\theta, \hat{\theta}) \in \Theta^2. \quad (2.A.28)$$

A sufficient condition for (2.A.27) to hold is thus:

$$\int_{\theta}^{\hat{\theta}} G_{\theta}(\sigma^{sb}(\tilde{\theta})|\tilde{\theta})d\tilde{\theta} \geq G(\sigma^{sb}(\hat{\theta})|\hat{\theta}) - G(\sigma^{sb}(\hat{\theta})|\theta), \quad \forall(\theta, \hat{\theta}) \in \Theta^2. \quad (2.A.29)$$

Condition (2.A.29) amounts to

$$\begin{aligned} \int_{\theta}^{\hat{\theta}} [G_{\theta}(\sigma^{sb}(\tilde{\theta})|\tilde{\theta}) - G_{\theta}(\sigma^{sb}(\hat{\theta})|\tilde{\theta})]d\tilde{\theta} &\geq 0 \\ \int_{\theta}^{\hat{\theta}} \int_{\sigma^{sb}(\hat{\theta})}^{\sigma^{sb}(\tilde{\theta})} g_{\theta}(\sigma|\tilde{\theta})d\sigma d\tilde{\theta} &\geq 0 \end{aligned}$$

It immediately follows that (2.A.29) is satisfied when σ^{sb} is non-decreasing and $g_{\theta}(\sigma|\theta) \leq 0$ for $\sigma \in [\sigma^{sb}(\hat{\theta}), \sigma^{sb}(\theta)]$ if $\hat{\theta} \geq \theta$ (or $\sigma \in [\sigma^{sb}(\theta), \sigma^{sb}(\hat{\theta})]$ if $\theta \geq \hat{\theta}$). This last condition holds since Assumption 2.2 with the definition of $\sigma^*(\theta)$ in (2.5.13) altogether imply that

$$\frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} < -\frac{f(\theta)}{F(\theta)} < 0, \quad \forall \theta \in [\theta^*, \bar{\theta}], \forall \sigma \leq \sigma^{sb}(\theta). \quad (2.A.30)$$

It remains to prove that σ^{sb} is non-decreasing, let us first differentiate (2.5.13) to get

$$\frac{\partial}{\partial \sigma} \left(\frac{g_{\theta}(\sigma^{sb}(\theta)|\theta)}{g(\sigma^{sb}(\theta)|\theta)} \right) \dot{\sigma}^{sb}(\theta) = -\frac{d}{d\theta} \left(\frac{f(\theta)}{F(\theta)} \right) - \frac{\partial}{\partial \theta} \left(\frac{g_{\theta}(\sigma^{sb}(\theta)|\theta)}{g(\sigma^{sb}(\theta)|\theta)} \right). \quad (2.A.31)$$

Observe that the coefficient on the left-hand side is non-negative from Assumption 2.2 while the right-hand side, is non-negative from Assumption 2.5 and (2.5.13).¹⁶ Hence, $\dot{\sigma}^{sb}(\theta) \geq 0$ as requested.

We now demonstrate that $U^{sb}(\theta)$ is decreasing and thus the conjecture that (2.4.1) can be replaced by (2.5.3) is correct. It is true when

$$q^{bm}(\theta) \geq -BG_{\theta}(\sigma^{sb}(\theta)|\theta), \quad \forall \theta \in \Theta;$$

a condition that follows from (2.6.5).

¹⁶Indeed, we have

$$-\frac{d}{d\theta} \left(\frac{f(\theta)}{F(\theta)} \right) = \frac{\frac{d}{d\theta} \left(\frac{F(\theta)}{f(\theta)} \right)}{\left(\frac{F(\theta)}{f(\theta)} \right)^2}$$

and

$$-\frac{\partial}{\partial \theta} \left(\frac{g_{\theta}(\sigma^{sb}(\theta)|\theta)}{g(\sigma^{sb}(\theta)|\theta)} \right) = -\frac{\partial}{\partial \theta} \left(\frac{1}{\frac{g(\sigma^{sb}(\theta)|\theta)}{g_{\theta}(\sigma^{sb}(\theta)|\theta)}} \right) = \frac{\frac{\partial}{\partial \theta} \left(\frac{g(\sigma^{sb}(\theta)|\theta)}{g_{\theta}(\sigma^{sb}(\theta)|\theta)} \right)}{\left(\frac{g(\sigma^{sb}(\theta)|\theta)}{g_{\theta}(\sigma^{sb}(\theta)|\theta)} \right)^2}.$$

From (2.5.13), it follows that $\left(\frac{g(\sigma^{sb}(\theta)|\theta)}{g_{\theta}(\sigma^{sb}(\theta)|\theta)} \right)^2 = \left(\frac{F(\theta)}{f(\theta)} \right)^2$, and thus:

$$-\frac{d}{d\theta} \left(\frac{f(\theta)}{F(\theta)} \right) - \frac{\partial}{\partial \theta} \left(\frac{g_{\theta}(\sigma^{sb}(\theta)|\theta)}{g(\sigma^{sb}(\theta)|\theta)} \right) = \frac{\frac{\partial}{\partial \theta} \left(\frac{F(\theta)}{f(\theta)} + \frac{g(\sigma^{sb}(\theta)|\theta)}{g_{\theta}(\sigma^{sb}(\theta)|\theta)} \right)}{\left(\frac{F(\theta)}{f(\theta)} \right)^2} \geq 0,$$

which is non-negative by Assumption 2.5.

□

PROOF OF PROPOSITION 2.6. We use Jullien (2000)'s and Martimort and Stole (2022)'s techniques to characterize the solution (denoted as with a superscript r') to $(\mathcal{P}^{r'})$. We thus associate to the participation constraint (2.4.1), a probability measure μ defined over the Borel subsets of Θ with an associated adjoint function, $M : \Theta \rightarrow [0, 1]$, defined by $M(\underline{\theta}) = 0$, $M(\bar{\theta}^+) = 1$ and

$$M(\theta) = \int_{[\underline{\theta}, \theta)} \mu(d\tilde{\theta}) \quad \text{for } \theta > \underline{\theta}.$$

Necessary and sufficient conditions (thanks to strict concavity of the maximand) are then stated in terms of this probability measure which serves to first express a complementary slackness condition:

$$\int_{\underline{\theta}}^{\bar{\theta}} U^{r'}(\tilde{\theta}) \mu(d\tilde{\theta}) = 0. \quad (2.A.32)$$

This condition states that the probability measure μ has support $\Omega = \{\theta : U^{r'}(\theta) = 0\}$. When $\bar{\Omega} \neq \emptyset$, we may differentiate $U^{r'}(\theta) = 0$ with respect to θ and, using (2.6.2), get (2.6.14).

Second, the optimality conditions write as

$$(q^{r'}(\theta), b^{r'}(\theta, \sigma)) \in \arg \max_{(q, b(\cdot))} \left\{ f(\theta) (S(q) - \theta q - \mathbb{E}_{\sigma} (b(\sigma)|\theta)) - (F(\theta) - M(\theta)) \left(q + \mathbb{E}_{\sigma} \left(b(\sigma) \frac{g_{\theta}(\sigma|\theta)}{g(\sigma|\theta)} | \theta \right) \right) \right\} \quad \text{for a.e. } \theta \in \Theta. \quad (2.A.33)$$

The first-order optimality condition with respect to output q is given by (2.6.8). Because $M(\theta) \geq 0$, it immediately follows that (2.6.9) holds.

Given the linearity of the objective function on the right-hand side of (2.A.33) in $b(\cdot)$, the optimality condition with respect to the amount of money burned is given by then bang-bang solution (2.6.10). Observe that (2.6.14) requires $\sigma^{r'}(\theta) > \underline{\sigma}$ for all $\theta \in \bar{\Omega} \neq \emptyset$. Thus, (2.6.13) must necessarily hold.

Finally, when $q^{r'}(\theta) \leq q^{fb}(\theta)$ implies $M(\theta) \leq F(\theta)$. Hence, (2.6.12) implies

$$\frac{g_{\theta}(\sigma^{r'}(\theta)|\theta)}{g(\sigma^{r'}(\theta)|\theta)} = -\frac{f(\theta)}{F(\theta) - M(\theta)} < -\frac{f(\theta)}{F(\theta)} = \frac{g_{\theta}(\sigma^{sb}(\theta)|\theta)}{g(\sigma^{sb}(\theta)|\theta)}, \quad \forall \theta \in [\theta^*, \bar{\theta}]. \quad (2.A.34)$$

Assumption 2.2 then implies (2.6.12).

INCENTIVE COMPATIBILITY. The proof of that (2.4.5) hold provided that $q^{r'}$ is non-increasing and $\sigma^{r'}$ non-decreasing is similar to that made in the Proof of Proposition 2.5 and is thus omitted. □

2.B. A REVELATION PRINCIPLE

We provide here a proof of the Revelation Principle that is used in the main text. It follows quite closely the one formulated in Appendix B in [Dequiedt and Martimort \(2015\)](#). A general mechanism in this case consists in a message spaces $\mathcal{M}_A$ and $\mathcal{M}_P$ respectively for the agent and the principal, and an allocation rule x that maps the set $\mathcal{M}_A \times \mathcal{M}_P$ into the set $\Delta(\mathcal{Q} \times \mathcal{T} \times \mathcal{B})$ of probability distributions over $\mathcal{Q} \times \mathcal{T} \times \mathcal{B}$, where we remind the reader that $\mathcal{Q}$ is the set of possible quantities and $\mathcal{T}$ is the set of possible transfers. The agent's strategy maps Θ into the set $\Delta(\mathcal{M}_A)$ of probability distributions over $\mathcal{M}_A$. The principal's strategy maps $\Sigma \times \mathcal{M}_a$ into the set $\Delta(\mathcal{M}_P)$ of probability distributions over $\mathcal{M}_P$. Fixing the mechanism allocation x , a message pair (m_A, m_P) induces the following utilities for principal and agent respectively:

$$\begin{aligned}\mathcal{V}(x(m_A, m_P)) &= \int_{\mathcal{Q} \times \mathcal{T} \times \mathcal{B}} v(q, t + b) dx(q, t, b | m_A, m_P) \\ \mathcal{U}(x(m_A, m_P); \theta) &= \int_{\mathcal{Q} \times \mathcal{T} \times \mathcal{B}} u(q, t; \theta) dx(q, t, b | m_A, m_P)\end{aligned}$$

The timing is the following:

1. The principal proposes a mechanism $\{\mathcal{M}_A, \mathcal{M}_P, x\}$
2. The agent and the principal respectively learn θ and σ .
3. The agent decides whether to accept the mechanism. In the affirmative case, the agent sends a message $m_A \in \mathcal{M}_A$. Otherwise, both get their reservation values that are normalised to zero.
4. The principal observes the message m_a , and then send a message $m_P \in \mathcal{M}_P$.
5. The allocation rule x inputs (m_A, m_P) and outputs $(q, t, b) \in \mathcal{Q} \times \mathcal{T} \times \mathcal{B}$.
6. Payoffs are realised, and the interaction ends.

The appropriate equilibrium concept for the induced game after acceptance is that of Perfect Bayesian Equilibrium. Thus, after observing σ and m_A , the principal updates her beliefs about θ according to Bayes' rule whenever applicable. Similarly, the agent updates his beliefs about σ after observing θ . The message strategies must be optimal for each player:

$$\begin{aligned}\forall \sigma \in \Sigma : \forall m_A \in \mathcal{M}_A : \forall m_P \in \text{supp}(m_P^*(m_A, \sigma)) : m_P \in \arg \max_{\hat{m}_P \in \mathcal{M}_P} \mathbb{E}_\theta (\mathcal{V}(x(m_A, \hat{m}_P)) | m_A, \sigma) \\ = \arg \max_{\hat{m}_P \in \mathcal{M}_P} \mathcal{V}(x(m_A, \hat{m}_P))\end{aligned}\tag{2.B.1}$$

$$\forall \theta \in \Theta : \forall m_A \in \text{supp}(m_A^*(\theta)) : m_A \in \arg \max_{\hat{m}_A \in \mathcal{M}_A} \mathbb{E}_\sigma \left(\int_{\mathcal{M}_P} \mathcal{U}(x(\hat{m}_A, m_P); \theta) dm_P^*(m_P | \hat{m}_A, \sigma) \middle| \theta \right)\tag{2.B.2}$$

Where the second line follows because we are in a private values context in which θ does not enter principal's utility. Moreover, because we are in the continuation game after acceptance, we must have:

$$\forall \theta \in \Theta : \forall m_A \in \text{supp}(m_A^*(\theta)) : \mathbb{E}_\sigma \left(\int_{\mathcal{M}_P} \mathcal{U}(x(m_A, m_P); \theta) dm_P^*(m_P | \hat{m}_A, \sigma) \mid \theta \right) \geq 0 \quad (2.B.3)$$

We now consider the direct revelation mechanism $\{\Theta, \Sigma, \tilde{x}\}$, where Θ is the set of reports of the agent, Σ is the set of reports of the principal, and $\tilde{x}$ is the allocation rule defined as:

$$\forall \hat{\theta} \in \Theta : \forall \hat{\sigma} \in \Sigma : \tilde{x}(\hat{\theta}, \hat{\sigma}) = x(m_A^*(\hat{\theta}), m_P^*(m_A^*(\hat{\theta}), \hat{\sigma})) \quad (2.B.4)$$

PROPOSITION 2.7. *Any allocation rule achieved in a continuation equilibrium (m_A^*, m_P^*) with the offer and acceptance of $(\mathcal{M}_A, \mathcal{M}_P, x)$ can be replicated with a direct revelation mechanism $\{\Theta, \Sigma, \tilde{x}\}$ satisfying (2.B.4), acceptance and truthful reporting:*

$$\forall \theta \in \Theta : \mathbb{E}_\sigma (\mathcal{U}(\tilde{x}(\theta, \sigma); \theta) \mid \theta) \geq 0 \quad (2.B.5)$$

$$\forall \theta \in \Theta : \mathbb{E}_\sigma (\mathcal{U}(\tilde{x}(\theta, \sigma); \theta) \mid \theta) = \max_{\hat{\theta} \in \Theta} \mathbb{E}_\sigma \left(\mathcal{U}(\tilde{x}(\hat{\theta}, \sigma); \theta) \mid \theta \right) \quad (2.B.6)$$

$$\forall \theta \in \Theta : \forall \sigma \in \Sigma : \mathcal{V}(\tilde{x}(\theta, \sigma)) = \max_{\hat{\sigma} \in \Sigma} \mathcal{V}(\tilde{x}(\theta, \hat{\sigma})) \quad (2.B.7)$$

Proof. We use the following notations:

$$\begin{aligned} \mathcal{U}(x(m_A^*(\theta), m_P^*(m_A^*(\theta), \sigma)); \theta) &= \int_{\mathcal{M}_A} \int_{\mathcal{M}_P} \mathcal{U}(x(m_A, m_P); \theta) dm_P^*(m_P | m_A, \sigma) dm_A^*(m_A | \theta) \\ \mathcal{V}(x(m_A^*(\theta), m_P^*(m_A^*(\theta), \sigma))) &= \int_{\mathcal{M}_A} \int_{\mathcal{M}_P} \mathcal{V}(x(m_A, m_P)) dm_P^*(m_P | m_A, \sigma) dm_A^*(m_A | \theta) \end{aligned}$$

Condition (2.B.5) follows straightforwardly from the combination of (2.B.3) and (2.B.4):

$$\mathbb{E}_\sigma (\mathcal{U}(\tilde{x}(\theta, \sigma); \theta) \mid \theta) \underbrace{=}_{\text{by (2.B.4)}} \mathbb{E}_\sigma (\mathcal{U}(x(m_A^*(\theta), m_P^*(m_A^*(\theta), \sigma)); \theta) \mid \theta) \underbrace{\geq}_{\text{by (2.B.3)}} 0$$

Condition (2.B.6) is implied by (2.B.2) and (2.B.4):

$$\begin{aligned} \mathbb{E}_\sigma (\mathcal{U}(\tilde{x}(\theta, \sigma); \theta) \mid \theta) &\underbrace{=}_{\text{by (2.B.4)}} \mathbb{E}_\sigma (\mathcal{U}(x(m_A^*(\theta), m_P^*(m_A^*(\theta), \sigma)); \theta) \mid \theta) \\ &\underbrace{\geq}_{\text{by (2.B.2)}} \mathbb{E}_\sigma \left(\mathcal{U}(x(m_A^*(\hat{\theta}), m_P^*(m_A^*(\hat{\theta}), \sigma)); \theta) \mid \theta \right) \\ &\underbrace{=}_{\text{by (2.B.4)}} \mathbb{E}_\sigma \left(\mathcal{U}(\tilde{x}(\hat{\theta}, \sigma); \theta) \mid \theta \right) \end{aligned}$$

where we took $\hat{m}_A \in \text{supp}(m_A^*(\hat{\theta}))$ in the second line. Finally, condition (2.B.7) is implied by

conditions (2.B.1) and (2.B.4):

$$\begin{aligned}
\mathcal{V}(\tilde{x}(\theta, \sigma)) &\stackrel{\text{by (2.B.4)}}{=} \mathcal{V}(x(m_A^*(\theta), m_P^*(m_A^*(\theta), \sigma))) \\
&\stackrel{\text{by (2.B.1)}}{\geq} \mathcal{V}(x(m_A^*(\theta), m_P^*(m_A^*(\theta), \hat{\sigma}))) \\
&\stackrel{\text{by (2.B.4)}}{=} \mathcal{V}(\tilde{x}(\theta, \hat{\sigma}))
\end{aligned}$$

where we took $\hat{m}_P \in \text{supp}(m_P(m_A^*(\theta), \hat{\sigma}))$ in the second line. □

2.C. PARTICIPATION CONSTRAINTS

Propositions 2.3 and 2.4 are valid provided that $\dot{U}(\theta) < 0$ for all $\theta < \bar{\theta}$, and that global IC constraints are all met. However, the informativeness of the signal may be strong enough to bring the rents of some positive mass of types to zero. Thus, in a similar fashion to that in Proposition 2.6, we consider a partially relaxed problem where the participation constraint writes as (2.4.1), a pure state constraint, while incentive compatibility constraints (2.4.5) are still reduced to the Envelope Condition (2.4.7). The solution to this problem is now characterized.

PROPOSITION 2.8. *The solution to the relaxed problem $(\mathcal{P}^{r'})$ entails the following properties.*

1. *There exists a probability measure μ whose support is $\Omega = \{\theta : U^{r'}(\theta) = 0\}$ such that*

$$\begin{cases} q^{r'}(\theta, \sigma) = 0 & \text{if } 1 + \frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta) - M(\theta)}{f(\theta)} \leq 0 \text{ and } F(\theta) \geq M(\theta), \\ S'(q^{r'}(\theta)) = \theta + \frac{\frac{F(\theta) - M(\theta)}{f(\theta)}}{1 + \frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta) - M(\theta)}{f(\theta)}}, & \text{otherwise.} \end{cases} \quad (2.C.1)$$

where $M(\theta) = \int_{[\underline{\theta}, \theta)} \mu(d\tilde{\theta})$.

2. *Money is burned following very bad news only*

$$b^{r'}(\theta, \sigma) = \begin{cases} B & \text{if } 1 + \frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)} \frac{F(\theta) - M(\theta)}{f(\theta)} \leq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (2.C.2)$$

3. *And*

$$\mathbb{E}_\sigma(q^{r'}(\theta, \sigma)|\theta) = \mathbb{E}_\sigma\left((S(q^{r'}(\theta, \sigma)) - \theta q^{r'}(\theta, \sigma) - b^{r'}(\theta, \sigma)) \frac{g_\theta(\sigma|\theta)}{g(\sigma|\theta)} \middle| \theta\right), \quad \forall \theta \in \overset{\circ}{\Omega}. \quad (2.C.3)$$

PROOF OF PROPOSITION 2.8. The proof replicates that of Proposition 2.6, and is therefore omitted.

The main difference stems from the objective function of the Principal's problem:

$$(q^{r'}(\theta, \sigma), b^{r'}(\theta, \sigma)) \in \arg \max_{(q, b)} \left\{ f(\theta) g(\sigma | \theta) (S(q) - \theta q - b) \right. \\ \left. + [F(\theta) - M(\theta)] [-qg(\sigma | \theta) + g_\theta(\sigma | \theta) (S(q) - \theta q - b)] \right\} \text{ for a.e. } \theta \in \Theta. \quad (2.C.4)$$

□

Chapter 3

Forward Guidance with Limited Commitment and Reputation

Javier González-Morín and Emil Mortensen

ABSTRACT

We argue that reputational concerns and limited commitment help overcome the forward guidance puzzle. Specifically, we study monetary policy in a New Keynesian model with a recurrent Zero Lower Bound where 1) the central bank is unable to commit to future policy, and 2) has private information about its discount factor giving rise to type-based reputation. Limited commitment imposes an upper bound on the length of sustainable promises while reputational concerns imply that forward guidance remains a useful policy tool at the Zero Lower Bound. Quantitative exploration shows that there is a non-monotonic relationship between the potency of forward guidance and the persistence of the Zero Lower Bound when the central bank cannot commit, while it is monotonically decreasing with full commitment. We find that the optimal length of forward guidance is 5 quarters with limited commitment compared to just 1 quarter with full commitment.

3.1. INTRODUCTION

The welfare cost of the Zero Lower Bound (ZLB) depends on the effectiveness of the unconventional monetary policy tools that the central bank can resort to, among which forward guidance. In the baseline New Keynesian model, forward guidance is excessively expansionary: the effect on current output grows without bounds in the promised length of a zero-interest rate path – the so called "forward guidance puzzle" ([Del Negro, Giannoni and Patterson, 2012](#)). Hence, any recession can be abated by a sufficiently generous promise about future interest rates ([McKay, Nakamura and Steinsson, 2016](#)).

In this paper, we take seriously one important, but overlooked, aspect of forward guidance: time inconsistency. We abandon the strong assumption that central banks can commit to a time-inconsistent policy, and we assume that forward-looking households understand the incentives of the central bank to keep its promise. We argue that this is sufficient to overcome the forward guidance puzzle.

We make two departures from a baseline New Keynesian model with recurrent ZLB events. First, we add limited commitment to capture the idea that central banks have an incentive to promise an expansionary future path of interest rates at the ZLB, but as soon as the economy leaves the ZLB, they would like to fulfill their actual mandate instead of sticking to the promised path of interest rates. Second, we add type-based reputation that we model by introducing private information of the central banker about her preferences. We assume that a central banker is only in charge of the central bank for a finite period, while the central bank is infinitely-lived. The central banker can be of two types: she may either be 1) a forward-looking one who wants to maximize the lifetime utility of infinitely-lived households, or 2) a myopic central banker who only cares about the utility of the households when she is in charge of the central bank. Modeling type-based reputation allows us to study how the effectiveness of forward guidance evolves over time as a function of the central bank's past policy decisions. In this environment, we first derive some theoretical properties of forward guidance under limited commitment. We then develop new insights by exploring the effect of forward guidance under limited commitment quantitatively, and we see how this effect compares to a model with full commitment.

In our model, the key decision faced by a central banker is whether to honor a past promise to keep interest rates low. The central banker trades off the cost of an excess expansion upon exiting the ZLB to honor a past promise against signaling its willingness to keep promises, which builds reputation. When ZLB events are recurrent, reputation is valuable to the central bank as it determines the effectiveness of forward guidance, and hence the severity of recessions in future ZLB episodes. An economy with a myopic central banker will not face an excess expansion upon exiting the ZLB while a forward-looking central banker may be willing to incur this cost to gain reputation.

We prove that the effect of forward guidance is less potent once agents factor in the incentives of the central bank to keep its promise. In particular, we show that there are realizations of the demand shock such that no promise is ever kept. We further prove that there exists a maximum duration of forward guidance that can be sustained in equilibrium. Any promise of a longer duration is not credible and hence futile. As a corollary, the marginal effect of increasing the promised duration of zero-interest is negative at this length.

We then calibrate the model to the US economy to explore the quantitative features of our model. We find that the effect of forward guidance is much mitigated once agents factor in the incentives of the central bank to keep its promises. The effect on output gap of promising one quarter of forward guidance is only half as powerful under limited commitment compared to full commitment. As agents realize that more expansionary promises are harder to keep, this mitigation is increasing in the promised

duration of keeping the interest rate at zero. For instance, the effect of promising forward guidance for three quarters is a 0.55 percentage points increase in the output gap under limited commitment while it is 1.4 percentage points under full commitment. For our calibration, the maximum sustainable length of forward guidance is five quarters, which is also the optimal duration of forward guidance, while the optimal length is one quarter with full commitment.

The incentives to keep a promise are linked to the persistence of ZLB episodes, resulting in a non-monotonic relationship between this persistence and the effect of forward guidance. More persistent ZLB episodes decrease the probability of exiting to a more expansionary state where interest rates are still kept at zero. This reduces the effect of forward guidance. Yet, the persistence of the ZLB under limited commitment also determines the value of having a high reputation. The more persistent the ZLB and therefore the more severe the drop in output and inflation, the more important it is for the central bank to have a high reputation implying that they will be more likely to honor its past promise. With full commitment, this latter effect is absent and therefore the effectiveness of forward guidance is monotonically decreasing in the persistence of the ZLB.

We then move on to compare our paper with the discounted Euler literature which is one of the more prominent attempts of solving the forward guidance puzzle. The main difference is that attenuation in our paper works though expected interest rates not responding fully to announcements about forward guidance while fully credible changes in expectations are still as potent as in the standard New Keynesian model. One possible way to distinguish our paper from that literature would be to look at how the yield curve responds to policy announcements. In our framework, we would not expect the yield curve to respond perfectly to forward guidance which is consistent with the results in [Campbell, Ferroni, Fisher and Melosi \(2019\)](#) that the FED has only limited influence on expected future interest at long horizons. We also relate our paper to [Bhattarai, Eggertsson and Gafarov \(2023\)](#) which show that quantitative easing serves as a commitment device for forward guidance and hence is complementary. We find that these two policies can actually be substitutes as quantitative easing is already reducing the expected future interest rates.

The paper is organized as follows. Section [3.2](#) relates our paper to the existing literature. Section [3.3](#) introduces forward guidance in a New Keynesian model augmented with limited commitment and reputation. Section [3.4](#) presents the main results. Section [3.5](#) explores some quantitative features of the model. Section [3.6](#) relates our model to the literature on the discounted Euler equation. Section [3.7](#) extends the model to study the interaction between forward guidance and quantitative easing. Section [3.8](#) concludes. Appendix [3.A](#) contains all the proofs.

3.2. RELATED LITERATURE

Our paper is related to four strands of literature: forward guidance, time inconsistency of monetary policy, reputation, and loose commitment.

Forward guidance has been the subject of a large and growing literature. [Eggertsson and Wood-](#)

ford (2003) shows that the possibility of engaging in forward guidance may significantly reduce the welfare cost of the ZLB which is consistent with the view expressed in Bernanke (2020), while Bilbiie (2019) solves for the optimal level of forward guidance in a New Keynesian model with full commitment. Nevertheless, the large impact of forward guidance predicted by the plain vanilla NK model has been questioned by Del Negro, Giannoni and Patterson (2012), which has given rise to the well-known "forward guidance puzzle." This puzzle has spurred a literature trying to solve it where a common approach has been to mute the importance of expected future output on current output by assuming some kind of bounded rationality as in Farhi and Werning (2019), Gabaix (2020) and García-Schmidt and Woodford (2019) or lack of common knowledge as in Angeletos and Lian (2018). Another approach is to assume incomplete markets as McKay, Nakamura and Steinsson (2016), yet the degree to which market incompleteness provides a solution to the forward guidance puzzle depends on the cyclicity of risk as remarked by Werning (2015), and elaborated further by Acharya and Dogra (2020). One common implication of these papers is that they give rise to a "discounted" Euler equation which attenuates the effect of higher expected future output. Nakata, Ogaki, Schmidt and Yoo (2019) show that this implies the prediction that to obtain the same effect of forward guidance as in the standard model, interest rates should be kept low for longer time.

Since the seminal contribution of Kydland and Prescott (1977) and Barro and Gordon (1983), it has been known that optimal policies may be time inconsistent, which implies that there are gains to commitment. Forward guidance is a perfect example of this. The ability of committing to a low future interest rate at the ZLB can enhance welfare but the policy is time-inconsistent: once the economy exits the zero lower bound, the central bank would be better off by setting a higher interest rate than promised to stabilize current economic outcomes. Related to this observation, Barthélemy and Mengus (2018) argue that central banks should increase inflation prior to the ZLB to make forward guidance more credible while Bhattarai, Eggertsson and Gafarov (2023) shows how quantitative easing can serve as a commitment device to keep future interest rates low. Following Stokey (1991) and Chari and Kehoe (1990), we study forward guidance in a model where central banks cannot commit to future policy, and the usefulness of forward guidance therefore depends on the extent to which it constitutes a sustainable plan in the sense that honoring past promises is sequentially optimal. Similar to Walsh (2018) and Nakata (2018), we exploit the recurrence of ZLB events, which provides incentives to honor past promises.

Another strand of literature that is relevant to our paper is reputation as pioneered by Kreps and Wilson (1982) and Milgrom and Roberts (1982). King and Lu (2022) estimate how the reputation of the FED has evolved assuming that the FED has private information about its discount factor, and Lu, King and Pasten (2016) solves for optimal reputation building in a New Keynesian model. DAVIS and Kirpalani (2020) study a model where the government can be of a commitment type or an optimizing type. The optimizing type chooses policy sequentially which also happens to be the case in our framework for the myopic type. Contrary to them, we don't assume that some type can perfectly commit, and the subject of study is different as they study how the effect of fiscal rules depend on the reputation of the government. Loisel (2008) shows how the inflation bias and the stabilization bias can be overcome by a reputation-concerned central bank if the punishment length is at least a few years.

This can induce the central bank to implement otherwise time-inconsistent optimal monetary policy in the NK model.

Finally, our paper is also related to the literature on loose commitment as developed in [Debortoli and Nunes \(2010\)](#), [Debortoli and Lakdawala \(2016\)](#), and [Debortoli, Maih and Nunes \(2014\)](#). These papers assume that the policy maker re-optimizes each period with some exogenous probability, while in our model the probability that a promise is kept is endogenous to the promise given. This endogeneity is key to solving the forward guidance puzzle.

The paper which is most closely related to ours is [Nakata \(2018\)](#), which uses a trigger strategy to sustain the optimal Ramsey plan in the tradition of [Chari and Kehoe \(1990\)](#). Our main departure from [Nakata \(2018\)](#) is that credibility arises as an equilibrium outcome of a signaling game with reputation, rather than being imposed through an exogenous punishment mechanism. This distinction has several important consequences. First, while the effect of forward guidance is constant across ZLB episodes in [Nakata \(2018\)](#), it varies over time on the equilibrium path in our model, as the probability that a promise is kept is an equilibrium object that evolves with the central bank’s track record of promise-keeping. This probability increases upon promise-keeping and decreases upon deviations. Second, our model implies that the first ZLB episode a central bank faces is categorically different from subsequent ones: with no prior ZLB history to update on, households rely solely on their prior, making forward guidance least effective precisely when it is first deployed. This provides a natural interpretation of the evidence in [Del Negro, Giannoni and Patterson \(2012\)](#), who document that the forward guidance puzzle is most severe in 2008-2009 — the Federal Reserve’s inaugural ZLB episode — consistent with our model’s prediction rather than constituting a puzzle within our framework. Finally, the time-variation in credibility is consistent with the empirical evidence in [King and Lu \(2022\)](#), who estimate that the Fed’s reputation for commitment has varied substantially over time — exactly the kind of variation that our model predicts should translate into heterogeneous forward guidance effectiveness across ZLB episodes.

3.3. A MODEL OF FORWARD GUIDANCE

We consider a plain vanilla New Keynesian model where demand shocks can bring the economy to the ZLB. Following [Eggertsson and Woodford \(2003\)](#) the central bank engages in forward guidance once the nominal interest rate reaches the lower bound of zero. We first define this policy with full commitment, and we then re-define it under the assumption of limited commitment. Forward guidance under limited commitment is modeled as a reputational game where the central banker can be of two types: myopic or forward-looking. Uncertainty about the discount factor captures the idea that forward guidance is a policy entailing a dynamic trade off and hence valued differently by the two types of central bankers. One natural way of thinking about these types is that there is uncertainty about how accommodative the central banker is to pressure from politicians to stabilize current economic outcomes.

3.3.1. New Keynesian model with a ZLB

We consider a standard New Keynesian model summarized by the New Keynesian Phillips curve and the IS curve following the text book treatment by [Woodford \(2003\)](#) and [Galí \(2015\)](#).

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa y_t \quad (\text{PC})$$

$$y_t = \mathbb{E}_t[y_{t+1}] - \frac{1}{\eta}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n) \quad (\text{IS})$$

where κ is the slope of the New Keynesian Phillips curve, i.e. it measures the sensitivity of inflation to the output gap, and η is the coefficient of relative risk aversion. The natural rate of interest (r_t^n) follows a finite-state Markov process that is assumed to be time-homogeneous and regular. Following the New Keynesian literature, we interpret a change to this rate as a demand shock.

As is standard in the New Keynesian literature, we need to close the model by making an assumption about how the interest rate is determined. We assume that monetary policy is determined by a central banker with the following quadratic loss function

$$L(\{y_t\}_{t=0}^\infty, \{\pi_t\}_{t=0}^\infty) = \sum_{t=0}^\infty \delta^t (\lambda y_t^2 + \pi_t^2), \quad \lambda = \frac{\kappa}{\theta}.$$

where the notation follows [Woodford \(2003\)](#), except that we allow the central banker to have a discount factor δ that may differ from that of the representative household. When $\delta = \beta$, this loss function represents an approximation to the level of expected utility of the household under some regularity conditions for an appropriately chosen weight on the output gap (λ).¹

We assume that a central banker is only in charge for one period, and that she cannot commit future central bankers. In addition, in this baseline model with no Forward Guidance, the central banker is assumed to design policy according to optimal discretion. That is, she re-optimizes the interest rate at each period taking expectations as given. This is equivalent to restrict the chosen policy to be stationary Markov in r_t^n . As a result, the central banker faces a sequence of static problems and therefore solves:

$$\min_{i_t, y_t, \pi_t} \lambda y_t^2 + \pi_t^2 \quad (\text{DP})$$

subject to (PC), (IS) and

$$i_t \geq 0, \quad (\text{ZLB})$$

where we impose that the central bank is constrained to choose a non-negative nominal interest rate, which gives rise to ZLB events.² We notice that the actual discount factor of the central banker does

¹The conditions are that the disturbances are small enough, the outcomes are close enough to the allocation around which the expansion is taken, and finally that distortions are small enough ([Woodford, 2003](#)).

²For this to be the case, we require the support of r_t^n to be rich enough.

not enter in the loss minimization problem due to discretionary restriction. In what follows, we assume that a solution to [DP](#), also called discretionary outcome, exists.³

3.3.2. Forward Guidance with Full Commitment

Once the constraint $i_t \geq 0$ is binding (i.e. the ZLB is hit), we assume that the central bank engages in forward guidance with full commitment which we define as a *binding* promise on future interest rates. This promise specifies a commitment to keeping the interest rate at zero for τ^{FC} periods.

DEFINITION 3.1. *Forward guidance with full commitment is an exogenous integer $\tau^{FC} \in \mathbb{N}$ that prescribes the binding duration of a zero nominal interest rate path after exiting the ZLB.*

We thus consider a form of Odyssean forward guidance where the central bank gives a promise about future policy rates rather than Delphic forward guidance where the central gives a prediction about future output and inflation. We define the effect of forward guidance as follows:

DEFINITION 3.2. *Let $y_{ZLB}(\tau^{FC}, r_t^n)$ be the level of the output gap at the ZLB under forward guidance. The impact response of output to forward guidance is $IR_y(\tau^{FC}, r_t^n) = y_{ZLB}(\tau^{FC}, r_t^n) - y_{ZLB}(0, r_t^n)$. Forward guidance is called effective if $IR_y(\tau^{FC}, r_t^n) > 0$.*

Hence, the impact response is the change in output at the ZLB when there is a promise of keeping the interest at zero for τ^{FC} periods.

A well-known result emphasized by for instance [Del Negro, Giannoni and Patterson \(2012\)](#), [Farhi and Werning \(2019\)](#) and [García-Schmidt and Woodford \(2019\)](#) is that the response of the output gap grows without bounds in τ^{FC} . This has become known as the "forward guidance puzzle." The [IS](#) curve illustrates the root of this puzzle. Future increases in expected output increase one-to-one output today. Hence, any changes in future interest rates that are believed with probability one can have arbitrarily large effects on current output.

3.3.3. Reputation and Limited Commitment

We add type-based reputation and limited commitment to the New Keynesian model with a ZLB. To model type-based reputation, we assume that the central banker may have a different discount factor than households as the central banker faces a finite mandate as she is replaced each period while households live forever. In particular, the discount factor of the central banker is $\delta \in \{0, \beta\}$. We assume that this discount factor is private information to the central banker with common prior $\mu_0 \in \Delta(\{0, \beta\})$. In words, we imagine that monetary policy is conducted by a central banker who can be of two types: 1) she can be a myopic one who only cares about consumer welfare during her mandate, or 2) she can be a forward-looking central banker who cares about consumer welfare even when she is not in charge.

³Existence of the discretionary outcome requires consistency between the two regimes associated with the ZLB. A candidate equilibrium must induce a partition of the state space such that the ZLB binds exactly in the states where the implied unconstrained nominal rate is non-positive, and is slack otherwise. In the quantitative section, we verify that this consistency condition is satisfied under our parameterization.

We assume that a mandate lasts for one period causing the myopic central banker to have a discount factor of zero while the forward-looking one has a discount factor of β .

We suppose that the central banker is replaced by someone with the same type as herself with probability $1 - \gamma$, while she is replaced by someone with a different type with probability $\gamma < \frac{1}{2}$. This gives rise to stochastic and persistent types, which is meant to capture stability of the central bank as an institution.

Each central banker engages in forward guidance at the ZLB under limited commitment. We define this as follows:

DEFINITION 3.3. *Forward guidance with limited commitment is an exogenous integer $\tau \in \mathbb{N}$ that prescribes a non-binding promised duration of a zero nominal interest rate path after exiting the ZLB.*

Hence, we consider whether some exogenously given τ can be sustained as an equilibrium. The definition of the impact response to forward guidance under full commitment extends naturally to the case under limited commitment.

As there is limited commitment, the central banker faces the strategic choice of whether to stick to the promised zero interest rate path or deviate to discretion in any of the τ periods upon exiting the ZLB. Following [Maskin and Tirole \(2001\)](#), we restrict attention to stationary Markovian strategies such that the central banker bases her decision on the following state:

DEFINITION 3.4. *A state is a triplet (r_t^n, μ_t, P_t) such that:*

- $r_t^n \in \mathbb{R}$ is the natural interest rate at time t
- $\mu_t \in \Delta(\{0, \beta\})$ is the private agents' belief distribution over types at time t
- $P_t \in \mathcal{P} = \{0, 1, \dots, \tau\}$ is the number of periods for which forward guidance is honored at calendar time t .

The first dimension is the natural interest rate r_t^n . The second dimension is the private sector assessment about the distribution of types, μ_t , namely reputation. The third dimension P_t is the number of times that a promise has been kept at calendar time t . For instance, $P_t = 0$ means that so far a promise has not been kept at time t , while $P_t = 3$ means that at time t , the central bank has kept the promise in period $t - 2$, $t - 1$, and t . We restrict attention to stationary Markovian equilibria in which both strategies and continuation beliefs depend on past history only through this state. This equilibrium-selection device yields a tractable recursive formulation, as inflation and output gap are pinned down by the natural interest rate, reputation and monetary policy status.

For each state of the economy, the central banker needs to decide whether or not to honor a past promise. Hence, we define strategies as follows:

DEFINITION 3.5. Fix a promise τ . For each calendar period t and type δ , a strategy σ prescribes the probability of transitioning from the promise period $P_{t-1} \in \mathcal{P}$ to the promise period $P_t = P_{t-1} + 1 \in \mathcal{P}$ with natural real interest rate r_t^n and reputation μ_{t-1} :

$$\begin{aligned} \sigma : \{0, \beta\} \times \mathbb{R} \times \Delta(\{0, \beta\}) \times \mathcal{P} &\longrightarrow [0, 1] \\ (\delta, r_t^n, \mu_{t-1}, P_{t-1}) &\longmapsto \mathbb{P}_\sigma(P_t | P_{t-1}, r_t^n, \mu_{t-1}, \delta) \end{aligned}$$

We only consider pure strategies, so the possible transition probabilities are either zero or one. Instead, under full commitment $\sigma \equiv 1$.

All actions are publicly observed, so the private sector will learn about the likelihood of each possible type, and the central bank perfectly tracks the evolution of agent's beliefs. The prior distribution $\mu_0 \in \Delta(\{0, \beta\})$ is common knowledge. The private sector is composed of Bayesian agents, such that beliefs after promise keeping evolve according to:

$$\mu_t(\delta; r_t^n, \mu_{t-1}, P_t) = \sum_{\hat{\delta}} \frac{\sigma(\hat{\delta}, r_t^n, \mu_{t-1}, P_t) \cdot \mu_{t-1}(\hat{\delta})}{\sum_{\tilde{\delta}} \sigma(\tilde{\delta}, r_t^n, \mu_{t-1}, P_t) \cdot \mu_{t-1}(\tilde{\delta})} \cdot \Gamma(\hat{\delta}, \delta)$$

where

$$\Gamma(\hat{\delta}, \delta) = \begin{cases} \Gamma(\hat{\delta}, \delta) = 1 - \gamma & \text{if } \hat{\delta} = \delta \\ \Gamma(\hat{\delta}, \delta) = \gamma & \text{if } \hat{\delta} \neq \delta \end{cases}$$

for $\delta, \hat{\delta} \in \{0, \beta\}$, where $\gamma < \frac{1}{2}$.

Whenever there is no promise in place, reputation mechanically evolves following the transition probabilities, as there is no strategic choice being made by the central banker, and hence households do not update from policy.

3.3.4. Equilibrium

An equilibrium in our limited commitment-reputation model is to constitute both a macroeconomic equilibrium, in the sense that the macroeconomic variables are given by the New Keynesian model, and a strategic equilibrium such that the strategies of the central bankers together with the private sector's beliefs constitute a Perfect Bayesian Equilibrium.

DEFINITION 3.6. An equilibrium is a tuple (σ, μ, y, π, i) such that:

- Given (y, π, i) , the pair (σ, μ) is a Perfect Bayesian Equilibrium:
 - Beliefs μ are consistent with the strategies σ of the central bankers on the equilibrium path.
 - Strategies σ of the central bankers are sequentially rational given the beliefs μ .
- Given strategies and beliefs (σ, μ) , output gap y , inflation π , and nominal interest rates i satisfy (IS) and (PC).

3.3.5. Computation of the Equilibrium

To solve for the macroeconomic equilibrium, we first notice that for each pair (σ, μ) there is a subset $Z_{\sigma, \mu} \subset \mathbb{R}$ such that the ZLB binds if and only if $r_t^n \in Z_{\sigma, \mu}$. At the ZLB, output gap and inflation are given by the New Keynesian Phillips curve and the IS curve, imposing that the nominal interest rate is zero. We report the full set of equations characterizing the macroeconomic equilibrium in Appendix 3.C. Below we report the expression for the IS curve at the ZLB.

$$y_t^Z = \mathbb{E}[y_{t+1}] + \frac{1}{\eta}(\mathbb{E}[\pi_{t+1}] + r_t^n)$$

where expected output gap is given by⁴

$$\mathbb{E}[y_{t+1}] = \mathbb{E} \left[\underbrace{y_{t+1}^Z \mathbb{1}\{r_{t+1}^n \in Z_{\sigma, \mu}\}}_{(a)} + \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma, \mu}\} \left(\underbrace{y_{t+1}^P \mathbb{E}_\delta^{\mu_t}[\sigma(\delta)]}_{(b)} + \underbrace{y_{t+1}^D \mathbb{E}_\delta^{\mu_t}[1 - \sigma(\delta)]}_{(c)} \right) \right]$$

The operator $\mathbb{E}$ denotes the expectation over r_{t+1}^n given r_t^n , $\mathbb{E}_\delta^{\mu_t}$ denotes the expectation with respect to δ using distribution μ_t . Z denotes the *zero lower bound*, P denotes *promise keeping* and D denotes *discretion*. Hence, y_t^P is output under promise keeping, and y_t^D is output in discretion. We emphasize that output gap and inflation depend on the state of the economy but we have suppressed it for the sake of exposition.

To better describe how to calculate the expectation, we decompose expected output gap into three terms. The term (a) corresponds to output gap whenever the ZLB is binding again. Terms (b) and (c) refer to a realization of the natural rate of interest such that the ZLB is not binding. Term (b) is output gap under promise keeping, i.e. the output gap that prevails if the central bank chooses $i_{t+1} = 0$, multiplied by the probability that the central bank keeps its promise, $\mathbb{E}_\delta^{\mu_t}[\sigma(\delta)]$. For notational convenience, we have suppressed the dependence of this probability on the state of the economy. Finally, term (c) is output gap in normal times, which is multiplied by the probability that the central bank breaks the promise, $\mathbb{E}_\delta^{\mu_t}[1 - \sigma(\delta)]$.

Having explained how to solve for the macroeconomic equilibrium at the ZLB, we now explain how to solve for the macroeconomic equilibrium in normal times, i.e. for any shock $r_t^n \notin Z_{\sigma, \mu}$ and $P_t = 0$. As the central bank follows optimal discretion, monetary policy is set to minimize the instantaneous loss function such that the first order condition reads:

$$-\kappa\pi_t = \lambda y_t$$

Output gap and inflation are then given by the New Keynesian Phillips curve and the IS curve. The full system of equations is reported in Appendix 3.C.

⁴Expected inflation can be obtained similarly.

Finally, the economy can also be in promise keeping, i.e. it is away from the ZLB ($r_t^n \notin Z_{\sigma,\mu}$) and there is a promise in place ($P_t > 0$). The central banker then decides whether to keep the promise or not for another period as captured by the strategy σ . Finally, the economy can be in the last period of promise keeping out of the ZLB, in which $P_t = \tau$ and no strategic choice is made. This follows as the economy simply transitions to either discretion or the ZLB depending on the realization of the natural interest rate. If the economy were to hit the ZLB in the sense that the desired interest rate is negative, we assume that the central bank renews the promise of length τ . This helps us gain a recursive structure.

ASSUMPTION 3.1. *If $P_t \geq 0$, and $r_{t+1}^n \in Z_{\sigma,\mu}$, then $P_{t+1} = 0$.*

The problem of the central bank given the state of the economy at the ZLB is given by

$$L_t^Z(\delta) = \lambda(y_t^Z)^2 + (\pi_t^Z)^2 + \delta \mathbb{E} \left[\underbrace{L_{t+1}^Z(\delta) \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\}}_{(a)} + \min_{x \in \{0,1\}} \left\{ \underbrace{x L_{t+1}^P(\delta)}_{(b)} + \underbrace{(1-x) L_{t+1}^D(\delta)}_{(c)} \right\} \mathbb{1}\{r_{t+1} \notin Z_{\sigma,\mu}\} \right]$$

The loss at the ZLB depends on the current realization of output and inflation and the discounted expected future loss. The future expected loss can be decomposed into 3 terms. Term (a) reflects the loss of being at the ZLB, while term (b) reflects the loss of promise keeping. Finally, term (c) is the loss from discretion. The problems in normal times and in promise keeping are carefully described in Appendix 3.C.

The above intertemporal loss function plays an important role once the economy exits the ZLB, and the central banker decides whether to honor a promise or to break it and conduct discretionary monetary policy. The trade off she faces is the following: on the one hand, by keeping the promise, inflation and the output gap are too expansionary compared to what could have been achieved by setting the interest rate according to optimal discretion. All else equal, this provides the central banker with an incentive to renege on her promise. On the other hand, by keeping the promise the central banker signals its ability to keep promises in future ZLB events. Consequently, the private sector's assessment about her type evolves over time as a function of the behavior of the central banker and the state of the economy.

3.4. SOLUTION TO THE FORWARD GUIDANCE PUZZLE

A New Keynesian model augmented with limited commitment and reputation does not suffer from the forward guidance puzzle. While some level of forward guidance at the ZLB remains useful, we find that an unreasonably expansionary promise cannot be sustained as an equilibrium strategy.

When the central banker decides whether to keep the promise or conduct optimal discretionary policy, she faces the following trade off: on the one hand, keeping a promise causes an excess expansion

today. On the other hand, keeping the promise helps the central bank gain reputation which is useful to stabilize output gap and inflation in future ZLB events. A first result is the following:

PROPOSITION 3.1. *For any state $s = (r_t^n, \mu_t, P_t)$, there exists $\bar{\lambda}(s) > 0$ such that if $\lambda \in [0, \bar{\lambda}(s)]$, then the myopic central banker optimally breaks any promise $\tau > 0$ at state s . Equivalently,*

$$\sigma(0, r_t^n, \mu_t, P_t) = 0 \quad \text{whenever} \quad \lambda \leq \bar{\lambda}(r_t^n, \mu_t, P_t).$$

For a large $\lambda > 0$, even the myopic central banker may keep her promise as reputational gains have a direct impact on current economic outcomes. In fact, in this case it is possible to find values of the parameters such that the myopic central banker will be willing to keep the promise.⁵ Going forward we assume that λ is small enough such that the myopic central banker never keeps her promise.

Another result is that there are states of the world where the reputation gains are not important enough to outweigh the cost of an excess expansion:

PROPOSITION 3.2. *There exists a sufficiently high realization $\hat{R}$ in the support of r_t^n such that, in any equilibrium, for all states (r_t^n, μ_t, P_t) with $r_t^n \geq \hat{R}$, and for both types $\delta \in \{0, \beta\}$, we have*

$$\sigma(\delta, r_t^n, \mu_t, P_t) = 0.$$

The intuition is that for a very positive demand shock, the expansion that the economy would undergo with a zero interest rate would be too costly. It would therefore be preferable to deviate to discretion to mitigate the boom in output gap and inflation. In fact, while the loss of the central banker is always bounded under optimal discretion for any realization of the demand shock, the loss goes to infinity as the realization increases if the interest rate is kept at zero.

We further find that the existence of even a minimal cost of keeping the promise is enough to ensure the existence of a pooling equilibrium where neither type keeps the promise:

PROPOSITION 3.3. *For any promise τ , there exist equilibrium strategies σ such that for any δ , r_t^n , μ_t and P_t , we have that $\sigma(\delta, r_t^n, \mu_t, P_t) = 0$.*

This is due to the fact that pessimistic off-path beliefs render deviations to promise keeping futile as the reputation channel is muted.

Combined with Proposition 3.1, we can therefore restrict attention to semi-separating equilibria where the forward-looking central banker keeps her promise in some states of the world only, and the myopic central banker never keeps her promise. There can be no other equilibrium in which forward guidance is effective.

While an excess expansion provides the central banker with a reason to renege on past promises, the reputation channel can be strong enough to provide her with an adequate incentive to honor them.

⁵However, we find that with our preferred parameterization, the myopic central banker never keeps her promise.

To better understand how this channel works, suppose that the forward-looking central banker keeps her promise in some states of the world. In that case, the effect of forward guidance increases in the probability that agents assign to the central banker being of the forward-looking type. This, in turn, provides the central banker with a reason to indeed keep the promise. For a strong enough reputation channel, there exists a semi-separating equilibrium where the forward-looking central banker keeps her promise in some states of the world. We verify that, with our preferred parameterization of the model, such an equilibrium indeed exists, i.e. promises are kept in states of the world where the realization of the demand is sufficiently moderate. We are now ready to state the main result of the paper: for a too generous promise, the effect of forward guidance is zero.

PROPOSITION 3.4. *There exists $\bar{\tau} \in \mathbb{N}$ such that for all $\tau > \bar{\tau}$, no equilibrium can sustain promise-keeping.*

We thus find that there is a maximum level of forward guidance that is sustainable in the sense that promises are kept in some states of the world. Any promise with a longer duration will not be believed and hence it will fail to stabilize economic outcomes at the ZLB.

Intuitively, the cost of keeping a promise is increasing in its length. A generous promise, if believed, implies a significant expansion once the economy exists the ZLB. Hence, agents understand that a promise cannot be kept if it is too generous, and therefore keeping it does not constitute a strategic equilibrium. The only equilibrium is therefore pooling where promises are never kept, rendering forward guidance ineffective. From Proposition 3.4, we immediately have the following corollary:

COROLLARY 3.1. *The marginal impact of forward guidance is negative at $\tau = \bar{\tau}$:*

$$IR_y(\bar{\tau} + 1) - IR_y(\bar{\tau}) < 0$$

Once a promise goes from being sustainable to unsustainable, the effect of forward guidance drops to zero, and hence the marginal effect of increasing the duration is negative. This corollary is an example of how our model is different from the literature giving rise to a "discounted Euler equation" where the marginal effect of forward guidance is always weakly positive. We further discuss differences with that literature in section 3.6.

We provide a graphical representation of Proposition 3.4 and Corollary 3.1 in Figure 3.1. We find that there exist semi-separating equilibria where the forward-looking central banker keeps her promise for moderate demand shocks for a duration of the promise being no higher than five periods corresponding to one year and a quarter.

Despite forward guidance being a useful policy tool in the semi-separating equilibrium, the effect is much mitigated compared to the full commitment case. This happens for two reasons. First, households realize that the myopic central banker will never keep her promise. Second, they realize that there are




Figure 3.1: Mitigation of forward guidance effect for $\mu_0(\beta) = 0.6$.

states of the world where it is too costly for any central banker to keep her promise regardless of her type.

3.5. THE EFFECT OF FORWARD GUIDANCE

For the quantitative exploration of the model, we will assume that the loss function takes the following form:

$$L(\{y_t\}_{t=0}^{\infty}, \{\pi_t\}_{t=0}^{\infty}) = \sum_{t=0}^{\infty} \delta^t (\lambda(y_t - \bar{y})^2 + \pi_t^2), \quad \delta \in \{0, \beta\}$$

where $\bar{y} > 0$ is the targeted output gap. The parameter $\bar{y}$ captures that there may be inefficiencies in the economy such that the central seeks to stabilize output at a level which is higher than the natural one and thereby target a positive output gap. We include a positive output gap target to counteract the deflationary forces caused by the ZLB.

We assume that the natural rate of interest follows an AR(1) process:

$$r_t^n = (1 - \rho)\bar{r} + \rho r_t^n + \varsigma \xi_t, \quad \xi_t \sim \mathcal{N}(0, 1)$$

We discretize the AR(1) process by the [Tauchen \(1986\)](#) method where we choose three values for the shock to the natural interest rate. We choose three to have a realization such that the ZLB binds, a realization where the forward-looking central banker keeps the promise, as well as a realization where it is excessively costly for any central banker to keep the promise. This is sufficient to illustrate the main points of the paper.

We choose to parameterize the model such that one period corresponds to one quarter. We calibrate the persistence of the demand shock (ρ), the variance of the demand shock (ς^2), and the targeted output gap ($\bar{y}$), while we choose values that we consider standard in the New Keynesian literature for the remaining parameters. We display all of the values in [Table 3.1](#).

Parameter	Value	Description	Target
β	0.99	Discount factor	4% annual interest rate
κ	0.02	Slope of Phillips Curve	Walsh (2018)
$\bar{y}$	1%	Output gap target	Average output gap of zero
η	2.0	Relative risk aversion	McKay, Nakamura and Steinsson (2016)
λ	0.003	Weight on output gap	Walsh (2018)
$\mu_0(\beta)$	0.6	Prior probability of high type	King and Lu (2022)
γ	0.03	Transition probability of types	8 years mandate
ς	0.014	Standard deviation of ξ_t	Std. deviation of output gap
ρ	0.62	Auto correlation of r_t^n	Auto correlation of inflation
$\bar{r}$	1%	Natural real interest rate	$-\log(\beta)$

Table 3.1: Parameter values

We calibrate the standard deviation of the innovation (ς) such that the model with no forward guidance matches the standard deviation in US output gap in the period 1987-2018. Similarly, we calibrate the persistence (ρ) of the demand shock to the auto correlation of US inflation. We retrieve data from FRED on real GDP in the US over the period 1987-2018. We obtain a standard deviation of output gap of 1.87, and an auto correlation of 0.54, at a quarterly level in the period 1987-2018. The implied values for the standard deviation for the innovation (ς) and for the persistence of the demand shock (ρ) are 1.4 and 0.62, respectively.

We choose to calibrate the output target to match a zero average output gap. If we did not include a positive target for the output gap, the model would predict a too low average output gap due to the deflationary bias introduced by the ZLB as explained by [Eggertsson \(2006\)](#).

For the remaining parameters, we choose values that we consider standard in the literature. First, as we consider a quarterly model, we choose β , the discount factor of private agents, to be 0.99, and hence the value of $\bar{r}$, i.e. the natural rate of interest is equal to $-\log(\beta)$, which corresponds to a steady state interest rate of around 4 % a year. We choose a value of the slope of the Phillips curve (κ) to be 0.02. For the coefficient of relative risk aversion (η), we choose a value of 2.

3.5.1. Reputation

The existence of a semi-separating equilibrium implies that the effect of forward guidance depends on reputation. The higher the probability agents assign to the central banker being the forward-looking type, the larger the effect of forward guidance on output as shown in Figure 3.2. We also report the effect of forward guidance under full commitment in red for comparison.




Figure 3.2: Reputation and forward guidance. In red, the effect of forward guidance under full commitment; in blue, the effect of forward guidance under limited commitment. The darker the blue, the higher the belief $\mu_0(\beta)$, ranging from 0 to 1 with increments of 0.1.

From Figure 3.2 we see that for $\mu_0(\beta) = 0$, forward guidance is almost futile as a policy tool as the households expect the current central banker to not keep the promise, and assign very low probability ($\gamma = 0.03$) to a change in the type of the central banker. It is then monotonically increasing in the belief that the central banker is the forward-looking type. The effect of forward guidance with limited commitment is mitigated compared to full commitment, even if households believe that the central banker is the forward-looking type (i.e. $\mu_0(\beta) = 1$). This occurs as households realize that (1) there are realizations of the natural interest rate such that it is too costly to keep the promise for the forward-looking type, and (2) types are mean-reverting. This mitigation effect is not very strong for $\tau = 1$. The intuition is that given that the economy is at the ZLB, it is considered very likely that the economy will not be in an excessive boom in the following period. That is, the realization of the demand shock is likely to be only moderate once the economy exits the ZLB. We see a stronger mitigation effect

for larger values of τ when we compare limited commitment to full commitment. The reason is that for higher values of τ , it is more likely that the economy will be hit by a very positive demand shock in which case even the forward-looking central banker will have to deviate from her promise rendering forward guidance less credible.

To illustrate the dynamics of the model, we consider a situation where households initially hold a belief of 0.5 that the central banker is forward-looking, and the economy is in normal times for the first 4 periods while in period 5 the ZLB hits and the central banker engages in forward guidance with $\tau = 1$. For the purpose of illustration, we assume that the ZLB only binds in period 5. In Figure 3.3, we show how beliefs evolve depending on whether the central banker keeps her promise or not.




Figure 3.3: Reputation convergence with $\mu_0(\beta) = 0.5$.

We see that immediately after the promise keeping decision of the central banker, there is a large change in beliefs. Following this change, beliefs then converge back to the stationary distribution in the absence of further ZLB periods.

We simulate our model for 100 periods in a path where the type of the central bank happens to not change with the aim of showing how reputation evolves depending on the type of the central banker, and how the reputation of the central banker influences economic outcomes. In particular, we plot the value of inflation and how it evolves together with reputation in a model with $\tau = 3$ and $\mu_0(\beta) = 0.5$.




Figure 3.4: A simulation of reputation and inflation with $\tau = 3$ and $\mu_0(\beta) = 0.5$.

We see that having a high reputation is associated to a smaller drop in inflation at the ZLB. Interestingly, if we consider inflation at the ZLB that hits the economy in periods 60 and 75, respectively, we notice that the drop is mitigated in period 75 for the forward-looking type (yellow line) while it is exacerbated for the myopic type (purple line). This happens as the occurrence of a ZLB in the recent past has given agents a good idea of the type of the central banker.

3.5.2. Forward guidance and the persistence of the ZLB

We find that the persistence of the demand shock plays an important role in determining the sustainability of forward guidance as well as its effect on the output gap at the ZLB. To show that, we vary the autoregressive coefficient (ρ) while at the same time adjusting the variance of the innovation ς^2

such that the variance of the natural interest rate is kept fixed.⁶

In Figure 3.5 we show the effect on the output gap at the ZLB of promising a zero interest rate for a time horizon of τ periods as a function of the persistence of the demand shock. We also show the effect of forward guidance with full commitment in red for comparison.




Figure 3.5: Forward guidance and the persistence of the natural rate of interest

We see that there is a non-monotonic relation between the effect of forward guidance and the persistence of the demand shock for some values of τ when the central bank cannot commit, while the effect of forward guidance is smoothly decreasing in the persistence of the demand shock in a model with full commitment.

Intuitively, a higher persistence of the demand shock implies that the economy is less likely to exit the ZLB in the following period which in turn means that the economy is less likely to transition to the promise keeping state. This explains why the effect of forward guidance is decreasing in the persistence of the demand shock with full commitment.

With limited commitment, there is no longer a monotonic relationship between the persistence of the demand shock and the effect of forward guidance. For a very low persistence of the shock, we find that a promise of keeping the interest rate at 0 for one quarter after the economy exits the ZLB

⁶We perform a similar exercise for the standard deviation of the demand shock, and we report the results in appendix 3.D.

has a zero impact on the output gap at the ZLB. This happens as a low persistence of the demand shock renders the ZLB less severe due to the fact that the economy is likely to exit in the following period, causing expectations about future output gap and inflation to be rather expansionary. Due to the fact that the ZLB is not very severe, the value of having a high reputation, allowing the central banker to stabilize economic outcomes at the ZLB, is modest. Hence, the central banker prefers stabilizing economic outcomes once out of the ZLB by deviating to discretion instead of building reputation.

We see that for a persistence of around 0.51, the effect of forward guidance on output gap jumps from 0 % to 0.25 % when $\tau = 1$. The reason is that as the persistence increases, the ZLB gets more severe implying that the central banker faces stronger incentives to keep her promise to gain reputation. After this jump, the effect of forward guidance then smoothly declines in the persistence of the demand shock for the same reasons as under full commitment: the higher the persistence, the less likely that the economy will transition to the promise keeping state in the following period. We see that eventually the effect of forward guidance drops to zero. This is due to the fact that for a too high persistence, the promise keeping state gets too expansionary implying that the central banker will have to deviate from her promise as it is too costly to keep it.

We find that the higher the value of τ , the smaller the set of values of the persistence such that forward guidance can be sustained as an equilibrium strategy. The intuition is that longer promises are associated to larger expansions implying that the central banker needs stronger incentives to keep the promise, i.e. a very bad ZLB while at the same the persistence cannot be too high as the promise keeping state would then be too expansionary.

3.5.3. Welfare loss

The presence of a ZLB causes a welfare loss as the central bank is unable to stabilize inflation and output gap through adjustments to the nominal interest rate. Instead, the central bank is forced to rely on unconventional policy measures such as forward guidance to stabilize the economy ([Eggertsson and Woodford, 2003](#)).

We compute the loss associated to the ZLB by applying the discounted loss function of the forward-looking central banker which approximates household welfare. Following [Galí \(2015\)](#), we can interpret this loss as the fraction of steady state consumption up to an additive constant that the household is willing to give up to avoid fluctuations in the economy. We find that the loss is 0.003 % of steady state consumption if the central bank cannot rely on forward guidance. It is not surprising that we do not find important quantitative welfare losses of the ZLB as we follow the approach by [Lucas \(1987\)](#), and it is well known that the estimated welfare losses of business cycle fluctuations are very small.

We then report the relative loss at the ZLB for central bank that can conduct forward guidance. We compare a central bank that has only limited commitment to a central bank of full commitment. The relative losses are reported in Table [3.2](#).

Scenario	$\tau = 1$	$\tau = 2$	$\tau = 3$	$\tau = 4$	$\tau = 5$
Loss with FC	95.6 %	95.8 %	105.4 %	132.8 %	191.9 %
Loss with LC	94.4 %	88.8 %	83.8 %	80.1 %	78.7 %

Table 3.2: Welfare losses relative to ZLB with no FG. Initial belief distribution is $\mu_0(\beta) = 0.6$.

In general, we find that the relative losses are smaller once forward guidance is available to the central bank. We find that a central bank which is forced to commit to a zero interest rate should optimally choose a length of $\tau = 1$. On the contrary, a central bank that does not have access to a commitment technology should optimally choose a length of $\tau = 5$ corresponding to five quarters.

We therefore find that not having access to a commitment technology causes the central bank to give a longer promise. The intuition is two-fold. First, as agents realize that the central bank may deviate from its promise, the expansionary effect of a promise at the ZLB is reduced under limited commitment. Second, as the central bank has not committed, it can choose to deviate from the given promise in a state where it is very costly to keep it.

3.6. RELATION TO DISCOUNTED EULER EQUATION

[McKay, Nakamura and Steinsson \(2016\)](#) emphasizes that the highly forward looking nature of the Euler equation is at the core of the forward guidance puzzle. A large literature solving the this puzzle can be summarized by the so-called "discounted Euler equation" which takes the following form:

$$y_t = \Omega \mathbb{E}_t[y_{t+1}] - \frac{\chi}{\eta} (i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n)$$

and can be iterated forward to obtain:

$$y_t = -\frac{\chi}{\eta} \sum_{j=0}^{\infty} \Omega^j \mathbb{E}_t [i_{t+j} - \pi_{t+j+1} - r_{t+j}^n]$$

In the standard New Keynesian model $\Omega = \chi = 1$. The condition $\Omega < 1$ dampens the effect of future expected interest rates on current economic outcomes which explains why forward guidance about interest rates in the distant future is mitigated with a discounted Euler equation compared to the standard model.

However, we find it paradoxical that agents make consumption-savings decisions looking infinitely into the future while failing to understand the incentives of the central bank to keep a time-inconsistent promise. The ability of agents to understand the incentives of the central bank is what we add to the standard model.

The difference between our limited commitment-reputation model and this literature is that we do not obtain attenuation of forward guidance through discounting of the Euler equation, i.e. we have $\Omega = 1$. Instead, we obtain attenuation of forward guidance by assuming that agents understand the

incentives of the central bank to keep or renege on a promise. Hence, our attenuation works through the lack of effect that a promise has on expectations regarding the future interest rate path.

In Table 3.3, we compare the output gap and inflation that pertain at the ZLB in the standard model, the discounted Euler equation model, and our model with limited commitment and reputation.

Model		Output gap	Inflation
Standard NK	($\Omega = 1.00, \chi = 1.00$)	-6.49%	-0.33%
Discounted Euler	($\Omega = 0.97, \chi = 0.75$)	-4.48%	-0.18%
LC and reputation	($\Omega = 1.00, \chi = 1.00$)	-6.49%	-0.33%

Table 3.3: Recession at the ZLB with no forward guidance

A discounted Euler equation attenuates the fall in output and inflation at the ZLB compared to the standard model as stressed by [McKay, Nakamura and Steinsson \(2016\)](#). Instead, our model with limited commitment and reputation does not impact economic outcomes at the ZLB in the absence of forward guidance compared to the standard NK model. This is due to the fact that the discounted Euler literature obtains mitigation of a ZLB by reducing the front-loading effect of future drops in inflation and output gap.

We show how the economy responds to a promise of keeping the interest rate at 0 for τ periods in Figure 3.6 in the three different models.




Figure 3.6: Discounted Euler equation model

Both a model with a discounted Euler equation and our limited commitment-reputation model provide attenuation of the effect of forward guidance, but an important difference between the two models is

that with the former, the marginal effect of giving a more generous promise is always weakly positive while in the latter we obtain that the marginal effect may be negative if the promise induces agents to lose faith in the ability of the central bank to keep its promise. Hence, central bank communication may be more expansionary by giving a humble promise that can be kept compared to a more expansionary one, which agents expect central banks to renege on.

While an empirical investigation of forward guidance to discriminate between lack of commitment or behavioral agents is beyond the scope of this paper, we notice that our paper is consistent with the notion that central banks have only limited influence on expected interest rates far out in the future. This is consistent with the results in [Campbell, Ferroni, Fisher and Melosi \(2019\)](#) that the FED has limited ability to affect expectations at long horizons.

3.7. EXTENSION: FORWARD GUIDANCE AND QUANTITATIVE EASING

[Bhattarai, Eggertsson and Gafarov \(2023\)](#) show that quantitative easing (an increase in short run government debt) stimulates the economy by inducing the central bank to choose a lower interest rate in the future. They informally discuss how quantitative easing therefore can serve as a commitment device to forward guidance and hence work to strengthen the effect of it. We now extend our quantitative model to study the interaction between forward guidance and government debt in our limited commitment-reputation model.

We assume that there is a level of government spending at the ZLB that is pure waste in the sense that it does not affect aggregate demand nor consumer welfare. Whatever is spent at the ZLB is financed by short term debt that has to be paid back in the following period.

[Bhattarai, Eggertsson and Gafarov \(2023\)](#) microfound a loss function that is increasing in the level of taxation as tax payments are associated to a decrease in welfare due to their distortionary nature. We then follow their approach and assume a loss function of the following form:

$$L(\{y_t\}_{t=0}^{\infty}, \{\pi_t\}_{t=0}^{\infty}, \{T_t\}_{t=0}^{\infty}) = \sum_{t=0}^{\infty} \delta^t \left[\lambda \cdot (y_t - \bar{y})^2 + \pi_t^2 + \lambda_T \cdot T_t^2 \right]$$

The term T denotes taxes that are raised to repay government debt raised in the previous period. In particular, $T_t = (1 + i_{t-1}) \cdot G_{t-1}$. The model is otherwise identical to the one described in [Section 2.2](#). We choose the same calibration as in the baseline model, and we put the weight on taxes $\lambda_T = 0.003$.

In [Table 3.4](#) we report the interest rate chosen by the central bank at the median realization of the demand shock in the model with and without government debt but with no forward guidance. We find that having higher government debt reduces the interest rate chosen by the central bank when they conduct discretionary policy. This result replicates the finding of [Bhattarai, Eggertsson and Gafarov \(2023\)](#).

Debt level		i_t
$G_t = 0$	(baseline model)	0.66 %
$G_t = 0.03$	(model with debt)	0.35 %
$G_t = 0$	(model with debt)	0.67 %

Table 3.4: Interest rates under discretionary policy with no forward guidance.

We find that the interest rate in discretion is 0.66 % in the baseline model where government debt is always zero. We then compute the interest rate in a model where the government issues debt at the ZLB. The interest rate is lower in discretion (0.35%) when the government has a positive level of debt. The intuition is that the central bank tries to minimize the tax payments of the private agents while still taking into account how the interest rate influences output and inflation. As a consequence of this, the zero lower bound is less severe in such a model which causes the interest rate to be higher (0.67%) in discretionary periods with no debt as expectations about future outcomes are more expansionary.

We then study the consequences of forward guidance in the model with government debt. We display in Figure 3.7 the increase in output at the ZLB of promising an interest rate of zero in the period where the economy exists the ZLB.




Figure 3.7: Forward guidance with government debt

From the graph, it follows that the effect of forward guidance is muted once there is government debt. The reason for this is that the presence of government debt reduces the optimal interest rate of the central bank in discretion. This renders a promise of low interest rates in the future superfluous as private agents already expect the central bank to keep rates low. We therefore obtain that short term government debt and forward guidance are substitute policies whereas [Bhattarai, Eggertsson and Gafarov \(2023\)](#) argue that they are complements.

3.8. CONCLUSION

We showed how a standard New Keynesian model augmented with limited commitment and type-based reputation bounds the effect of forward guidance. An important implication of the model is that the marginal effect of the promised horizon becomes negative when the promise is judged unsustainable by private agents. We demonstrated a strong interaction between the persistence of the ZLB, and the sustainability and effect of forward guidance. We conjecture that our model could be useful for studying other time inconsistent policies.

BIBLIOGRAPHY

- Acharya, S. and Dogra, K. (2020). Understanding HANK: Insights From a PRANK. *Econometrica*, 88(3), 1113–1158.
- Angeletos, G. M. and Lian, C. (2018). Forward Guidance without Common Knowledge. *American Economic Review*, 108(9), 2477–2512.
- Angeletos, G. M. and Sastry, K. A. (2020). Managing Expectations: Instruments Versus Targets*. *The Quarterly Journal of Economics*, 136(4), 2467–2532.
- Barro, R. J. and Gordon, D. B. (1983). Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12(1), 101–121.
- Barthélemy, J. and Mengus, E. (2018). The signaling effect of raising inflation. *Journal of Economic Theory*, 178, 488–516.
- Bassetto, M. (2019). Forward guidance: Communication, commitment, or both?. *Journal of Monetary Economics*, 108, 69–86.
- Bernanke, B. S. (2020). The New Tools of Monetary Policy. *American Economic Review*, 110(4), 943–83.
- Bhattarai, S., Eggertsson, G. B., and Gafarov, B. (2023). Time Consistency and Duration of Government Debt: A Model of Quantitative Easing. *Quarterly Journal of Economics*, Forthcoming.
- Bilbiie, F. O. (2019). Optimal Forward Guidance. *American Economic Journal: Macroeconomics*, 11(4), 310–45.
- Boneva, L., Harrison, R., and Waldron, M. (2018). Threshold-based forward guidance. *Journal of Economic Dynamics and Control*, 90, 138–155.
- Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3), 383–398.
- Campbell, J. R., Ferroni, F., Fisher, J. D., and Melosi, L. (2019). The limits of forward guidance. *Journal of Monetary Economics*, 108, 118–134.
- Chari, V. V. and Kehoe, P. J. (1990). Sustainable Plans. *Journal of Political Economy*, 98(4), 783–802.

- Clayton, C. and Schaab, A. (2022). A Theory of Dynamic Inflation Targets. Toulouse School of Economics (TSE), TSE Working Papers.
- Debortoli, D. and Nunes, R. (2010). Fiscal policy under loose commitment. *Journal of Economic Theory*, 145(3), 1005–1032.
- Debortoli, D., Maih, J., and Nunes, R. (2014). Loose Commitment in Medium-Scale Macroeconomic Models: Theory and Applications. *Macroeconomic Dynamics*, 18(1), 175–198.
- Debortoli, D. and Lakdawala, A. (2016). How Credible Is the Federal Reserve? A Structural Estimation of Policy Re-optimizations. *American Economic Journal: Macroeconomics*, 8(3), 42–76.
- Del Negro, M., Giannoni, M. P., and Patterson, C. (2012). The Forward Guidance Puzzle. *FRB of New York Staff Report*, 3133–58.
- Dovis, A. and Kirpalani, R. (2020). Fiscal Rules, Bailouts, and Reputation in Federal Governments. *American Economic Review*, 110(3), 860–88.
- Eggertsson, G. and Woodford, M. (2003). The Zero Bound on Interest Rates and Optimal Monetary Policy. *Brookings Papers on Economic Activity*, 34(1), 139–235.
- Eggertsson, G. (2006). The Deflation Bias and Committing to Being Irresponsible. *Journal of Money, Credit and Banking*, 38, 283–321.
- Farhi, E. and Werning, I. (2019). Monetary Policy, Bounded Rationality, and Incomplete Markets. *American Economic Review*, 109(11), 3887–3928.
- Gabaix, X. (2020). A Behavioral New Keynesian Model. *American Economic Review*, 110(8), 2271–2327.
- Galí, J. (2015). Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications Second edition. Princeton University Press.
- García-Schmidt, M. and Woodford, M. (2019). Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis. *American Economic Review*, 109(1), 86–120.
- Hodrick, R. and Prescott, E. C. (1997). Postwar U.S. Business Cycles: An Empirical Investigation. *Journal of Money, Credit, and Banking*, 29, 1–16.
- King, R. G. and Lu, Y. K. (2022). Evolving Reputation for Commitment: The Rise, Fall and Stabilization of US Inflation. National Bureau of Economic Research, Inc, NBER Working Papers, No. 30763.
- Kreps, D. M. and Wilson, R. (1982). Reputation and imperfect information. *Journal of Economic Theory*, 27(2), 253–279.
- Kydland, F. E. and Prescott, E. C. (1977). Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, 85(3), 473–491.
- Loisel, O. (2008). Central bank reputation in a forward-looking model. *Journal of Economic Dynamics and Control*, 32(11), 3718–3742.

- Lu, Y. K., King, R. G., and Pasten, E. (2016). Optimal reputation building in the New Keynesian model. *Journal of Monetary Economics*, 84, 233–249.
- Lucas, R. E. (1987). Models of business cycles. Oxford: Basil Blackwell.
- Marcet, A. and Marimon, R. (2019). Recursive Contracts. *Econometrica*, 87(5), 1589–1631.
- Maskin, E. and Tirole, J. (2001). Markov Perfect Equilibrium I. Observable Actions. *Journal of Economic Theory*, 100, 191–219.
- McKay, A., Nakamura, E., and Steinsson, J. (2016). The Discounted Euler Equation: A Note. *NBER Working paper*.
- McKay, A., Nakamura, E., and Steinsson, J. (2016). The Power of Forward Guidance Revisited. *American Economic Review*, 106(10), 3133–58.
- Michelacci, C. and Paciello, L. (2019). Ambiguous Policy Announcements. *The Review of Economic Studies*, 87(5), 2356–2398.
- Milgrom, P. and Roberts, J. (1982). Predation, reputation, and entry deterrence. *Journal of Economic Theory*, 27(2), 280–312.
- Nakata, T. (2018). Reputation and liquidity traps. *Review of Economic Dynamics*, 28, 252–268.
- Nakata, T., Ogaki, R., Schmidt, S., and Yoo, P. (2019). Attenuating the forward guidance puzzle: Implications for optimal monetary policy. *Journal of Economic Dynamics and Control*, 105, 90–106.
- Stokey, N. L. (1991). Credible public policy. *Journal of Economic Dynamics and Control*, 15(4), 627–656.
- Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2), 177–181.
- Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.
- Walsh, C. E. (2018). Simple Sustainable Forward Guidance at the ELB. *Working paper*.
- Werning, I. (2015). Incomplete Markets and Aggregate Demand. National Bureau of Economic Research, Working Paper, No. 21448.
- Woodford, M. (2003). Interest and prices. Princeton Univ. Press.

Appendices

3.A. MAIN PROOFS

Proof of Proposition 3.1. The myopic central banker only compares the static loss $\lambda y_t^2 + \pi_t^2$. If she breaks her promise, output and inflation are given by:

$$\begin{aligned}\pi_D &= \beta \pi_D^e + \kappa y_D \\ 0 &= \lambda y_D + \kappa \pi_D\end{aligned}$$

Where (π_D^e, y_D^e) are inflation and output expectations when the central banker breaks the promise (they are fixed once the promise is broken). The IS curve is irrelevant, because the shock is absorbed by the nominal interest rate (and this can be done as we assume the ZLB is not binding). Solving for the central banker's first-order condition and Phillips Curve:

$$\begin{aligned}\pi_D &= \pi_D^e \frac{\beta \lambda}{\lambda + \kappa^2} \\ y_D &= -\pi_D^e \frac{\beta \kappa}{\lambda + \kappa^2}\end{aligned}$$

The static loss is:

$$\ell_D = \lambda y_D^2 + \pi_D^2 = \frac{\lambda \beta^2}{\lambda + \kappa^2} (\pi_D^e)^2$$

The loss decreases to zero as $\lambda \rightarrow 0$. The nominal interest rate should be positive:

$$i_D = r + \pi_D^e + \rho y_D^e + \frac{\rho \beta \kappa}{\lambda + \kappa^2} \pi_D^e > 0$$

If she keeps her promise, output and inflation are given by:

$$\begin{aligned}y_Z &= y_Z^e - \frac{1}{\rho}(0 - r - \pi_Z^e) \\ \pi_Z &= \beta \pi_Z^e + \kappa y_Z\end{aligned}$$

Where (π_Z^e, y_Z^e) are inflation and output expectations when the central banker keeps the promise.

Solving:

$$\begin{aligned} y_Z &= y_Z^e + \frac{\pi_Z^e}{\rho} + \frac{r}{\rho} \\ \pi_Z &= \beta\pi_Z^e + \kappa y_Z^e + \frac{\kappa}{\rho}\pi_Z^e + \frac{\kappa}{\rho}r \end{aligned}$$

So the loss is:

$$\ell_Z = \lambda \left(y_Z^e + \frac{\pi_Z^e}{\rho} + \frac{r}{\rho} \right)^2 + \left(\beta\pi_Z^e + \kappa y_Z^e + \frac{\kappa}{\rho}\pi_Z^e + \frac{\kappa}{\rho}r \right)^2$$

This loss does not vanish for $\lambda \rightarrow 0$:

$$\lim_{\lambda \rightarrow 0} \ell_Z = \left(\beta\pi_Z^e + \kappa y_Z^e + \frac{\kappa}{\rho}\pi_Z^e + \frac{\kappa}{\rho}r \right)^2$$

which is generically positive.

Finally, the map $\ell_D(\lambda) - \ell_Z(\lambda)$ is continuous in λ , so $\ell_D(0) < \ell_Z(0)$ guarantees that there exists some $\lambda^* > 0$ such that $\ell_D(\lambda) < \ell_Z(\lambda)$ for all $\lambda \in [0, \lambda^*)$. $\square$

Proof of Proposition 3.2.

Consider the loss difference between keeping the promise and breaking it at state (r_t^n, μ_t, P_t) :

$$(\ell_t^K - \ell_t^D) + \delta(L_{t+1}^K - L_{t+1}^D).$$

Since continuation values are well defined and the state space is compact, there exists a constant $C < \infty$ such that

$$\delta(L_{t+1}^D - L_{t+1}^K) \leq C$$

for all states.

Under promise-keeping, $i_t = 0$. For sufficiently high realizations of r_t^n , the resulting output gap and inflation are sufficiently expansionary that

$$\ell_t^K - \ell_t^D > C.$$

Therefore, for such realizations,

$$\ell_t^K + \delta L_{t+1}^K > \ell_t^D + \delta L_{t+1}^D,$$

so breaking the promise is strictly optimal.

Hence there exists a sufficiently high realization $\hat{R}$ such that for all $r_t^n \geq \hat{R}$, all μ_t , all P_t , and both $\delta \in \{0, \beta\}$,

$$\sigma(\delta, r_t^n, \mu_t, P_t) = 0.$$

$\square$

Proof. Consider the loss difference between keeping the promise and breaking it at state (r_t^n, μ_t, P_t) :

$$\left(\ell_t^K + \delta L_{t+1}^K\right) - \left(\ell_t^D + \delta L_{t+1}^D\right).$$

Since r_t^n has finite support, per-period losses are finite in every state, and $\delta < 1$, continuation losses are uniformly bounded. Hence there exists a constant $C < \infty$ such that

$$\delta(L_{t+1}^D - L_{t+1}^K) \leq C$$

for all states.

Under promise-keeping, $i_t = 0$. Therefore, for realizations of r_t^n sufficiently high in the support, the current-period loss from promise-keeping exceeds the current-period loss from discretion by more than C , that is,

$$\ell_t^K - \ell_t^D > C.$$

For all such realizations, breaking the promise is strictly optimal.

Hence there exists a sufficiently high realization $\hat{R}$ such that for all $r_t^n \geq \hat{R}$, all μ_t , all P_t , and both $\delta \in \{0, \beta\}$,

$$\sigma(\delta, r_t^n, \mu_t, P_t) = 0.$$

□

Proof of Proposition 3.3.

Given the pooling strategies and the off-path beliefs, the private sector assigns probability zero to promises being kept. Thus, promise-keeping induces no change in expected output gap and inflation, implying that there are no gains from deviating. □

Proof of Proposition 3.4.

Consider the first promise-keeping state reached after exiting the ZLB. At that state, the forward-looking type keeps the promise only if

$$\ell_t^{K,\tau} + \beta L_{t+1}^{K,\tau} \leq \ell_t^D + \beta L_{t+1}^D.$$

Equivalently,

$$\ell_t^{K,\tau} - \ell_t^D \leq \beta(L_{t+1}^D - L_{t+1}^{K,\tau}).$$

Since one-period losses are uniformly bounded under discretion and $\beta < 1$, continuation losses are uniformly bounded. Hence there exists a constant $C < \infty$ such that

$$\beta(L_{t+1}^D - L_{t+1}^{K,\tau}) \leq C$$

for all states and all τ .

Under promise-keeping, a longer promised duration τ implies a longer anticipated path of zero nominal interest rates. Therefore, at the first promise-keeping state, output and inflation under promise-keeping

become more expansionary as τ increases, so that the current loss difference $\ell_t^{K,\tau} - \ell_t^D$ increases with τ . Hence there exists $\bar{\tau} \in \mathbb{N}$ such that for all $\tau > \bar{\tau}$,

$$\ell_t^{K,\tau} - \ell_t^D > C.$$

For all such τ , the forward-looking type strictly prefers to break the promise already at the first promise-keeping state. Therefore, for all $\tau > \bar{\tau}$, no equilibrium can sustain promise-keeping. $\square$

Proof of Corollary 3.1.

Follows from Proposition 3.4. $\square$

3.B. INTEREST RATE UNDER OPTIMAL DISCRETION

When the central bank is not honoring a promise⁷, the interest rate is set to minimize the instantaneous loss function by taking expectations as given. We impose the constraint on the minimization problem that the interest rate cannot be negative. The interest rate chosen by the central banker will therefore be

$$i_t = \arg \min \ell(y_t, \pi_t)$$

subject to

$$\begin{aligned} \pi_t &= \beta \mathbb{E}_t[\pi_{t+1}] + \kappa y_t \\ y_t &= \mathbb{E}_t[y_{t+1}] - \frac{1}{\rho}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n) \\ i_t &\geq 0 \end{aligned}$$

Using the IS curve and the Phillips curve, we can rewrite the loss function of the central bank in terms of the nominal interest rate rather than inflation and output gap:

$$\ell(y_t, \pi_t) = \ell \left(\mathbb{E}_t[y_{t+1}] - \frac{1}{\rho}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n), \beta \mathbb{E}_t[\pi_{t+1}] + \kappa \mathbb{E}_t[y_{t+1}] - \kappa \frac{1}{\rho}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n) \right)$$

We denote by $\zeta_t \leq 0$ the multiplier associated to the constraint $i_t \geq 0$. The Lagrangian associated to the minimization problem of the central bank writes:

$$\begin{aligned} \mathcal{L}(i_t, \zeta_t) &= \left(\beta \mathbb{E}_t[\pi_{t+1}] + \kappa \mathbb{E}_t[y_{t+1}] - \kappa \frac{1}{\rho}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n) \right)^2 \\ &\quad + \lambda \left(\mathbb{E}_t[y_{t+1}] - \frac{1}{\rho}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n) - \bar{y} \right)^2 \\ &\quad + \zeta_t i_t \end{aligned}$$

⁷Either because no promise has been given or because the central bank decides to break it.

The first order condition is:

$$2\frac{\kappa}{\rho} \underbrace{\left(\beta \mathbb{E}_t[\pi_{t+1}] + \kappa \mathbb{E}_t[y_{t+1}] - \kappa \frac{1}{\rho} (i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n) \right)}_{=\pi_t} + 2\frac{\lambda}{\rho} \underbrace{\left(\mathbb{E}_t[y_{t+1}] - \frac{1}{\rho} (i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n) \right)}_{=(y_t - \bar{y})} = \zeta_t$$

$$\iff 2\frac{\kappa}{\rho} \pi_t + 2\frac{\lambda}{\rho} (y_t - \bar{y}) = \zeta_t$$

So that substituting back π_t and y_t we obtain

$$\kappa \pi_t + \lambda (y_t - \bar{y}) = \frac{\rho}{2} \zeta_t \leq 0$$

Whenever the ZLB is binding, the multiplier ζ_t will generically be strictly negative. On the contrary, whenever the ZLB is slack, the multiplier ζ_t is exactly zero, and the FOC writes:

$$\lambda (y_t - \bar{y}) + \kappa \pi_t = 0$$

The second order condition is:

$$2\frac{\kappa^2}{\rho^2} + 2\frac{\lambda}{\rho^2} > 0$$

So that the FOC characterizes the unique global minimum.

3.C. EQUILIBRIUM DESCRIPTION

We describe here in detail how the equilibrium is computed. For fixed strategies σ and beliefs μ , output gap, inflation and interest rates must constitute a solution to the New Keynesian model. The key differences are that interest rates are set to zero whenever a promise is kept, and that the expectations about future outcomes depend on the private sector assessment about the type of the central bank and how they behave.

3.C.1. At the ZLB

Whenever the economy is at the ZLB at time t , the nominal interest rate is set to zero, $i_t^Z = 0$. An Odyssean promise is given, and therefore the private sector forms an assessment about how likely it is that the promise will be kept. We have to consider three possibilities:

1. If the ZLB is binding again in the next period, then the promise is renewed. This corresponds to the event $\{r_{t+1}^n \in Z_{\sigma, \mu}\}$.
2. If the ZLB is not binding at time $t+1$, then each type δ of the central bank will keep the promise with probability $\sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)$. As the type of the central bank is private information, the private sector forms an expectation with respect to δ using the probability measure μ_t . Thus

conditional on exiting the ZLB, i.e. $r_{t+1}^n \notin Z_{\sigma,\mu}$, the expected probability of transitioning to the first promise-keeping period is $\mathbb{E}_\delta^{\mu_t}[\sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)]$.

3. Finally, if the ZLB is not binding at time $t + 1$, the conditional probability of transitioning to normal times is $\mathbb{E}_\delta^{\mu_t}[1 - \sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)]$.

Output gap, inflation and interest rates therefore satisfy the following system:

$$\left\{ \begin{array}{l} \pi_t^Z = \beta \mathbb{E}_{r_{t+1}^n} \left[\begin{array}{l} \pi_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \\ + \pi_{t+1}^P \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_\delta^{\mu_t}[\sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)] \\ + \pi_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_\delta^{\mu_t}[1 - \sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)] \end{array} \right] + \kappa y_t^Z \\ \\ y_t^Z = \mathbb{E}_{r_{t+1}^n} \left[\begin{array}{l} y_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \\ + y_{t+1}^P \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_\delta^{\mu_t}[\sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)] \\ + y_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_\delta^{\mu_t}[1 - \sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)] \end{array} \right] \\ \\ -\frac{1}{\rho} \left(0 - \mathbb{E}_{r_{t+1}^n} \left[\begin{array}{l} \pi_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \\ + \pi_{t+1}^P \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_\delta^{\mu_t}[\sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)] \\ + \pi_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_\delta^{\mu_t}[1 - \sigma(\delta, r_{t+1}^n, \mu_t, P_t = 1)] \end{array} \right] - r_t^n \right) \\ \\ i_t^Z = 0 \end{array} \right.$$

For the ease of exposition we have suppressed the dependence on the current shock realization. A similar omission applies to the remaining cases.

3.C.2. Promise-keeping period with $P_t < \tau$

If the economy is away from the ZLB and a promise is to be kept at time t , the interest rate is set to zero, $i_t^P = 0$. In the next period, the ZLB could be binding, so a promise is renewed and the process starts over. If the economy stays away from the ZLB for one more period, a transition to the next period of promise keeping occurs with conditional probability $\mathbb{E}_\delta^{\mu_t}[\sigma(\delta, r_{t+1}^n, \mu_t, P_t + 1)]$. The promise is broken in next period with conditional probability $\mathbb{E}_\delta^{\mu_t}[1 - \sigma(\delta, r_{t+1}^n, \mu_t, P_t + 1)]$.

Output gap, inflation and interest rates therefore satisfy the following system:

$$\left\{ \begin{array}{l} \pi_t^P = \beta \mathbb{E}_{r_{t+1}^n} \left[\begin{array}{l} \pi_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \\ + \pi_{t+1}^P \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_{\delta}^{\mu_t} [\sigma(\delta, r_{t+1}^n, \mu_t, P_t + 1)] \\ + \pi_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_{\delta}^{\mu_t} [1 - \sigma(\delta, r_{t+1}^n, \mu_t, P_t + 1)] \end{array} \right] + \kappa y_t^P \\ \\ y_t^P = \mathbb{E}_{r_{t+1}^n} \left[\begin{array}{l} y_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \\ + y_{t+1}^P \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_{\delta}^{\mu_t} [\sigma(\delta, r_{t+1}^n, \mu_t, P_t + 1)] \\ + y_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_{\delta}^{\mu_t} [1 - \sigma(\delta, r_{t+1}^n, \mu_t, P_t + 1)] \end{array} \right] \\ \\ - \frac{1}{\rho} \left(0 - \mathbb{E}_{r_{t+1}^n} \left[\begin{array}{l} \pi_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \\ + \pi_{t+1}^P \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_{\delta}^{\mu_t} [\sigma(\delta, r_{t+1}^n, \mu_t, P_t + 1)] \\ + \pi_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} \cdot \mathbb{E}_{\delta}^{\mu_t} [1 - \sigma(\delta, r_{t+1}^n, \mu_t, P_t + 1)] \end{array} \right] - r_t^n \right) \\ \\ i_t^P = 0 \end{array} \right.$$

3.C.3. Last promise-keeping period ($P_t = \tau$)

If the economy is in the last promise-keeping period the interest rate is $i_t^P = 0$. There are only two possible cases in the following period: either the ZLB binds, or the ZLB does not bind and the central bank conducts discretionary policy. Output gap, inflation and interest rates therefore satisfy the following system:

$$\left\{ \begin{array}{l} \pi_t^P = \beta \mathbb{E}_{r_{t+1}^n} \left[\pi_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} + \pi_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \right] + \kappa y_t^P \\ \\ y_t^P = \mathbb{E}_{r_{t+1}^n} \left[y_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} + y_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \right] \\ \\ - \frac{1}{\rho} \left(0 - \mathbb{E}_{r_{t+1}^n} \left[\pi_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma,\mu}\} + \pi_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma,\mu}\} \right] - r_t^n \right) \\ \\ i_t^P = 0 \end{array} \right.$$

3.C.4. Normal times

Away from the ZLB and with no promise to be kept, the central bank conducts discretionary policy. Thus, the interest rate is set to satisfy the derived FOC:

$$-\kappa \pi_t^D = \lambda (y_t^D - \bar{y})$$

In the following period, the economy may either hit the ZLB or stay in discretion. Output gap, inflation and interest rates therefore satisfy the following system:

$$\begin{cases} \pi_t^D = \beta \mathbb{E}_{r_{t+1}^n} \left[\pi_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z(\sigma)\} + \pi_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z(\sigma)\} \right] + \kappa y_t^D \\ y_t^D = \mathbb{E}_{r_{t+1}^n} \left[y_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z(\sigma)\} + y_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z(\sigma)\} \right] \\ \quad - \frac{1}{\rho} \left(i_t^D - \mathbb{E}_{r_{t+1}^n} \left[\pi_{t+1}^D \cdot \mathbb{1}\{r_{t+1}^n \notin Z(\sigma)\} + \pi_{t+1}^Z \cdot \mathbb{1}\{r_{t+1}^n \in Z(\sigma)\} \right] - r_t^n \right) \\ -\kappa \pi_t^D = \lambda (y_t^D - \bar{y}) \end{cases}$$

3.C.5. Decisions of the central bank

Once we have computed output gap, inflation and interest rates for fixed (σ, μ) we need to verify that the strategies σ are optimal.

When the economy is at the ZLB, the strategies σ are optimal if they solve the following Bellman equation:

$$\begin{aligned} L_t^Z(\delta, r_t^n, \mu_t, 0) &= \ell \left(y_t^Z(r_t^n, \mu_t, 0), \pi_t^Z(r_t^n, \mu_t, 0) \right) \\ &+ \delta \mathbb{E}_{r_{t+1}^n} \left[\min_{x \in \{0,1\}} \left\{ x L_{t+1}^P(\delta, r_{t+1}^n, \mu_{t+1}, 1) \right. \right. \\ &\quad \left. \left. + (1-x) L_{t+1}^D(\delta, r_{t+1}^n, \mu_{t+1}, 0) \right\} \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma, \mu}\} \right. \\ &\quad \left. + L_{t+1}^Z(\delta, r_{t+1}^n, \mu_{t+1}, 0) \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma, \mu}\} \right] \end{aligned}$$

That is, the optimal action prescribes whether to transit to either discretion (break the promise) or promise-keeping. If the economy is in a promise keeping period with $P_t < \tau$, optimal strategies solve:

$$\begin{aligned} L_t^P(\delta, r_t^n, \mu_t, P_t) &= \ell \left(y_t^P(r_t^n, \mu_t, P_t), \pi_t^P(r_t^n, \mu_t, P_t) \right) \\ &+ \delta \mathbb{E}_{r_{t+1}^n} \left[\min_{x \in \{0,1\}} \left\{ x L_{t+1}^P(\delta, r_{t+1}^n, \mu_{t+1}, P_t + 1) \right. \right. \\ &\quad \left. \left. + (1-x) L_{t+1}^D(\delta, r_{t+1}^n, \mu_{t+1}, 0) \right\} \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma, \mu}\} \right. \\ &\quad \left. + L_{t+1}^Z(\delta, r_{t+1}^n, \mu_{t+1}, 0) \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma, \mu}\} \right] \end{aligned}$$

When the economy is in the last promise-keeping period, no decision is taken with respect to the transition, so the Bellman equation writes:

$$L_t^P(\delta, r_t^n, \mu_t, \tau) = \ell \left(y_t^P(r_t^n, \mu_t, \tau), \pi_t^P(r_t^n, \mu_t, \tau) \right) + \delta \mathbb{E}_{r_{t+1}^n} \left[L_{t+1}^D(\delta, r_{t+1}^n, \mu_{t+1}, 0) \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma, \mu}\} \right. \\ \left. + L_{t+1}^Z(\delta, r_{t+1}^n, \mu_{t+1}, 0) \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma, \mu}\} \right]$$

Finally, if the economy is in discretion, no strategic decision is made, so the Bellman equation is:

$$L_t^D(\delta, r_t^n, \mu_t, 0) = \ell \left(y_t^D(r_t^n, \mu_t, 0), \pi_t^D(r_t^n, \mu_t, 0) \right) + \delta \mathbb{E}_{r_{t+1}^n} \left[L_{t+1}^D(\delta, r_{t+1}^n, \mu_{t+1}, 0) \cdot \mathbb{1}\{r_{t+1}^n \notin Z_{\sigma, \mu}\} \right. \\ \left. + L_{t+1}^Z(\delta, r_{t+1}^n, \mu_{t+1}, 0) \cdot \mathbb{1}\{r_{t+1}^n \in Z_{\sigma, \mu}\} \right]$$

3.D. FORWARD GUIDANCE AND THE VARIANCE OF THE DEMAND SHOCK

We vary the standard deviation (ς) of the innovation ξ_t to assess its effect on credibility and power of forward guidance.

In Figure 3.8 we show the effect on the output gap at the ZLB of promising a zero interest rate for a time horizon of $\tau = 1$ periods as a function of the standard deviation of the innovation to the demand shock. We also show the effect of forward guidance with full commitment in red for comparison.




Figure 3.8: Forward guidance and the standard deviation of the innovation to the demand shock

We see that there is a range of values for the standard deviation ς for which FG can be sustained in equilibrium. In particular, we notice that for a standard deviation of less than 0.007, forward guidance cannot be sustained in equilibrium, as well as for a standard deviation of more than 0.023.

Conditional on being sustainable, we find that the higher the standard deviation, the smaller the effect of FG. This occurs as we consider promises of keeping the rate at zero, and interest rates are already low for high levels of the standard deviation, as can be seen from Figure 3.9. Thus, the promise is less expansionary by construction.




Figure 3.9: Nominal interest rates with medium realizations of the demand shock.