Games and Economic Behavior 95 (2016) 137–155

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Bounds on the welfare loss from moral hazard with limited liability [☆]

Felipe Balmaceda [a], Santiago R. Balseiro [b,*], José R. Correa [c], Nicolás E. Stier-Moses [d,e]

[a] Economics Department, Diego Portales University, Chile
[b] Fuqua School of Business, Duke University, USA
[c] Industrial Engineering Department, Universidad de Chile, Chile
[d] Graduate School of Business, Columbia University, USA
[e] School of Business, Universidad Torcuato Di Tella and CONICET, Argentina

ARTICLE INFO

Article history:
Received 26 October 2012
Available online 23 October 2015

JEL classification:
J41
J24
D21

Keywords:
Principal–agent problem
Moral hazard
Limited liability
Welfare loss
Price of anarchy

ABSTRACT

We study a principal–agent problem with discrete outcome and effort level spaces. The principal and the agent are risk neutral and the latter is subject to limited liability. Quantifying welfare loss as the ratio between the first-best social welfare and that arising from the principal's optimal pay-for-performance contract, we provide simple parametric bounds for problem instances with moral hazard. Relying on that, we compute the worst-case welfare loss ratio among all problem instances with a fixed number of effort and outcome levels as a function of the number of possible effort levels and the likelihood ratio evaluated at the highest outcome. As extensions, we look at linear contracts and at cases with multiple identical tasks. Our work constitutes an initial attempt to quantify the losses arising from moral hazard when the agent is subject to limited liability, and shows that these losses are non-negligible in the worst case.

1. Introduction

The principal–agent model with moral hazard has been the workhorse paradigm to understand many interesting economic phenomena where incentives play a crucial role, such as the theory of insurance under moral hazard (Spence and Zeckhauser, 1971), the theory of managerial firms (Alchian and Demsetz, 1972; Jensen and Meckling, 1979), optimal sharecropping contracts between landowners and tenants (Stiglitz, 1974), the efficiency wages theory (Shapiro and Stiglitz, 1984), financial contracting (Holmström and Tirole, 1997; Innes, 1990), and job design and multi-tasking (Holmström and Milgrom, 1991). Casual observation also suggests that moral hazard could be of practical importance. In fact, most sales workers are paid according to a fixed wage and either a bonus paid when a certain sales target is achieved or a commission rate over

☆ We greatly appreciate the comments of two anonymous referees that helped us refine the content and significantly improve the presentation of the paper. This work was partially supported by FONDECYT Chile through grants 1140140 and 1130671, by Instituto Milenio para la Investigación de Imperfecciones de Mercado y Políticas Públicas grant ICM IS130002, by the Instituto Milenio Sistemas Complejos de Ingeniería, by the Millenium Nucleus Information and Coordination in Networks ICM/FIC RC130003, and by CONICET Argentina Grant PIP 112-201201-00450CO, Convenio Cooperación CONICET–CONICYT Resolución 1562/14, and FonCyT Argentina Grant PICT 2012-1324.

* Corresponding author at: 100 Fuqua Drive, Fuqua School of Business, Duke University, Durham, NC, 27708.
E-mail addresses: felipe.balmaceda@udp.cl (F. Balmaceda), srb43@duke.edu (S.R. Balseiro), jcorrea@dii.uchile.cl (J.R. Correa), stier@gsb.columbia.edu (N.E. Stier-Moses).

http://dx.doi.org/10.1016/j.geb.2015.10.008
0899-8256/

total sales. Franchisees are also motivated by contracts that entail a fixed payment and an agreement about how to share profits or sales. Additionally, managerial contracts usually consist of a combination of fixed wages and payments that are conditioned on performance. In short, incentive contracts are ubiquitous to market economies.

Regardless of the reason for moral hazard, in most cases it entails welfare losses that remain as far as we know quantitatively uncounted for. This paper begins the task of quantifying the welfare losses implied by the existence of moral hazard in a principal–agent relationship with risk neutral individuals and limited liability. The main consequences of moral hazard are by now well understood and deeply rooted in the economics of information literature, thus the moral-hazard paradigm is ripe for a deeper analysis of the quantitative, rather than qualitative, consequences of it.

The setting we consider consists of a risk-neutral principal who hires a risk-neutral agent subject to limited liability to exert costly effort. The effort level and outcome space are discrete and finite, and effort influences the distribution of output. Because the agent's effort is not observable, the principal can only design contracts based on the agent's observable performance. If the principal wishes to induce the agent to choose a given effort level, he should reward the agent when the realization of output is most indicative of the desired effort level having been chosen, and he should punish him when a different outcome is observed. We assume throughout that the probability distribution of output, which is parameterized by the effort level, satisfies the *monotone likelihood ratio property* (**MLRP**).[1] This implies that the highest output is the most indicative that the highest effort level has been chosen and therefore the principal pays the agent a salary in excess of the limited-liability amount only when the highest outcome is observed, and pays the limited-liability amount otherwise. Because limited liability imposes a lower bound on the size of the punishment, the equilibrium contract leaves a limited-liability rent to the agent. As a result, the equilibrium contract might not maximize social welfare and the first-best outcome might not be attained; instead, the constrained contract will be second-best.[2]

Before going into the main results it is worthwhile to notice that the contract considered here is quite implausible in the empirically relevant case of a continuous-outcome space. Under the assumptions considered and a continuous-outcome space, the principal pays a bonus only when the highest outcome takes place, which is a measure-zero event. The incentive compatibility will require the principal to pay an infinite bonus in this state, which is unrealistic. This unappealing feature could be avoided, for example, by introducing risk aversion or by assuming that agents will behave perversely to reward systems that are not properly monotonic. The model considered here is still useful as a first step towards a more general study of the welfare loss arising from moral hazard in settings that are more plausible in the real world.

Finally, the model here can be interpreted in a very intuitive way as a situation where the principal is a monopsonist and the agent is a multi-product firm supplying the monopsonist.[3] In this context, states of nature correspond to potential markets, probabilities correspond to quantities sold in each market and the principal-monopsonist restrict itself to use linear-pricing contracts. The agent-supplier's outputs in the potential markets are joint outputs determined by his single effort level, and his production function is given by the vector of probability distribution conditional on the effort level chosen. Under this interpretation, the optimal contract follows from the standard intermediate microeconomics analysis of a monopsonist where the agent-supplier is a competitive price-taking firm and the principal-monopsonist pays the supplier a price equal to the marginal cost corresponding to the principal's target effort level. Compared with the case of a competitive input market, inefficiency is introduced because the principal-monopsonist buys a positive quantity from just one market and this is lower than the welfare-maximizing quantity. Therefore our results also quantify the efficiency loss of the competitive input market relative to a monopsonized input market.

1.1. Main contributions

To quantify the inefficiencies introduced by moral hazard and limited liability, we measure the welfare loss introduced by a given contract, and rely on the concept of *price of anarchy*. The latter refers to the worst-case welfare loss in a non-cooperative game, that is, the worst possible equilibrium welfare versus that of a socially-optimal solution. The idea of using worst-case analysis to study situations under competition was introduced by Koutsoupias and Papadimitriou (1999) and has gained followers over the last decade. The use of the price of anarchy as a metric of the welfare loss has been widely applied in economics to problems such as the study of competition and efficiency in congested markets (Acemoglu and Ozdaglar, 2007), games with serial, average and incremental cost sharing (Moulin, 2008), price and capacity competition (Acemoglu et al., 2009), Vickrey–Clarke–Groves mechanisms (Moulin, 2009), resource allocation problems (Kelly, 1997; Johari and Tsitsiklis, 2004), and congestion games (Roughgarden and Tardos, 2004; Correa et al., 2008). In our setting, we define the welfare-loss ratio as the ratio between the social welfare of a socially-optimal solution—the sum of the principal's and agent's payoffs when the first-best effort level is chosen—and that of the sub-game perfect equilibrium in which the principal offers the agent a performance-pay contract and then the agent chooses the effort level.

The goal of this paper is two-fold. First, we provide simple parametric bounds on the welfare-loss ratio in a given instance of the problem in the presence of limited-liability and moral hazard. These bounds allow one to directly quantify

[1] The **MLRP** assumption is pervasive in the principal–agent literature (Grossman and Hart, 1983; Rogerson, 1985).

[2] When the participation constraint—rather than the limited-liability constraint—binds, providing incentives is costless since the agent cares only about the expected compensation and the participation constraint binds on his expected payoff. Thus, we focus on the case in which the parameters are such that this does not occur.

[3] We thank an anonymous referee for this novel interpretation of the model.

the inefficiency in a given instance of the problem without the need to determine its first-best and second-best effort levels, and additionally, they shed some light on the structure of those problem instances with high welfare loss. Second, we study the worst-case welfare-loss ratio (i.e., the price of anarchy), which is defined as the largest welfare-loss ratio among all instances of the problem that satisfy our assumptions. In our model, an instance of the problem is given by the outcome vector, the vector of agent's costs of effort, and the probability distribution of outcomes for each effort level.

In order to obtain our bounds for the welfare-loss ratio, we assume throughout that the probability distribution of output, which is parameterized by the effort level, satisfies **MLRP**—under this property, a higher output is a better signal that the agent has chosen a higher effort level—and that the marginal cost per unit increase in the highest outcome probability is non-decreasing in the effort level (**IMCP**).[4] The latter assumption ensures that local incentive compatibility constraints are sufficient to induce an effort level, and is weaker than the well-known *convexity of the distribution function condition* (**CDFC**) that is also related to the idea of decreasing marginal returns to effort (see, for instance, Rogerson, 1985 and Mirrlees, 1999). Furthermore, it is assumed that the sequence of prevailing social welfare levels, as effort levels increase, is quasi-concave (**QCSW**); i.e., social welfare is single-peaked in the effort level.

Under our assumptions, inefficiency is introduced when the socially-optimal effort level is high and the principal finds it optimal to induce a low effort level. This happens when the net social gain of increasing the expected output is smaller than the rent given to the agent when inducing a higher effort level. Inducing a higher effort level is costly for the principal because of the agent's limited-liability constraint; the principal cannot severely punish the agent when a bad outcome is realized. Therefore the principal is restricted to rewarding the agent when a good outcome is realized, which may not be enough to compensate for the opportunity cost incurred when a bad outcome is observed.

An unbounded welfare loss would arise if the social welfare of the high effort level were arbitrarily larger than that of the low effort level and, regardless, the principal prefers to induce the lower effort level. Because social welfare is the sum of the principal's and the agent's utilities, the latter implies that the agent is capturing most of the social welfare at the high effort level. Our results preclude an arbitrarily large welfare loss because the principal's utility at the high level cannot be arbitrarily small when the social inefficiency at the low effort level is arbitrarily large. Under assumptions **MLRP**, **IMCP** and **QCSW**, we establish that for any instance of the problem, the welfare-loss ratio is bounded from above by a simple formula involving the probabilities of the highest possible outcome. The results arise from the fact that **MLRP** implies that the principal pays a bonus only when the highest outcome is observed.[5]

Subsequently, we show that the worst-case welfare-loss ratio among all problem instances with a fixed number of effort and outcome levels that satisfy our assumptions is equal to the number of effort levels E. As a consequence, the social welfare of a subgame perfect equilibrium is guaranteed to be at least $1/E$ of that of the social optimum. We prove that the worst-case is attained by a family of problem instances in which the highest outcome probability increases at a geometric rate with the effort level. A disadvantage of the previous bound is that it grows unboundedly with the number of effort levels. Thus, we study the worst-case welfare-loss ratio among problem instances with an arbitrary number of effort levels and the likelihood ratio of the highest outcome bounded from above by $r \geq 1$. Here, r is defined as the ratio of the highest outcome probability when the highest effort level is exerted to that when the lowest effort level is exerted. We show that the worst-case welfare-loss ratio within the previous family is equal to $1 + \ln(r)$ independently of the number of effort levels. Our results suggest that moral hazard is more problematic in situations where the agent's available actions are more numerous, or when the informational problem is such that the highest outcome is much more likely under the highest than under the lowest effort level.

As an extension to the basic model, we study the welfare loss when contracts are restricted to be linear. This is motivated by the prevalence of simple contracts in real life (see, e.g., Salanié, 2003). Surprisingly, we show that a similar bound on the welfare loss holds when the principal is restricted to choose linear contracts, and that the worst-case welfare-loss ratio is again equal to the number of effort levels E. Our results provide bounds on the welfare loss but do not shed light on whether, for the same instance of the problem, the restriction to linear contracts increases or decreases welfare loss. In another extension, we study the welfare loss when there are multiple identical and independent tasks, and for each task the agent chooses between two effort levels. We give a simple bound on the welfare loss involving the probabilities that all tasks are successful when the agent exerts a high and a low effort level, and find that the worst-case welfare-loss ratio is 2, regardless of how many tasks the agent has to work on.

In related work, Demougin and Fluet (1998) derive the optimality of the bonus contract in the setting considered here; i.e., with a discrete outcome space. Kim (1997) presents conditions that guarantee the existence of a bonus contract that achieves a first-best allocation under limited liability but for a continuous outcome and effort space. Our results extend this work by quantifying the impact on efficiency when the first-best allocation is not achieved. In this context, our work should be viewed as a preliminary step in a broader agenda of how to quantify the welfare loss of moral hazard in different settings. Along those lines, Babaioff et al. (2009, 2012) study a principal–agent problem with an approach similar to ours. They introduce a combinatorial agency problem with multiple agents performing two-effort–two-outcome tasks, and study the combinatorial structure of dependencies between agents' actions and the worst-case welfare loss for a number of different

[4] IMCP stands for *Increasing Marginal Cost of Probability*. In the monopsonist-supplier interpretation of the principal–agent problem, assumption **IMCP** simply requires that the agent-supplier's marginal cost of output in the highest gross-revenue market is increasing in the effort level.

[5] In addition to this, the extended version of the current paper shows that our results are robust by relaxing some assumptions of the model presented here (Balmaceda et al., 2012).

classes of action dependencies. They show that this loss may be unbounded for technologies that exhibit complementarities between agents, while it can be bounded from above by a small constant for technologies that exhibit substitutabilities between agents. In contrast, our model deals with a single agent and its complexity lies in handling more sophisticated tasks, rather than the interaction between tasks.

The rest of the paper is organized as follows. In Section 2, we introduce the model with its main assumptions and present some preliminary results that will prove useful in the rest of the paper. Section 3 presents our main results on bounds for the welfare loss. Section 4 extends our results in several directions, while Section 5 concludes with some remarks and future directions of study.

2. The principal–agent model

2.1. The basic setup

We consider a risk-neutral principal and a risk-neutral agent in a setting with $E \geq 2$ effort levels and $S \geq 2$ outcomes.[6] The agent chooses an effort level $e \in \mathcal{E} \triangleq \{1, \ldots, E\}$ and incurs a personal nonnegative cost of c_e. Effort levels are sorted in increasing order with respect to costs; that is, $c_e < c_f$ if $e < f$. Thus, a higher effort level demands more work from the agent. We denote by $\mathbf{c} = (c_1, \ldots, c_E)$ the vector of agent's costs. The task's output depends on a random state of nature $s \in \mathcal{S} \triangleq \{1, \ldots, S\}$ whose distribution in turn depends on the effort level chosen by the agent. Each state has an associated nonnegative dollar amount that represents the principal's revenue. We denote the vector of outputs indexed by state by $\mathbf{y} = (y^1, \ldots, y^S)$. Without loss of generality, the outputs are sorted in increasing order: $y^s < y^t$ if $s < t$; hence, the principal's revenues are higher under states with a larger index. Finally, we let $\pi_e^s > 0$ be the common-knowledge probability of state $s \in \mathcal{S}$ when the agent exerts effort level $e \in \mathcal{E}$. The probability mass function of the outcome under effort level e is given by $\boldsymbol{\pi}_e = (\pi_e^1, \ldots, \pi_e^S)$. To simplify the exposition we set $c_0 = 0$ and $\boldsymbol{\pi}_0 = \mathbf{0}$. An instance of the principal–agent problem is characterized by the tuple $I = (\boldsymbol{\pi}, \mathbf{y}, \mathbf{c})$, and we denote by $\mathbb{I}^{E,S} \subset \mathbb{R}^{E\times S} \times \mathbb{R}^S \times \mathbb{R}^E$ the set of all valid problem instances with E effort levels and S outcomes, that is, $\boldsymbol{\pi} = \{\boldsymbol{\pi}_e\}_{e=1}^E$ is the matrix of probability vectors, and the outputs $\mathbf{y}$ and costs $\mathbf{c}$ are nonnegative and increasing.

Because the agent's chosen effort level e cannot be observed by the principal, he cannot write a wage contract based on it. However, the principal can write a contract that conditions payments on the output. The timing is as follows. First, the principal makes a take-it-or-leave-it offer to the agent that specifies a state-dependent wage schedule $\mathbf{w} = (w^1, \ldots, w^S)$. The contract is subject to a *limited liability* (LL) constraint specifying that the wage must be nonnegative in every possible state. The LL constraint excludes contracts in which the agent ends up paying back to the principal. Second, the agent decides whether to accept or reject the offer, and if accepted, then he chooses the effort level $e \in \mathcal{E}$, before learning the realized state, and incurs a personal cost c_e. Third, the random state $s \in \mathcal{S}$ is realized, the agent is paid the wage w^s and the principal collects the revenue y^s. The agent should accept the contract if it satisfies an *individual rationality* (IR) constraint specifying that the contract must yield an expected utility to the agent greater than or equal to that of choosing the *outside option*, which is assumed for the sake of simplicity equal to zero.[7] Furthermore, he will choose the effort level $e \in \mathcal{E}$ that maximizes his expected payoff $\boldsymbol{\pi}_e \mathbf{w} - c_e$; that is, the difference between the expected wage and the cost of the effort level chosen.

The principal's problem consists on choosing a wage schedule $\mathbf{w}$ and an induced effort level e for the agent that solve the following problem:

$$u^{\mathrm{P}} \triangleq \max_{e\in\mathcal{E},\mathbf{w}} \boldsymbol{\pi}_e(\mathbf{y} - \mathbf{w}) \tag{1}$$

$$\text{s.t. } \boldsymbol{\pi}_e \mathbf{w} - c_e \geq 0, \tag{IR}$$

$$e \in \arg\max_{f\in\mathcal{E}} \left\{\boldsymbol{\pi}_f \mathbf{w} - c_f\right\}, \tag{IC}$$

$$\mathbf{w} \geq 0. \tag{LL}$$

The objective function measures the difference between the principal's expected revenue and payment, which is his expected utility. Constraints (IR) and (LL) were described earlier. The *incentive compatibility* (IC) constraint guarantees that the induced effort level maximizes the agent's utility. As is standard in the literature, we assume throughout the paper that when the agent is indifferent between two or more effort levels, he always picks the one preferred by the principal. Results can be extended to hold under strong IC without further gain in intuition and much more cumbersome mathematical derivations.

Following Grossman and Hart (1983), one can equivalently formulate the principal's problem as

[6] The extended version of the current paper considers other assumptions, including continuous effort levels and some other relaxations (Balmaceda et al., 2012).

[7] The extended version of the current paper generalizes these results by incorporating a nonzero outside option, and provides bounds that capture the dependence on this parameter (Balmaceda et al., 2012).

$$u^{\mathrm{P}} = \max_{e \in \mathcal{E}} \{\boldsymbol{\pi}_e \mathbf{y} - z_e\},$$

where z_e is the minimum expected payment incurred by the principal that induces the agent to exert effort level e. We denote by $u_e^{\mathrm{P}} \triangleq \boldsymbol{\pi}_e \mathbf{y} - z_e$ the principal's maximum expected utility when effort level e is induced, and by $\mathcal{E}^{\mathrm{P}}$ the set of optimal effort levels for the principal. Hence, $u^{\mathrm{P}} = \max_{e\in\mathcal{E}}\{u_e^{\mathrm{P}}\}$ and $\mathcal{E}^{\mathrm{P}} = \arg\max_{e\in\mathcal{E}}\{u_e^{\mathrm{P}}\}$. We denote by $u_e^{\mathrm{A}} \triangleq z_e - c_e$ the agent's expected utility when effort level e is induced.

Exploiting the fact that the effort level set is finite, we write the IC constraint explicitly to obtain the *minimum payment linear program* corresponding to e, which is independent of the output $\mathbf{y}$:

$$(\mathrm{MPLP}_e) \qquad z_e = \min_{\mathbf{w}\in\mathbb{R}^S} \boldsymbol{\pi}_e \mathbf{w} \tag{2}$$

$$\text{s.t. } \boldsymbol{\pi}_e \mathbf{w} - c_e \geq 0, \tag{IR}$$

$$\boldsymbol{\pi}_e \mathbf{w} - c_e \geq \boldsymbol{\pi}_f \mathbf{w} - c_f \quad \forall f \in \mathcal{E} \setminus e, \tag{IC}$$

$$\mathbf{w} \geq 0. \tag{LL}$$

We say that the principal *implements* effort level $e \in \mathcal{E}$ when the wage schedule $\mathbf{w}$ is consistent with the agent choosing effort level e. For a fixed effort level e, constraints (IR), (IC), and (LL) characterize the set of feasible wages that implement e. The principal will choose a wage schedule belonging to that set that achieves z_e by minimizing the expected payment $\boldsymbol{\pi}_e \mathbf{w}$. We denote an effort level as *implementable* if it is attainable under some wage schedule, that is, its set of feasible wages is nonempty.

From the perspective of the system, the social welfare when implementing effort level $e \in \mathcal{E}$ is given by $u_e^{\mathrm{SW}} \triangleq \boldsymbol{\pi}_e \mathbf{y} - c_e$, which corresponds to the sum of the principal's utility $\boldsymbol{\pi}_e \mathbf{y} - z_e$ and the agent's utility $z_e - c_e$. This follows because preferences are quasi-linear and wages are a pure transfer of wealth between the principal and the agent. The optimal social welfare is given by

$$u^{\mathrm{SO}} \triangleq \max_{e\in\mathcal{E}} \{u_e^{\mathrm{SW}}\},$$

and the set of first-best efficient effort levels is given by $\mathcal{E}^{\mathrm{SO}} \triangleq \arg\max_{e\in\mathcal{E}}\{u_e^{\mathrm{SW}}\}$. Notice that since z_e is the optimal objective value of MPLP_e, we have by the agent's IR condition that $z_e \geq c_e$. Thus, the agent's utility is always non-negative, and the social welfare is at least the principal's utility; i.e., $u_e^{\mathrm{SW}} \geq u_e^{\mathrm{P}}$ for all effort levels $e \in \mathcal{E}$.

2.2. Assumptions

We assume that the probability distributions $\{\boldsymbol{\pi}_e\}_{e\in\mathcal{E}}$ satisfy the well-known *monotone likelihood-ratio property*.[8]

Assumption 1 ***(MLRP).*** The distributions $\{\boldsymbol{\pi}_e\}_{e\in\mathcal{E}}$ satisfy $\dfrac{\pi_e^s}{\pi_f^s} > \dfrac{\pi_e^t}{\pi_f^t}$ for all states $s > t$ and effort levels $e > f$.

The **MLRP** is pervasive in the economics of information literature, and in particular in the principal–agent literature (see, e.g., Grossman and Hart, 1983; Rogerson, 1985). This property ensures that the higher the observed level of output, the more likely it is that the agent exerted a higher effort level.

Any distribution satisfying **MLRP** also satisfies the weaker *first-order stochastic dominance* (**FOSD**) property. Let $F_e^s \triangleq \sum_{s'=1}^{s} \pi_e^{s'}$ be the cumulative distribution function for effort level e. Rothschild and Stiglitz (1970) proved that for a fixed outcome the cumulative distribution function is non-increasing in the effort level under **MLRP**; or equivalently, $F_f^s \geq F_e^s$ for all states s and effort levels $f < e$. A simple consequence of this that plays an important role in our derivations is that probabilities for the highest outcome S are sorted in increasing order with respect to effort levels; i.e., $\pi_f^S < \pi_e^S$ for $f < e$. Note that in the case of two outcomes, **MLRP** and **FOSD** are equivalent.

The following ratios are central in the analysis that follows so we refer to them explicitly. We let the ratio of marginal cost to marginal probability of the highest outcome be equal to

$$m_e \triangleq \frac{c_e - c_{e-1}}{\pi_e^S - \pi_{e-1}^S} \text{ for all } e \in \mathcal{E}.$$

In the case of two outcomes and an arbitrary number of effort levels, our results hold under the Assumption **IMP** that requires that all effort levels are implementable (see Section B.2 of Supplementary Material). For the general case (arbitrary number of outcomes and effort levels) we impose the following two additional assumptions.

[8] Our results also hold under the weaker assumption that the largest likelihood ratio is given by the highest outcome, that is, $\pi_e^S/\pi_f^S > \pi_e^s/\pi_f^s$ for all states $s < S$ and effort levels $e > f$. We thank an anonymous referee for suggesting this point.

Table 1
Assumptions required for our results to hold.

Assumption name	Description	Assumption number	Case it applies to
MLRP	Monotone Likelihood Ratio Property	1 (§2.2)	general
IMCP	Increasing Marginal Cost of Probability	2 (§2.2)	> 2 outcomes
QCSW	Quasi-Concave Social Welfare	3 (§2.2)	> 2 outcomes
IMP	All efforts level are IMPlementable	4 (Supplementary §B.2)	2 outcomes
FOSD	First-Order Stochastic Dominance	(§4)	extensions
IMP^{L}	**IMP** with Linear contracts	(§4.1)	general linear

Assumption 2 *(**IMCP**).* The ratio of marginal cost to marginal probability of the highest outcome is non-decreasing with the effort level; i.e., $m_e \leq m_{e+1}$ for all effort levels $1 \leq e < E$.

This assumption establishes that the marginal cost of increasing the probability of outcome S from π^S_{e-1} to π^S_e, expressed in terms of per-unit increase in probability, is non-decreasing with the effort level. This assumption is satisfied, for instance, when the marginal cost of effort is constant and there are decreasing marginal returns to effort. It is also satisfied when the marginal cost of effort is increasing and there are constant marginal returns to effort. Hence, this assumption is the natural analogue of the standard convexity of the distribution function assumption (**CDFC**).[9] In fact, one can easily show that **CDFC** together with convexity of the cost function implies **IMCP**. Our condition is weaker than **CDFC** since we only impose a restriction on the likelihoods for the highest outcome while the latter requires convexity of the distribution for every outcome.

The assumptions **CDFC** together with **MLRP** are standard conditions for the validity of the first-order approach (see, for instance, Rogerson, 1985, Mirrlees, 1974, 1999, 1976, Laffont and Martimort, 2001 and Salanié, 2003).[10] This is meant to ensure that critical points are global maxima, making local conditions enough to characterize the global optima for the agent's effort level. Here, assumption **IMCP** plays a similar role. That is, an effort level satisfying the local incentive compatibility constraint is a global maximum in the agent's optimization problem. This allows one to order the effort levels in a manner that facilitates the computations upon deviations.[11] Furthermore, **IMCP** implies that any effort level is *implementable* since the incentive constraints are ordered.

Assumption 3 *(**QCSW**).* The sequence of prevailing social welfare levels under increasing effort levels is quasi-concave.

This assumption establishes that the social welfare is unimodal, or equivalently for some effort level e^{SO}, social welfare is monotonically non-decreasing for effort levels $e \leq e^{SO}$ and monotonically non-increasing for effort levels $e \geq e^{SO}$. This technical assumption excludes the possibility of local maxima in the social welfare and greatly simplifies our results.

To summarize, Table 1 lists the assumptions needed for the different results we show.

2.3. Optimal contract

A critical step needed to derive our results involves characterizing the minimum expected payment z_e incurred by the principal when inducing an effort level $e \in \mathcal{E}$. From the theory of linear programming, we know that the optimal wage schedule is given by an extreme point of the constraint set. In general, the IC constraints of many other effort levels $f \in \mathcal{E} \setminus e$ may be binding at the optimal wage schedule, not allowing one to compactly characterize the minimum expected payment. While not needed,[12] assumption **IMCP** greatly simplifies the analysis since it implies that the IC constraint relative to the immediately lower effort level is necessary and sufficient to implement effort level e. That is, to implement the effort level $e > 1$, the principal only needs to check that the agent's utility for the effort level at hand dominates that of the previous effort level $e - 1$. Missing proofs are available in the appendix.

Proposition 1. *Assume that **MLRP** and **IMCP** hold. For every effort level $e \in \mathcal{E}$ the following holds:*

1. *The minimum expected payment that the principal makes to the agent when he induces effort level e is $z_e = \pi^S_e m_e$.*
2. *The optimal contract pays $w^s_e = 0$ for all $s < S$ and $w^s_e = m_e$ for $s = S$.*

[9] The assumption **CDFC** requires that the distribution function is convex with respect to the effort level for any given fixed outcome; that is, $F^s_{e+1} - F^s_e \geq F^s_e - F^s_{e-1}$ for all outcomes s and effort levels $1 < e < E$.

[10] See, Jewitt (1988) for a less stringent but much less used condition ensuring that the first-order approach is valid. One advantage of Jewitt's conditions is that they are closer to the economic notion of decreasing returns.

[11] However as Rogerson (1985, p. 1362) points out, "if output is determined by a stochastic production function with diminishing returns to scale in each state of nature, the implied distribution function over output will not, in general, exhibit the **CDFC**." Thus, the **CDFC** in general requires more than diminishing returns (see Conlon, 2009 for a more detailed discussion).

[12] For details, see the extended version of the current paper (Balmaceda et al., 2012).

The optimal wage schedule in Proposition 1 rewards the agent only when the highest outcome is realized. Specifically, the principal will only reward the highest outcome (due to the **MLRP**), and he will impose the maximum penalty for all other outcomes. Because of the agent's liability limit, this maximum penalty is equal to zero.

It is worthwhile to say a few words on the structure of the optimal contract. While the optimal payment scheme is very intuitive, from an empirical point of view it is implausible to pay a bonus only when the highest outcome is observed. In fact, if we extend the model to the empirically-relevant case of a continuous outcome space, the optimal contract under our assumptions will require the principal to pay an infinite bonus when the highest outcome is observed, which occurs with probability zero.[13] The literature on financial contracts with moral hazard and limited liability avoids this unappealing feature by considering monotonicity constraints, in which case a threshold contract (i.e., debt contract) is optimal (Innes, 1990; Matthews, 2001; Poblete and Spulber, 2012). However, within the literature on moral hazard with limited liability that deals with employment contracts, monotonicity constraints are usually disregarded and many papers focus on the two outcome case, in which case the bonus-type contract derived here is the natural solution. Our paper deals with bounds on the welfare loss that relate to this latter literature, and our results are useful as a first step towards a more general study of the welfare loss from moral hazard in more plausible settings.

To close this section we show that a simple consequence of the previous result is that the agent's utility is non-decreasing with respect to the effort level.

Corollary 1. *Assume that **MLRP** and **IMCP** hold. The agent's utility u_e^{A} is nondecreasing in the effort level implemented by the principal; that is, $u_e^{\mathrm{A}} \le u_{e+1}^{\mathrm{A}}$ for all $1 \le e < E$.*

This implies that in order to induce higher effort levels the principal must increase the agent's revenue at a higher rate than the increase in cost experienced by the agent. As we will discuss in more detail later, this creates an incentive for the principal to induce a lower-than-optimal effort level.

2.4. Monopsony interpretation

The principal–agent model can be alternatively interpreted as a situation where the principal is a monopsonist and the agent is a multi-product firm supplying the monopsonist. In this context there is no uncertainty: the S possible states of nature correspond to S potential markets and π_e^s is the deterministic quantity sold in market $s \in \mathcal{S}$. The agent-supplier's outputs in the S markets are joint outputs determined by his single effort level e, and his production function is given by the vector $\boldsymbol{\pi}_e = (\pi_e^1, \dots, \pi_e^S)$. The principal-monopsonist is assumed to offer a linear price w^s and has a gross revenue y^s per-unit of output π_e^s in market $s \in \mathcal{S}$. As a result, the agent-supplier's total profit is given by $\boldsymbol{\pi}_e \mathbf{w} - c_e$ and the principal-monopsonist's total profit is $\boldsymbol{\pi}_e(\mathbf{y} - \mathbf{w})$.

A priori the problem of determining the optimal prices and effort level seems challenging because the agent-supplier's joint outputs for the S markets are determined by a single effort level. Since players are risk neutral, however, the **MLRP**[14] implies that the principal-monopsonist only rewards the output of market S (the market with the highest gross revenue) and pays nothing for the outputs of the other markets. Hence, the agent-supplier problem simplifies to that of a single-output competitive firm facing competitive price w^S for output π_e^S supplied at total cost c_e and marginal cost m_e. Because of **IMCP**, marginal costs are increasing so local conditions are sufficient for global optimality and the agent-supplier chooses the greatest effort level $e \in \mathcal{E}$ such that $m_{e+1} > w^S \ge m_e$. It follows that the principal-monopsonist's cost-minimizing contract to induce effort level $e \in \mathcal{E}$ is given by $w_e^S = m_e$ and $w_e^s = 0$ for all other markets $s < S$. As a result, the principal pays the agent-supplier the marginal cost of production at effort level e, as in Proposition 1. The principal then becomes a single-market monopsonist who offers price m_e to his supplier and whose production function is $\boldsymbol{\pi}_e \mathbf{y}$, so that its *marginal revenue product* is

$$\mathrm{MRP}_e \triangleq \frac{\boldsymbol{\pi}_e \mathbf{y} - \boldsymbol{\pi}_{e-1} \mathbf{y}}{\pi_e^S - \pi_{e-1}^S}. \tag{3}$$

Later we will show that the principal-monopsonist implements an inefficient effort level, which is consistent with the fact that the quantity demanded by the monopsonist is inefficiently low.

3. Bounding the welfare loss

The goal of this section is two-fold. First, we aim to provide simple parametric bounds on the welfare-loss ratio of a given instance of the problem in the presence of limited-liability and moral hazard under the previous assumptions. Second,

[13] In the case of a continuous outcome space, Kim (1997) showed that, under the assumption that the first-best effort is achievable, there exists an optimal contract that pays a fixed wage plus a bonus when the outcome exceeds a given threshold. If the first-best is not achievable, as it happens here, then an optimal contract would consist of an infinite payment in the zero-probability event that the highest outcome occurs.

[14] In this context **MLRP** imposes that the ratio of output between high gross-revenue markets and low gross-revenue markets is increasing in the agent-supplier's effort level.

we study the worst-case welfare-loss ratio among all problem instances satisfying our assumptions, which is commonly referred to as the *Price of Anarchy* in the computer science literature (Nisan et al., 2007).

3.1. The welfare loss

For a given instance of the problem $I = (\boldsymbol{\pi}, \mathbf{y}, \mathbf{c}) \in \mathbb{I}^{E,S}$ we quantify the welfare-loss ratio, denoted by $\rho(I)$, as the ratio of the social welfare under the socially-optimal effort level to that of the socially worst second-best effort level; that is,

$$\rho(I) \triangleq \frac{u^{\mathrm{SO}}}{\min_{e \in \mathcal{E}^{\mathrm{P}}} u_e^{\mathrm{SW}}}. \tag{4}$$

Note that the minimum in the denominator corresponds to taking a worst-case perspective which reflects that, among the possibly multiple outcomes, one cannot know which one the principal will choose. The welfare-loss ratio is clearly at least one because the social welfare of any sub-game perfect equilibrium cannot be larger than the socially-optimal one.

The *worst-case welfare-loss ratio* of a class of problem instances $\mathcal{I} \subseteq \mathbb{I}^{E,S}$ is defined as the supremum of the welfare-loss ratios over all problem instances in the class[15]:

$$\rho(\mathcal{I}) \triangleq \sup_{I \in \mathcal{I}} \rho(I). \tag{5}$$

In particular, we are interested in the worst-case welfare-loss ratio for the class of problem instances $\mathbb{I}^{E,S}_{MIQ} \subset \mathbb{I}^{E,S}$ with E effort levels and S outcome levels satisfying assumptions **MLRP**, **IMCP** and **QCSW**.

It is worth noting that, without limited liability, moral hazard does not affect the efficiency of the optimal contract. That is, if the principal and the agent are risk-neutral and there is no limited-liability constraint, the minimum expected payment z_e incurred by the principal when inducing a implementable effort level e is c_e (see, e.g., Laffont and Martimort, 2001, p. 154). Thus, without the constraint (LL), the principal implements the socially-optimal effort level and he fully captures all social surplus, leaving no rent to the agent. As a consequence, the worst-case welfare-loss ratio is 1.

In the monopsonist interpretation of our model the first-best coincides with the competitive equilibrium case where the principal-monopsonist acts as if he is price taking, and he determines the optimal effort level by equating the marginal revenue product MRP_e to the marginal cost of the output m_e. Therefore the welfare-loss ratio $\rho(I)$ of a problem instance I gives the loss of efficiency of the monopsonized industry relative to the competitive industry.

3.2. The main result

Situations with high inefficiency may arise under our assumptions when the socially-optimal effort level is high, but the optimal choice for the principal is to induce a low effort level. In view of Corollary 1, inducing a high effort level is costly for the principal because the agent's limited-liability rent increases, so the agent's payment must increase at a higher rate than the agent's cost. As a result the principal may find it optimal to induce a low effort level because the net social gain of increasing the expected output is smaller than the rent to the agent when inducing the high effort level.

We bound the welfare-loss ratio when the effort level is chosen from a discrete set that takes $E \geq 2$ possible values, and give a bound on the welfare-loss ratio that involves the sum of the *relative marginal returns to effort* given by

$$\mathrm{RMRE}_e \triangleq \frac{\pi_e^S - \pi_{e-1}^S}{\pi_e^S},$$

for effort levels $1 < e \leq E$. A salient feature of this bound, called RMRE from now on, is that it provides some structural insights into problem instances with high welfare loss. Additionally, we show that welfare-loss ratio is bounded from above by $1 + \ln\left(\pi_E^S / \pi_1^S\right)$, which depends on the probabilities of the lowest and highest effort levels, but not on the total number of effort levels. An advantage of this last bound is that, in some problem instances, it may achieve a well-behaved limit as we increase the number of effort levels. Overall, the next result is appealing because it provides upper bounds on the inefficiency that only depend on the basic parameters of the problem instance at hand, without having to compute first or second-best contracts.[16]

Theorem 1. *Let $I = (\boldsymbol{\pi}, \mathbf{y}, \mathbf{c}) \in \mathbb{I}^{E,S}_{MIQ}$ be an instance of the problem satisfying assumptions* ***MLRP****,* ***IMCP*** *and* ***QCSW****. Then, the welfare-loss ratio is bounded from above by*

[15] While the worst-case inefficiency corresponding to a family of problem instances $\rho(\mathcal{I})$ has been dubbed the 'price of anarchy,' the computer science literature refers to the welfare-loss ratio $\rho(I)$ corresponding to a problem instance $I \in \mathbb{I}^{E,S}$ shown in (4) by the 'coordination ratio.' Further, note that the price of anarchy for a maximization problem such as the one we work with in this article is often defined as the inverse of the ratio in (5). We define the ratio in this way so ratios and welfare losses point in the same direction. Other definitions such as the relative gap between the first-best and second-best solutions are possible, but we keep the standard one for consistency.

[16] In Section B.1 of Supplementary Material we provide a similar bound in terms of the agent's personal cost.

$$\rho(I) \leq 1 + \sum_{e=2}^{E} \frac{\pi_e^S - \pi_{e-1}^S}{\pi_e^S} \leq 1 + \ln \frac{\pi_E^S}{\pi_1^S}.$$

The message of Theorem 1 is that the welfare loss cannot be arbitrarily large, because when the social inefficiency at the lowest effort level is arbitrarily large, the principal's utility at the highest effort level cannot be smaller than his utility at the lowest effort level. To fix ideas let us consider the case of two effort levels. An unbounded welfare loss would imply that the social welfare of the high effort level is arbitrarily larger than that of the low effort level. Inefficiency is introduced if the principal still prefers to induce the lower effort, which implies that his utility derived for implementing the higher effort is even smaller. Because social welfare is the sum of the principal's and agent's utility, the latter implies that the agent is capturing all of the increase in social welfare at the high effort level, which our result precludes. Indeed, we will later show that $u_e^{\mathrm{P}} \geq u_e^{\mathrm{SW}} - u_{e-1}^{\mathrm{SW}}$, which implies that when the welfare loss of implementing the lower effort level (i.e., $u_e^{\mathrm{SW}} - u_{e-1}^{\mathrm{SW}}$) is large, the principal's utility for the higher effort level (i.e., u_e^{P}) is large too, thus giving the principal an incentive to implement the higher (and efficient) effort level.

In order to provide some intuition for our main result we discuss the proof of the weaker result $\rho(I) \leq E$, which follows from writing the RMRE bound as $E - \sum_{e=2}^{E} \pi_{e-1}^S / \pi_e^S$. First note that because the utility of the agent increases with his effort level (from Corollary 1), the principal has no incentive to implement any effort level higher than $e^{\mathrm{SO}} = \max \mathcal{E}^{\mathrm{SO}}$ defined as the largest of the socially-optimal effort levels. We obtain the bound by lower bounding the social welfare under the optimal effort level for the principal, which we denote by e^*.

Let e be any effort level bigger than e^*, but no bigger than the welfare-maximizing effort level e^{SO}. We first bound the principal's utility in terms of the difference in social welfare as $u_e^{\mathrm{P}} \geq u_e^{\mathrm{SW}} - u_{e-1}^{\mathrm{SW}}$. Since under the cost-minimizing contract implementing effort level e, the agent is indifferent between e and $e-1$, we obtain that the agent's utility is

$$u_e^{\mathrm{A}} = m_e \pi_e^S - c_e = m_e \pi_{e-1}^S - c_{e-1}. \tag{6}$$

Observe that the right-hand side is close to the welfare when effort $e-1$ is implemented. We know that the marginal cost is bounded from above by the marginal revenue product, that is,

$$m_e \leq \mathrm{MRP}_e. \tag{7}$$

This holds because $u_{e-1}^{\mathrm{SW}} \leq u_e^{\mathrm{SW}}$ since $e < e^{\mathrm{SO}}$ and u_e^{SW} is quasi-concave from **QCSW**. Denoting the principal-monopsonist *average revenue product* by $\mathrm{ARP}_e \triangleq \boldsymbol{\pi}_e \mathbf{y} / \pi_e^S$, remarkably **MLRP** implies that

$$\mathrm{MRP}_e \leq \mathrm{ARP}_{e-1}. \tag{8}$$

In other words, the marginal revenue product when effort e is exerted is lower than the average revenue product when effort $e-1$ is chosen. To see this, first note that ARP_e is decreasing in the effort level, that is,

$$\mathrm{ARP}_e = \frac{\boldsymbol{\pi}_e}{\pi_e^S} \mathbf{y} \leq \frac{\boldsymbol{\pi}_{e-1}}{\pi_{e-1}^S} \mathbf{y} = \mathrm{ARP}_{e-1},$$

because **MLRP** implies that $\pi_e^s / \pi_e^S \leq \pi_{e-1}^s / \pi_{e-1}^S$ for any outcome $s \leq S$ and $\mathbf{y} \geq 0$. Therefore the marginal revenue product satisfies

$$\mathrm{MRP}_e = \frac{\pi_e^S \mathrm{ARP}_e - \pi_{e-1}^S \mathrm{ARP}_{e-1}}{\pi_e^S - \pi_{e-1}^S} \leq \mathrm{ARP}_{e-1} \frac{\pi_e^S - \pi_{e-1}^S}{\pi_e^S - \pi_{e-1}^S} = \mathrm{ARP}_{e-1},$$

as claimed. Combining equations (6), (7), and (8) it follows that

$$u_e^{\mathrm{A}} \leq \mathrm{ARP}_{e-1} \pi_{e-1}^S - c_{e-1} = \boldsymbol{\pi}_{e-1} \mathbf{y} - c_{e-1} = u_{e-1}^{\mathrm{SW}}.$$

That is, the agent's welfare from effort level e is bounded by the social welfare from $e-1$. Because $u_e^{\mathrm{SW}} = u_e^{\mathrm{P}} + u_e^{\mathrm{A}}$ from quasilinearity we obtain

$$u_e^{\mathrm{P}} = u_e^{\mathrm{SW}} - u_e^{\mathrm{A}} \geq u_e^{\mathrm{SW}} - u_{e-1}^{\mathrm{SW}}, \tag{9}$$

as we claimed.

The social welfare of the second-best effort level can be bounded from below by the difference in social welfare between effort levels e and $e-1$ as follows

$$u_{e^*}^{\mathrm{SW}} \geq u_{e^*}^{\mathrm{P}} \geq u_e^{\mathrm{P}} \geq u_e^{\mathrm{SW}} - u_{e-1}^{\mathrm{SW}},$$

where we used the facts that social welfare is at least the principal's utility, effort level e^* is optimal for the principal and inequality (9). Rearranging and then adding over $e = e^* + 1, \ldots, e^{\mathrm{SO}}$ results in the following

$$u^{\mathrm{SO}} - u_{e^*}^{\mathrm{SW}} \leq (e^{\mathrm{SO}} - e^*)u_{e^*}^{\mathrm{SW}},$$

which implies that $\rho(I) \leq E$ as we claimed.

This derivation makes it clear that the result hinges on (i) the optimal contract implementing effort level e leaves the agent indifferent between e and $e-1$; (ii) $m_e \leq \mathrm{MRP}_e$ since social welfare is increasing in e for all $e \leq e^{\mathrm{SO}}$; (iii) the fact that $\mathrm{MRP}_e \leq \mathrm{ARP}_{e-1}$, that is, the marginal revenue product of effort e is lower than the average revenue product of effort $e-1$; and (iv) the fact that the principal prefers e^* to any $e > e^*$ so that the second-best social welfare $u_{e^*}^{\mathrm{SW}}$ is at least the principal utility at e. The RMRE bound on the statement of the theorem follows similarly by replacing inequality (7) by the equality $m_e = \mathrm{MRP}_e - (u_e^{\mathrm{SW}} - u_{e-1}^{\mathrm{SW}})\pi_e^S/\mathrm{RMRE}_e$.

We finish this section by discussing the welfare loss in terms of the monopsony interpretation of the model. Theorem 1 shows that the welfare-loss ratio of the monopsony relative to the competitive market is at most $1 + \ln(\pi_E^S/\pi_1^S)$, that is, the welfare-loss ratio is bounded from above by the ratio of the highest possible output in market S to the lowest possible output in the same market. Thus the bound on the welfare-loss ratio depends on the sensitivity to the input level of the output in the market with highest gross revenue.

3.3. Instances with worst-case welfare loss

In remainder of this section we discuss the tightness of the bounds on the welfare-loss ratio in terms of the number of effort levels and in terms of the likelihood ratio of the highest outcome.

The previous derivation shows that the welfare-loss ratio for any problem instance $I \in \mathbb{I}_{MIQ}^{E,S}$ satisfies $\rho(I) \leq E$. The RMRE bound in Theorem 1 provides some insight into the structure of the worst-case instance of the problem. For the bound to be close to the number of effort levels, the RMRE should be as large as possible, which suggests that the likelihood of the highest outcome should increase at a geometric rate with the effort level and the likelihood ratio π_E^S/π_1^S should grow unbounded. In Appendix A.7 we provide a family of problem instances with a fixed number of effort levels E exhibiting these properties, whose welfare-loss ratio is arbitrarily close to the upper bound E. This implies that the latter bound is tight, and the worst-case welfare-loss ratio is $\rho(\mathbb{I}_{MIQ}^{E,S}) = E$.

For a fixed number of effort levels the bound on the welfare-loss given by $1 + \ln\left(\pi_E^S/\pi_1^S\right)$ is not tight in general, except in the trivial case when the welfare-loss ratio is one. However, if we are allowed to increase the number of effort levels arbitrarily, while holding fixed the likelihood ratio $r = \pi_E^S/\pi_1^S$, then it is possible to construct problem instances whose welfare-loss ratio converges to $1 + \ln(r)$ as the number of effort levels grows to infinity. In particular, consider the class of problem instances with an arbitrary number of efforts levels and likelihood ratio bounded from above by $r \geq 1$, as given by $\mathbb{I}_{MIQ}^{r,S} = \bigcup_{E \geq 2}\left\{I \in \mathbb{I}_{MIQ}^{E,S} : \pi_E^S/\pi_1^S \leq r\right\}$. Appendix A.8 shows that the worst-case welfare-loss ratio of this class is $\rho(\mathbb{I}_{MIQ}^{r,S}) = 1 + \ln(r)$.

4. Extensions

In this section we look at different extensions of our results, which illustrate the flexibility of our approach. First, we establish that the worst-case welfare-loss ratio does not change if we restrict the contracts to be linear. The motivation to look at this class of contracts arises from their prevalence in practice. Then, we show how our bounds can be tightened if the agent has to work on many identical tasks with two effort levels each. In Supplementary Appendix B.2 we show that in the case of two outcomes our bounds hold under weaker assumptions. Finally, in the extended version of the current paper we establish that our results generalize to arbitrary (potentially negative) costs for effort, to the case when the outside option has nonzero utility, and to more general limited-liability constraints, for which we can provide a more refined parametric bound that captures the dependence between the inefficiency of a contract and the lower bound on payments (Balmaceda et al., 2012).

4.1. Linear contracts

Linear contracts are simple to analyze and implement, are observed in many real-world settings, and have an appealing property, which is to create uniform incentives in the following sense. Think of $\mathbf{y}$ as aggregate output over a given period of time (say a year), and think of the agent taking several actions during this period (say one per day). In this setting, a non-linear contract may create unintended incentives over the course of the year, depending on how the agent has done so far. For instance, suppose the contract pays a bonus if the output exceeds a given target level. Given this contract, once the agent reaches the target, he will stop working. He will also stop working if he is far from reaching the target in a date close to the end of the year. Neither of these two things will occur when the agent faces a linear incentive contract.[17]

[17] A growing body of empirical evidence is consistent with the prediction that non-linear contracts create history-dependent incentives: see Healy (1985) on bonus plans with ceilings and floors, Asch (1990) and Oyer (1998) on bonuses tied to quotas, and Brown et al. (1996) and Chevalier and Ellison (1999) on how the convex relationship between mutual fund performance and assets under management caused risk-taking portfolio choices by ostensibly conservative funds. For a nice theoretical discussion about this point see Holmström and Milgrom (1987).

However, focusing on linear contracts is not free of problems. Mirrlees (1999) showed that the best linear contract, $\mathbf{w} = a + b\mathbf{y}$, is worse than various non-linear contracts. Why, then, are linear contracts so common in practice? A principal could pay a large premium for "simplicity" if he adopts a linear contract in settings where nonlinear contracts are optimal. A partial explanation arises by comparing the welfare loss when the principal attempts to motivate a risk-neutral worker subject to limited liability with a linear contract rather than an unrestricted contract. Naturally, this does not fully answer the question since this would necessitate a tight bound on the principal's profit loss. A plausible explanation of why firms do not deviate from linear contracts may be the fact that non-linear contracts are expensive to manage, lend themselves to gaming, and the ensuing worst-case bound is the same as that when the optimal nonlinear contract is used. In addition, linear contracts, when the outcome space is continuous, avoid the implausible feature that the contract considered here will pay an infinite bonus when the highest outcome is realized, which is a measure-zero event.

A linear contract is characterized by two parameters, an intercept a and a slope b, such that the wage at state $s \in \mathcal{S}$ is given by $w^s = a + by^s$. The parameters of the contracts are not restricted in any way other than by the limited-liability constraint. Given an instance of the problem $I \in \mathbb{I}^{E,S}$ we denote by $\rho^{\mathrm{L}}(I)$ the welfare-loss ratio when the principal implements the optimal linear contract. The next result bounds the welfare-loss ratio under linear contracts under assumptions **FOSD** and **IMP**$^{\mathrm{L}}$, where **IMP**$^{\mathrm{L}}$ requires that all effort levels $e \in \mathcal{E}$ are implementable by the principal via a linear contract.

Theorem 2. *Let $I = (\boldsymbol{\pi}, \mathbf{y}, \mathbf{c}) \in \mathbb{I}^{E,S}_{FIL}$ be an instance of the problem restricted to linear contracts satisfying assumptions **FOSD** and **IMP**$^{\mathrm{L}}$. Then, denoting the vector of all ones by $\mathbf{1}$, the welfare-loss ratio satisfies*

$$\rho^{\mathrm{L}}(I) \leq 1 + \sum_{e=2}^{E} \frac{(\boldsymbol{\pi}_e - \boldsymbol{\pi}_{e-1})(\mathbf{y} - y^1\mathbf{1})}{\boldsymbol{\pi}_e(\mathbf{y} - y^1\mathbf{1})}.$$

The bound on the welfare-loss ratio when restricted to linear contracts is similar to the RMRE bound of Theorem 1. However, since contracts are linear in the output, the relative marginal returns to efforts are a function of the expected normalized output $\boldsymbol{\pi}_e(\mathbf{y} - y^1\mathbf{1})$ instead of the probability of the highest outcome π_e^S. In the two-outcome case ($S = 2$) these bounds are equivalent because $\mathbf{y} - y^1\mathbf{1} = (0, y^2 - y^1)$, as expected since linearity imposes no restriction in the space of contracts. The proof of Theorem 2 proceeds by exploiting the fact that an instance of the problem $I \in \mathbb{I}^{E,S}$ restricted to linear contracts with an arbitrary number of outcomes and effort levels can be reduced to an unrestricted problem $\tilde{I} \in \mathbb{I}^{E,2}$ with the same number of effort levels but only *two* outcomes (one for each parameter). As a result we show that the welfare-loss ratio when restricted to linear contracts is equal to the unrestricted welfare-loss ratio of the reduced instance of the problem, or equivalently $\rho^{\mathrm{L}}(I) = \rho(\tilde{I})$.[18]

Next, we characterize the worst-case welfare-loss ratio when restricted to linear contracts of the class of instances $\mathbb{I}^{E,S}_{FIL} \subset \mathbb{I}^{E,S}$ with E effort levels and S outcome levels satisfying assumptions **FOSD** and **IMP**$^{\mathrm{L}}$. The previous theorem yields that $\rho^{L}(\mathbb{I}^{E,S}_{FIL}) \leq E$. When the original problem has two outcomes there is a one-to-one correspondence between any wage schedule and the two parameters of the linear contract.[19] As a result we obtain that the worst-case welfare-loss ratio restricted to linear contracts is $\rho^{L}(\mathbb{I}^{E,S}_{FIL}) = E$, as before.

Recall that in Section 3, we proved bounds for the welfare-loss ratio under non-linear contracts. Combining those bounds with Theorem 2, we have that the worst-case welfare-loss ratio under linear contracts coincides with that for the general case. Therefore, our results suggest that with respect to the *worst-case scenario*, restricting attention to linear contracts does not generate more inefficiency from a social point of view. Our results, however, do not shed light on whether, for a fixed instance of the problem I, the unrestricted welfare-loss ratio $\rho(I)$ and the welfare-loss ratio when restricted to linear contracts $\rho^{\mathrm{L}}(I)$ coincide or not.

To put our bounds in perspective we conclude by comparing the unrestricted RMRE bound of Theorem 1 with that parametric bound restricted to linear contracts of Theorem 2. Let the relative marginal return to effort for linear contracts be defined as

$$\mathrm{RMRE}^{\mathrm{L}}_e \triangleq \frac{(\boldsymbol{\pi}_e - \boldsymbol{\pi}_{e-1})(\mathbf{y} - y^1\mathbf{1})}{\boldsymbol{\pi}_e(\mathbf{y} - y^1\mathbf{1})},$$

for effort levels $1 < e \leq E$. If we strengthen **FOSD** to **MLRP** we can show that $\mathrm{RMRE}^{\mathrm{L}}_e \leq \mathrm{RMRE}_e$ for effort levels $e > 1$, and thus the same RMRE bound of Theorem 1 applies in the case of linear contracts.[20]

[18] In view of the result in the Section B.2 of Supplementary Material for instances of the problem with two outcomes, we only need **IMP** and **FOSD** for our results to hold, which, as we show, translate to the reduced instance of the problem. Assumption **IMP** is formally defined in Supplementary §B.2 and requires that all effort levels are implementable by the principal when he is not restricted to linear contracts. Alternatively, we could require that **IMCP** holds for the reduced instance of the problem instead of **IMP**, however, we believe that our assumption on the implementability of the effort levels is more natural for this setting.

[19] The instance of the problem in the Supplementary Material Section B.2 applies and the bound for welfare-loss ratio is tight.

[20] Letting $\mathbf{y}' = \mathbf{y} - y^1\mathbf{1}$, the condition $\mathrm{RMRE}^{\mathrm{L}}_e \leq \mathrm{RMRE}_e$ can be written after some algebra as $(\pi_e^S \boldsymbol{\pi}_{e-1} - \pi_{e-1}^S \boldsymbol{\pi}_e)\mathbf{y}' \geq 0$, which follows because $\mathbf{y}' \geq 0$ from outputs being increasing, and $\pi_e^S \boldsymbol{\pi}_{e-1} \geq \pi_{e-1}^S \boldsymbol{\pi}_e$ from **MLRP**.

4.2. Multiple tasks

Often principal–agent relationships require that the agent performs different tasks, each endowed with different actions. In this section, we consider a principal–agent relationship with multiple tasks, adopting a model proposed by Laux (2001). The principal is endowed with N identical and stochastically independent tasks. Each task has two possible outcomes: success or failure. The corresponding payoffs for the principal are $\overline{y}$ if the task is successful, and $\underline{y}$ in the case of failure, with $\overline{y} > \underline{y}$. For each task, the agent can exert two effort levels: either high or low. The high effort level entails a cost c_h for the agent, while the cost of the low effort level is c_l. Since a higher effort level demands more work, $c_h > c_l$. Finally, we denote by p_h the probability of success when the effort level is high, and by p_l the probability of success when the effort level is low. We assume that **FOSD** holds for each task, which means that $p_h > p_l$ (the higher the effort level, the greater the likelihood of success).

The principal hires an agent to perform the N tasks. Since tasks are identical, the principal offers a compensation that depends only on the number of tasks that end up being successful, denoted by $s \in \mathcal{S} = \{0, \dots, N\}$; the identity of each task is irrelevant. Hence, the agent is paid a wage w^s when s tasks turn out to be successful. The total revenue for the principal is thus $y^s = s\overline{y} + (N-s)\underline{y}$ for $s \in \mathcal{S}$. In view of the tasks' symmetric nature, the agent is indifferent between tasks and he is only concerned about the total number of tasks in which he exerts high effort level. We define the *aggregated effort level* $e \in \mathcal{E} = \{0, \dots, N\}$ as the number of tasks in which the agent works hard. Notice that, for notational simplicity, we adopt indices that start at zero for both effort levels and states. We assume that costs are additive, and linear in the number of tasks. Hence, the aggregate costs for the agent are $c_e = ec_h + (N-e)c_l$ for $e \in \mathcal{E}$. Finally, note that the probability of having s successful tasks, given that the agent works hard on e tasks, is given by

$$\pi_e^s = \sum_{i=0}^{s} \binom{e}{i} p_h^i (1-p_h)^{e-i} \binom{N-e}{s-i} p_l^{s-i} (1-p_l)^{N-e-s+i},$$

where we let $\binom{n}{k} = 0$ if $k > n$ for notational simplicity.

This model can be fully reduced to a principal–agent model with a single task, $N+1$ states and $N+1$ effort levels. To map the multiple-task model into the model of Section 2 we show that the aggregate instance of the problem satisfies **MLRP**, which implies that the larger the observed number of successful tasks, the more likely it is that the agent works hard in many tasks.

Lemma 1. *Assume that **FOSD** holds for each of the identical tasks, that is, $p_h > p_l$. When the principal hires one agent to perform N of these independent tasks, the probability distribution of the outcome satisfies **MLRP** for the aggregated problem.*

Assumption **IMCP** does not hold in this setting since the ratios m_e are non-increasing with the aggregate effort level. Laux (2001) shows that when the agent exerts the high effort level in all tasks (aggregate effort level N), the only binding constraint is the one that corresponds to choosing a low effort level for the N tasks (aggregate effort level 0). As a result, in equilibrium the principal induces either the lowest or highest aggregate effort level. Using this characterization of the optimal payments we can show that 2 is a tight bound for the welfare-loss ratio, regardless of the number of tasks. Note that assumption **QCSW** is always guaranteed to hold since social welfare is linear in the number of tasks for which the agent exerts a high effort level.

Theorem 3. *Assume that **FOSD** holds for each of the identical tasks, that is, $p_h > p_l$. In the principal–agent problem in which both players are risk-neutral and there are N identical and independent tasks, the welfare-loss ratio is bounded from above by $2 - (p_l/p_h)^N < 2$.*

5. Conclusions

This paper quantifies the welfare loss that arises from the principal's inability to observe the agent's effort level when there is limited liability. We have provided a simple parametric bound on the welfare-loss ratio involving the probabilities of the highest possible outcome. This bound leverages the structure of the optimal contract that pays a bonus only when the highest outcome is observed. The general structure of the bounds found in this paper suggests that the welfare loss in a principal–agent relationship depends on the set of effort levels available to an agent as well as the likelihood ratio of the probability that the highest outcome occurs when the highest versus the lowest effort level is considered. Furthermore, the paper provides an interpretation of the principal–agent model with limited liability as a monopsony model, which reduces the problem of determining the optimal contract to a standard intermediate microeconomics analysis. In this setting our results bound the efficiency loss of the competitive industry relative to the monopsonized industry, independently of the equilibrium input traded in the competitive and monopsonist markets.

The principal–agent model in its different forms has been used to explain many contractual arrangements such as sharing contracts, insurance contracts, managerial contracts, political relationships, etc. In addition, it has been used to provide an economic theory of the firm and a theory of organizational forms. Our results show that in these cases and in many others, the existence of an agency relationship with moral hazard may have nontrivial consequences in terms of welfare loss and thus the proper design of contracts and organizations to deal with moral hazard is of great practical importance.

Nonetheless, in spite of the relevance of our results, they open more questions than they answer. The most ambitious question is to quantify the welfare loss for each plausible instance of the problem with more general cost functions, monitoring technologies and utility functions. A feasible step in this direction is to consider a risk-averse agent. In this case the optimal contract is highly nonlinear, and thus its characterization in terms of the main parameters is a complex task. There is an exception to this, which is given by the linear agency model by Holmström and Milgrom (1987), which could be used as the starting point to investigate the question at hand. Another potential extension involves considering monotonicity constraints as done in the literature on optimal financial contracts with moral hazard. Monotonicity of the contract and of the principal's payoff with respect to the outcome leads to more reasonable threshold contracts, at the expense of a more complex analysis.

Appendix A. Proofs

A.1. Proof of Proposition 1

First, we prove the following lemma, which is rather straightforward but a necessary building block to prove Proposition 1.

Lemma 2. *Assume that* **IMCP** *holds, and let $m_{e,f} = (c_e - c_f)/(\pi_e^S - \pi_f^S)$ with the convention that $c_0 = \pi_0^S = 0$ (so that $m_{e,0} = c_e/\pi_e^S$). Then, for all $e \in \mathcal{E}$:*

(i) $m_e \geq m_{e,f}$, *for all $f = 0, \ldots, e-1$;*
(ii) $m_e \leq m_{f,e}$, *for all $f = e+1, \ldots, E$.*

Proof. For item (i) note from **IMCP** we have that $m_{g,g-1} \leq m_{e,e-1}$ for all $g \leq e$, which implies that $(c_g - c_{g-1})(\pi_e^S - \pi_{e-1}^S) \leq (c_e - c_{e-1})(\pi_g^S - \pi_{g-1}^S)$ because costs and the probabilities of the highest outcome are non-decreasing with respect to the effort level. Adding over $g = f+1, \ldots, e$ and collecting terms we obtain that $(c_e - c_f)(\pi_e^S - \pi_{e-1}^S) \leq (c_e - c_{e-1})(\pi_e^S - \pi_f^S)$, and the result follows.

For item (ii) we proceed similarly and get from **IMCP** that $(c_g - c_{g-1})(\pi_e^S - \pi_{e-1}^S) \geq (c_e - c_{e-1})(\pi_g^S - \pi_{g-1}^S)$ for $g > e$. Adding over $g = e+1, \ldots, f$ and collecting terms we obtain that $(c_f - c_e)(\pi_e^S - \pi_{e-1}^S) \geq (c_e - c_{e-1})(\pi_f^S - \pi_e^S)$, and the result follows. $\square$

With the lemma in hand, we are now ready to prove the proposition.

Proof. The proof proceeds by relaxing the principal's problem for effort level e to include *only* the IC constraint for effort levels lower than e, and then showing that the optimal solution of the relaxed problem is feasible for the original problem. The relaxed problem is[21]:

$$\begin{aligned}
&\min_{\mathbf{w}\in\mathbb{R}^S} \boldsymbol{\pi}_e \mathbf{w} \\
&\quad \text{s.t. } \boldsymbol{\pi}_e \mathbf{w} - c_e \geq 0, \\
&\qquad \boldsymbol{\pi}_e \mathbf{w} - c_e \geq \boldsymbol{\pi}_f \mathbf{w} - c_f, \ \ \forall f < e, \\
&\qquad \mathbf{w} \geq 0.
\end{aligned}$$

We claim that **MLRP** implies that there exists an optimal solution for the relaxed problem in which all contractual payments can be focused on the highest outcome, with payments of zero for all other outcomes. Suppose that there exists an optimal contract $\mathbf{w}_0$ with positive payments for some outcome $t < S$, and consider the new contract $\mathbf{w}_1$ that pays $w_1^S = w_0^S + \frac{\delta}{\pi_e^S}$, $w_1^t = w_0^t - \frac{\delta}{\pi_e^t}$, and $w_1^s = w_0^s$ for all $s \neq \{t, S\}$. The principal's cost of implementing this new contract is $\boldsymbol{\pi}_e \mathbf{w}_1 = \boldsymbol{\pi}_e \mathbf{w}_0$, and thus he is indifferent between these two. The IR constraint for $\mathbf{w}_1$ follows directly from the IR constraint of $\mathbf{w}_0$. The difference between contract $\mathbf{w}_1$ and $\mathbf{w}_0$, in terms of the incentives provided to the agent to choose effort level e over effort level $f < e$ is

$$(\boldsymbol{\pi}_e \mathbf{w}_1 - \boldsymbol{\pi}_f \mathbf{w}_1) - (\boldsymbol{\pi}_e \mathbf{w}_0 - \boldsymbol{\pi}_f \mathbf{w}_0) = \boldsymbol{\pi}_f (\mathbf{w}_0 - \mathbf{w}_1) = \left(\frac{\pi_f^t}{\pi_e^t} - \frac{\pi_f^S}{\pi_e^S} \right) \delta,$$

[21] Similar relaxations have been considered in the literature. In the continuous case, for example, Rogerson (1985) relaxes the IC constraints to require that the agent chooses an action at which his utility is at a stationary point.

since $\boldsymbol{\pi}_e \mathbf{w}_1 = \boldsymbol{\pi}_e \mathbf{w}_0$. The expression in parenthesis is positive because of **MLRP**. Hence, the new contract also satisfies the relaxed incentive constraint for $\delta \geq 0$. By picking $\delta = w_0^t \pi_e^t$ we can make the payment in state t equal to zero. Repeating the process for all states $t < S$ with positive payments we prove the claim.

Using the fact that $w^s = 0$ for $s < S$, the IC constraints become $w^S \geq m_{e,f}$ for $f < e$. From item (i) in Lemma 2 we can simplify the IC constraints to $w^S \geq m_e$ since all other constraints are redundant. This also gives the IR constraint. The optimal payment for the relaxed problem is the smallest possible value of w^S, which is given by $w^S = m_e$. Feasibility of the relaxed solution for the original problem follows from checking that the IC constraints for effort levels $f > e$ hold, which follows from item (ii) in Lemma 2 and **FOSD**. $\square$

A.2. Proof of Corollary 1

Proof. In order to implement $e+1$, the principal must offer the contract $w_{e+1}^S = m_{e+1}$ to the agent. Incentive compatibility implies that the agent prefers effort level $e+1$ to e under w_{e+1}^S. As a result we obtain:

$$u_{e+1}^{\mathrm{A}} = \pi_{e+1}^S m_{e+1} - c_{e+1} \geq \pi_e^S m_{e+1} - c_e \geq \pi_e^S m_e - c_e = u_e^{\mathrm{A}},$$

where the last inequality follows from **IMCP** and the last equality because $w_e^S = m_e$ is the optimal contract for effort level e. $\square$

A.3. Proof of Theorem 1

Proof. In order to prove the result we first show that the principal has no incentive to implement an effort level higher than $e^{\mathrm{SO}} = \max \mathcal{E}^{\mathrm{SO}}$, defined as the largest of the socially-optimal effort levels. Then we obtain the bound by lower bounding the social welfare under the optimal effort level for the principal, which we denote by e^*. We conclude by proving the logarithm bound for RMRE.

Step 1. Because the utility of the agent increases with his effort level, the principal has no incentive to implement any effort level higher than e^{SO}. Indeed, for $f > e^{\mathrm{SO}}$, Corollary 1 together with the fact that e^{SO} is the largest socially-optimal effort level imply that

$$u_{e^{\mathrm{SO}}}^{\mathrm{P}} - u_f^{\mathrm{P}} = u^{\mathrm{SO}} - u_f^{\mathrm{SW}} + u_f^{\mathrm{A}} - u_{e^{\mathrm{SO}}}^{\mathrm{A}} > 0.$$

Hence, effort levels larger than e^{SO} provide a suboptimal utility to the principal, and can be disregarded.

Step 2. Note that if $e^* = e^{\mathrm{SO}}$ then the welfare-loss ratio is one and the bound trivially holds. In the remainder of the proof we shall consider the case where $e^* < e^{\mathrm{SO}}$. Consider an effort level e such that $e^* < e \leq e^{\mathrm{SO}}$. The fact that the social welfare is at least the principal's utility, and Proposition 1 (because $e > 1$) imply that

$$u_{e^*}^{\mathrm{SW}} \geq u_{e^*}^{\mathrm{P}} \geq u_e^{\mathrm{P}} = \boldsymbol{\pi}_e \mathbf{y} - z_e = \boldsymbol{\pi}_e \mathbf{y} - \pi_e^S \frac{c_e - c_{e-1}}{\pi_e^S - \pi_{e-1}^S}. \tag{A.1}$$

We now use the fact that $c_e - c_{e-1} = \boldsymbol{\pi}_e \mathbf{y} - \boldsymbol{\pi}_{e-1} \mathbf{y} - u_e^{\mathrm{SW}} + u_{e-1}^{\mathrm{SW}}$ and obtain that the last expression is equivalent to

$$\frac{\boldsymbol{\pi}_{e-1} \pi_e^S - \boldsymbol{\pi}_e \pi_{e-1}^S}{\pi_e^S - \pi_{e-1}^S} \mathbf{y} + \frac{u_e^{\mathrm{SW}} - u_{e-1}^{\mathrm{SW}}}{\pi_e^S - \pi_{e-1}^S} \pi_e^S \geq \frac{\pi_e^S}{\pi_e^S - \pi_{e-1}^S} \left(u_e^{\mathrm{SW}} - u_{e-1}^{\mathrm{SW}} \right), \tag{A.2}$$

where the inequality follows from **MLRP** because $\boldsymbol{\pi}_{e-1} \pi_e^S \geq \boldsymbol{\pi}_e \pi_{e-1}^S$. Combining the inequalities in (A.1) and (A.2) implies that $(\pi_e^S - \pi_{e-1}^S)/\pi_e^S \geq (u_e^{\mathrm{SW}} - u_{e-1}^{\mathrm{SW}})/u_{e^*}^{\mathrm{SW}}$. Summing over the effort levels $e = e^* + 1, \ldots, e^{\mathrm{SO}}$ and rearranging terms we get that

$$\frac{u_{e^{\mathrm{SO}}}^{\mathrm{SW}}}{u_{e^*}^{\mathrm{SW}}} \leq 1 + \sum_{e=e^*+1}^{e^{\mathrm{SO}}} \frac{\pi_e^S - \pi_{e-1}^S}{\pi_e^S} \leq 1 + \sum_{e=2}^{E} \frac{\pi_e^S - \pi_{e-1}^S}{\pi_e^S},$$

where the second inequality follows from including all effort levels greater than one in the sum, and using the fact that the probabilities π_e^S are increasing with the effort level.

Step 3. Finally, using the fact that $\ln(x) \leq x - 1$ for all $x \geq 0$, we can bound each term of the RMRE bound by $-\ln\left(\pi_{e-1}^S / \pi_e^S\right)$, which yields

$$1 + \sum_{e=2}^{E} \frac{\pi_e^S - \pi_{e-1}^S}{\pi_e^S} \leq 1 - \sum_{e=2}^{E} \ln\left(\frac{\pi_{e-1}^S}{\pi_e^S}\right) = 1 + \ln\left(\frac{\pi_E^S}{\pi_1^S}\right). \qquad \square$$

A.4. Proof of Theorem 2

Proof. We prove the result in four steps. First, we take a principal–agent problem $I = (\boldsymbol{\pi}, \mathbf{y}, \mathbf{c}) \in \mathbb{I}^{E,S}$ restricted to linear contracts with E effort levels and S outcomes, and construct an unrestricted *reduced* problem $\tilde{I} = (\tilde{\boldsymbol{\pi}}, \tilde{\mathbf{y}}, \mathbf{c}) \in \mathbb{I}^{E,2}$ with E effort levels and 2 outcomes. Second, we show that for each effort level $e \in \mathcal{E}$ the linear program MPLP_e^{L} for instance of the problem I restricted to linear contracts is equivalent to the linear program MPLP_e for the reduced instance of the problem $\tilde{I}$. Third, we show that the reduced instance of the problem $\tilde{I}$ satisfies assumptions **FOSD** and **IMP**, and $\tilde{\mathbf{y}}$ is non-negative. Assumption **IMP** is formally defined in Supplementary §B.2 and requires that all effort levels are implementable by the principal when he is not restricted to linear contracts. Finally, we conclude by showing that the bound holds because assumptions **FOSD** and **IMP** are sufficient when there are 2 outcomes (see Supplementary §B.2).

Step 1. In showing the reduction, we write the contracts in terms of wage at the lowest outcome $d = a + by^1$, and consider the wages $w^s = d + b(y^s - y^1)$. This is valid because there are no constraints on the slope and intercept of the linear contract. Then the reduced instance of the problem is given by

$$\tilde{\mathbf{y}} \triangleq (y^1, 1),$$
$$\tilde{\boldsymbol{\pi}}_e \triangleq (1, \boldsymbol{\pi}_e \mathbf{y} - y^1), \quad \forall e \in \mathcal{E},$$

and additionally we refer to the reduced wage schedule as $\tilde{\mathbf{w}} \triangleq (d, b)$. The first entry in the reduced wage vector $\tilde{\mathbf{w}}$ represents the wage at the lowest outcome while the second represents the slope of the linear contract. The expected payment under effort level $e \in \mathcal{E}$ is $\boldsymbol{\pi}_e \mathbf{w} = d + b(\boldsymbol{\pi}_e \mathbf{y} - y^1) = \tilde{\boldsymbol{\pi}}_e \tilde{\mathbf{w}}$, and the expected output satisfies $\boldsymbol{\pi}_e \mathbf{y} = \tilde{\boldsymbol{\pi}}_e \tilde{\mathbf{y}}$, as expected.

Notice that the reduced vectors $\tilde{\boldsymbol{\pi}}_e$ for $e \in \mathcal{E}$ no longer sum up to one, and thus are not probability distributions. Surprisingly, this will not be important for our results. Notice that the reduced output may not be increasing, but again our results will hold in this case too. By interpreting the entries of the reduced vectors above as two outcomes, we show that the original problem and the reduced problem are equivalent.

Step 2. Fix an effort level $e \in \mathcal{E}$. For the second point, under linear contracts, the original minimum payment linear program corresponding to effort level e, MPLP_e^{L} can be written as

$$z_e = \min_{a,b} a + b\boldsymbol{\pi}_e \mathbf{y}$$
$$\text{s.t. } a + b\boldsymbol{\pi}_e \mathbf{y} - c_e \geq 0 \tag{A.3a}$$
$$b\boldsymbol{\pi}_e \mathbf{y} - c_e \geq b\boldsymbol{\pi}_f \mathbf{y} - c_f \quad \forall f \in \mathcal{E} \setminus e \tag{A.3b}$$
$$a + by^s \geq 0 \quad \forall s \in \mathcal{S}. \tag{A.3c}$$

We first show that given a feasible solution (a, b) of MPLP_e^{L}, then the reduced solution $\tilde{\mathbf{w}} = (d, b)$ with $d = a + by^1$ is feasible for MPLP_e of the reduced instance of the problem $\tilde{I}$ and achieves the same objective. We have seen in step 1 that the objective is preserved by the reduction. The original individual rationality constraint (A.3a) can now be written as $\tilde{\boldsymbol{\pi}}_e \tilde{\mathbf{w}} \geq c_e$, while the original incentive compatibility constraints (A.3b) can be written as $\tilde{\boldsymbol{\pi}}_e \tilde{\mathbf{w}} - c_e \geq \tilde{\boldsymbol{\pi}}_f \tilde{\mathbf{w}} - c_f$ for any effort level $f \neq e$. For outcome $s = 1$ the original limited-liability constraint (A.3c) implies that $d \geq 0$. For an effort level $e > 1$ we have by the incentive compatibility constraint (A.3b) with effort level $f < e$ that $b \geq (c_e - c_f)/(\boldsymbol{\pi}_e \mathbf{y} - \boldsymbol{\pi}_f \mathbf{y}) \geq 0$, where the last inequality follows from **FOSD**. For the lowest effort level $e = 1$, one always has that $d = c_1$ and $b = 0$ is an optimal solution. Hence, the constraint $b \geq 0$ does not eliminate any optimal solution. This completes this direction of the reduction.

For the opposite direction, we need to show that given a feasible solution $\tilde{\mathbf{w}} = (d, b)$ for MPLP_e of the reduced instance of the problem $\tilde{I}$, the solution (a, b) with $a = d - by^1$ is feasible for MPLP_e^{L} and achieves the same objective. Notice that (A.3a) and (A.3b) follow directly from the reduced problem's IR and IC constraints respectively. From the reduced problem's LL we know that $b \geq 0$ and $d \geq 0$. Hence for all $s \in \mathcal{S}$ we have that $a + by^s \geq a + by^1 = d \geq 0$, because the outputs are non-decreasing. Thus, (A.3c) holds and the reduction is complete.

Step 3. First, we prove that $\tilde{\boldsymbol{\pi}}$ satisfies **FOSD**. We require that $\tilde{\pi}_e^2 \leq \tilde{\pi}_f^2$ for $e \leq f$. This is equivalent to $\boldsymbol{\pi}_e \mathbf{y} \leq \boldsymbol{\pi}_f \mathbf{y}$, which holds because the original distribution $\boldsymbol{\pi}_e$ satisfies **FOSD** and the output is non-decreasing. A similar argument shows that $\tilde{\boldsymbol{\pi}}_e \geq 0$. Finally, the equivalence given in step 2 shows that an effort level $e \in \mathcal{E}$ is implementable in MPLP_e^{L} for instance of the problem I if and only if it is implementable in MPLP_e for the reduced instance of the problem $\tilde{I}$. This shows that the reduced instance of the problem $\tilde{I}$ satisfies **IMP**.

Step 4. Finally, notice that in the proof of Theorem 4 in Section B.2 of the Supplementary Material we do not use the fact that $\tilde{\mathbf{y}}$ is increasing, and the result still holds if $\tilde{\boldsymbol{\pi}}_e$ does not sum up to one. Thus all our results apply to the reduced problem. Putting everything together we obtain that

$$\rho^{\mathrm{L}}(I) = \rho(\tilde{I}) \leq 1 + \sum_{e=2}^{E} \frac{\boldsymbol{\pi}_e \mathbf{y} - \boldsymbol{\pi}_{e-1}\mathbf{y}}{\boldsymbol{\pi}_e \mathbf{y} - y^1} = 1 + \sum_{e=2}^{E} \frac{(\boldsymbol{\pi}_e - \boldsymbol{\pi}_{e-1})(\mathbf{y} - y^1\mathbf{1})}{\boldsymbol{\pi}_e(\mathbf{y} - y^1\mathbf{1})},$$

where the first equality follows from the reduction in step 2, and the inequality follows from Theorem 4 in Section B.2 of the Supplementary Material. □

A.5. Proof of Lemma 1

Proof. Let $\{X_e\}_{e\in\mathcal{E}}$ be a family of random variables, such that X_e is the random number of successes given that the agent works hard in e tasks. Then, X_e is the sum of e independent Bernoulli random variables with success probability p_h, and $N-e$ independent Bernoulli random variables with success probability p_l. Denote by $Y(p)$ a Bernoulli random variable with success probability p; i.e., $\mathbb{P}(Y(p)=1) = p = 1 - \mathbb{P}(Y(p)=0)$. Hence, we may write X_e as

$$X_e = \sum_{f=1}^{e} Y_f(p_h) + \sum_{f=e+1}^{N} Y_f(p_l) = \sum_{f=1}^{N} Y_f(p_f(e)),$$

where the functions $\{p_f(e)\}_{f\in\mathcal{E}}$ equal p_h if $f \leq e$, p_l otherwise. Notice that for all $f \in \mathcal{E}$ the functions $p_f(e)$ are non-decreasing in e. Ghurye and Wallace (1959) or more recently Huynh (1994) show that given any number of independent Bernoulli random variables Y_f with success probability $p_f(e)$ strictly between 0 and 1, and non-decreasing in e, then the sum $\sum Y_f$ satisfies the monotone likelihood-ratio property with respect to e. □

A.6. Proof of Theorem 3

Proof. First, Laux (2001) shows that in the case of multiple tasks the principal either implements the highest or the lowest aggregate effort levels. The minimum payments are given by $z_0 = Nc_l$ for the lowest effort level, and

$$z_N = \max\left(Nc_h, Np_h^N \frac{c_h - c_l}{p_h^N - p_l^N}\right)$$

for the highest effort level. Note that Corollary 1 applies and the agent's rent is non-decreasing with the effort level. Indeed, using $c_N = Nc_h$ and $c_0 = Nc_l$ we obtain that $z_N - c_N \geq 0 = z_0 - c_0$ as expected.

Second, we show that the social welfare is linear in the number of aggregated effort levels and thus satisfies the **QCSW** assumption. Referring to the number of successes given that the agent works hard in e tasks by the random variable X_e, its expected number is $\mathbb{E}[X_e] = ep_h + (N-e)p_l$. One may write the social welfare as

$$\begin{aligned} u_e^{\mathrm{SW}} &= \pi_e y - c_e = \mathbb{E}[X_e \overline{y} + (N - X_e)\underline{y}] - c_h e - c_l(N-e) \\ &= N\left[p_l\overline{y} + (1-p_l)\underline{y} - c_l\right] + e\left[(p_h - p_l)(\overline{y} - \underline{y}) - (c_h - c_l)\right], \end{aligned}$$

which is linear in effort level e. The first term can be interpreted as the baseline social welfare when the agent exerts a low effort level in all tasks, while the second term can be interpreted as the marginal contribution to social welfare of each additional task for which the agent exerts high effort level.

Third, because social welfare is linear, the largest socially-optimal effort level is either $e^{\mathrm{SO}} = 0$ or $e^{\mathrm{SO}} = N$. The principal has no incentive to implement any effort level higher than e^{SO}, because the utility of the agent is non-decreasing with his effort level. In the case when $e^{\mathrm{SO}} = 0$ we conclude that the optimal effort level for the principal is $e^* = 0$, the welfare-loss ratio is one, and the bound trivially holds. In the remainder of the proof we shall consider the case when $e^{\mathrm{SO}} = N$. If aggregate effort level N leaves no rent to the agent, or equivalently $z_N = c_N$, then it is optimal for the principal to induce the high effort level $e^* = N$. Inefficiency is thus introduced when the principal induces the low effort level $e^* = 0$, and $z_N = Np_h^N(c_h - c_l)/(p_h^N - p_l^N)$. In this case we apply the last steps of the proof of Theorem 1 to obtain inequalities (A.1) and (A.2), and conclude that the welfare-loss ratio is bounded above by

$$1 + \frac{p_h^N - p_l^N}{p_h^N} = 2 - \left(\frac{p_l}{p_h}\right)^N < 2.$$

This bound is tight. □

A.7. Tightness of the bound in terms of the number of effort levels

In this section we present problem instances satisfying the given assumptions and whose welfare-loss ratios are arbitrarily close to the upper bound in terms of the number of effort levels. More formally, we will define a family of problem instances $I^\delta \in \mathbb{I}^{E,2}_{MIQ}$ with $E > 1$ effort levels and $S = 2$ outcomes parameterized by $\delta \in (0,1)$ satisfying assumptions **MLRP**, **IMCP** and **QCSW** that verify $\lim_{\delta\to 0} \rho(I^\delta) = E$.

The RMRE bound on Theorem 1 provides some insight on the structure of the worst-case instance of the problem. For the bound to be close to the number of effort levels it should be the case that RMRE should be as large as possible. Using the fact that the arithmetic mean dominates the geometric mean (AM–GM inequality) and canceling terms we obtain that the RMRE bound is bounded from above by

$$1 + \sum_{e=2}^{E} \frac{\pi_e^S - \pi_{e-1}^S}{\pi_e^S} = E - (E-1)\Big(\frac{1}{E-1}\sum_{e=2}^{E}\frac{\pi_{e-1}^S}{\pi_e^S}\Big) \le E - (E-1)\left(\frac{\pi_E^S}{\pi_1^S}\right)^{-\frac{1}{E-1}}.$$

Because the equality between geometric and arithmetic means holds when all terms are constant, we obtain that the worst case likelihood of the highest outcome should increase at a geometric rate. Additionally, for the bound in terms of the number of effort levels to be tight it needs to be the case that the likelihood ratio π_E^S/π_1^S goes to infinity. In the remainder on this section we provide an instance of the problem exhibiting these properties.[22]

Fixing $0 < \delta < 1$, we let the probabilities of the outcomes associated to each effort level be $\boldsymbol{\pi}_e = \big(1 - \delta^{E-e}, \delta^{E-e}\big)$ for $e \in \mathcal{E}$. Notice that the likelihood of the highest outcome increases at a geometric rate of δ^{-1}. The probability distributions are such that effort level E guarantees a successful outcome with probability one, while the lower effort levels generate a failed outcome with high probability. Clearly, these distributions satisfy that $\pi_1^2 \le \ldots \le \pi_E^2$, and thus they satisfy **MLRP**. (Recall that in the case of two outcomes **MLRP** and **FOSD** are equivalent.)

Let the output be $\mathbf{y} = (0,1)$ so that the value of success is 1. To maximize the welfare loss, the worst-case instance should be such that the social optimal effort level is E, while the principal's optimal effort level is 1. One then can treat the principal as monopsonist as explained in Section 2.4. This suggests looking at the analogue of the *marginal factor cost*; i.e.,

$$\mathrm{MFC}_e \triangleq \frac{\pi_e^S m_e - \pi_{e-1}^S m_{e-1}}{\pi_e^S - \pi_{e-1}^S}. \tag{A.4}$$

In this notation the difference in utilities for the principal between two consecutive efforts levels is

$$u_e^{\mathrm{P}} - u_{e-1}^{\mathrm{P}} = (\pi_e^S - \pi_{e-1}^S)(1 - \mathrm{MFC}_e).$$

Let the marginal factor costs for $e = 2, \ldots, E$ equal to $\mathrm{MFC}_e = 1 + \epsilon$ with $\epsilon > 0$ to be determined later. This leads the principal to induce $e^* = 1$ because $u_e^{\mathrm{P}} < u_{e-1}^{\mathrm{P}}$ from **FOSD**.

Using $\pi_e^S = \delta^{E-e}$ and $\mathrm{MFC}_e = 1 + \epsilon$ in (A.4) yields the following recursive equation for the ratios of marginal cost to marginal probability of the highest outcome:

$$m_e = (1-\delta)(1+\epsilon) + \delta m_{e-1}. \tag{A.5}$$

Assume that the cost of the lowest effort level is $c_1 = 0$, so that $m_1 = c_1/\pi_1^S = 0$. The recursive equation (A.5) gives

$$m_e = (1 - \delta^{e-1})(1+\epsilon). \tag{A.6}$$

Note that we do not explicitly solve for the costs c_e because we do not need them to compute the utility and social welfare. If desired, they could be easily derived from the recurrence equation (A.6) and the initial value $c_1 = 0$.

Hence, **IMCP** holds because the ratio of marginal cost to marginal probability of the highest outcome is increasing with the effort level since $\delta < 1$. Furthermore costs are increasing with the effort level because $m_e = (c_e - c_{e-1})/(\pi_e^S - \pi_{e-1}^S) > 0$. Finally, the increase in social welfare from effort level $e-1$ to e is

$$u_e^{\mathrm{SW}} - u_{e-1}^{\mathrm{SW}} = (\pi_e^S - \pi_{e-1}^S)(1 - m_e), \tag{A.7}$$

where $1 - m_e$ is the social benefit over cost per unit increase in probability as the probability of the highest outcome S is increased from π_{e-1}^S to π_e^S. Because probabilities satisfy **FOSD** and m_e is increasing, it suffices to impose that $m_E < 1$ to obtain that **QCSW** holds and the social optimal level is $e^{\mathrm{SO}} = E$. This condition can be equivalently written as $\epsilon < \delta^{E-1}/(1-\delta^{E-1})$, which holds if we let $\epsilon = \delta^E$.

We next determine the social welfare at the first-best and second-best effort levels. The social welfare at $e = 1$ is equal to $u_1^{\mathrm{SW}} = \pi_1^S = \delta^{E-1}$. Using equation (A.6) in equation (A.7) we obtain that the difference in social welfare is $u_e^{\mathrm{SW}} - u_{e-1}^{\mathrm{SW}} = (1-\delta)\delta^{E-1}\big(1 + \epsilon - \epsilon\delta^{1-e}\big)$ and the social welfare at $e = E$ is

$$\begin{aligned} u_E^{\mathrm{SW}} &= u_1^{\mathrm{SW}} + \sum_{e=2}^{E}\left(u_e^{\mathrm{SW}} - u_{e-1}^{\mathrm{SW}}\right) = \big(1 + (E-1)(1-\delta)(1+\epsilon)\big)\delta^{E-1} - \epsilon(1-\delta)\delta^{E-1}\sum_{e=2}^{E}\delta^{1-e} \\ &= \big(1 + (E-1)(1-\delta)(1+\epsilon)\big)\delta^{E-1} - \epsilon\left(1 - \delta^{E-1}\right). \end{aligned}$$

[22] We thank an anonymous referee for suggesting this argument.

Thus, the welfare-loss ratio is given by

$$\rho(I^{\delta}) = \frac{u_E^{SW}}{u_1^{SW}} = 1 + (E-1)(1-\delta)(1+\epsilon) - \frac{\epsilon}{\delta^{E-1}}\left(1-\delta^{E-1}\right) \tag{A.8}$$
$$= 1 + (E-1)(1-\delta)(1+\delta^{E}) - \delta\left(1-\delta^{E-1}\right),$$

where the second equation follows from $\epsilon = \delta^E$. As a result $\lim_{\delta\to 0^+} \rho(I^\delta) = E$ since the second term converges to $E-1$ and the third term converges to zero. Therefore, the upper bound in terms of the number of effort levels is tight because we found a series of problem instances converging to a matching upper bound.

A.8. Tightness of the log bound

In this section we consider a family of problem instances I^E with likelihood ratio π_E^S/π_1^S equal to $r > 1$ whose welfare-loss ratio converges to $1 + \ln(r)$ as the number of effort levels grows to infinity. More formally, we will define a family of problem instances $I^E \in \mathbb{I}_{MIQ}^{r,2}$ with $E > 1$ effort levels and $S = 2$ outcomes satisfying assumptions **MLRP**, **IMCP** and **QCSW** that verify $\lim_{E\to\infty} \rho(I^E) = 1 + \ln(r)$.

Consider the same family of instances from Section A.7 with $\delta = r^{-1/(E-1)}$ so that the likelihood ratio of the highest outcome between the highest and lowest effort levels is given by

$$\frac{\pi_E^S}{\pi_1^S} = \frac{1}{\delta^{E-1}} = r.$$

Clearly we have $\delta \in (0,1)$ because $r > 1$. Letting $\epsilon = \delta^{E-1}/E = 1/(rE)$ we obtain from Section A.7 that the assumptions **MLRP**, **IMCP** and **QCSW** hold (because $0 < \epsilon < \delta^{E-1}/(1-\delta^{E-1})$); the optimal effort level for the principal is $e^* = 1$; and the socially optimal effort level is $e^{SO} = E$. Thus, using our expressions for δ and ϵ in equation (A.8) we obtain that the welfare-loss ratio is

$$\rho(I^E) = 1 + (E-1)\left(1 - r^{-1/(E-1)}\right)\left(1 + \frac{1}{rE}\right) - \frac{1}{E}\frac{r-1}{r},$$

which converges to $1+\ln(r)$ as $E \to \infty$ because $\lim_{n\to\infty} n\left(1 - \sqrt[n]{1/r}\right) = \ln(r)$. The limit follows from L'Hôspital's Rule, taking the derivative of the function r^{-x}.

Appendix B. Supplementary material

Supplementary material related to this article can be found online at http://dx.doi.org/10.1016/j.geb.2015.10.008.

References

Acemoglu, D., Bimpikis, K., Ozdaglar, A., 2009. Price and capacity competition. Games Econ. Behav. 66 (1), 1–26.

Acemoglu, D., Ozdaglar, A., 2007. Competition and efficiency in congested markets. Math. Oper. Res. 32 (1), 1–31.

Alchian, A.A., Demsetz, H., 1972. Production, information costs, and economic organization. Amer. Econ. Rev. 62 (5), 777–795.

Asch, B.J., 1990. Do incentives matter? The case of Navy recruiters. Ind. Lab. Relat. Rev. 43 (3), 89–106.

Babaioff, M., Feldman, M., Nisan, N., Winter, E., 2012. Combinatorial agency. J. Econ. Theory 147 (3), 999–1034.

Babaioff, M., Feldman, M., Nisan, N., 2009. Free-riding and free-labor in combinatorial agency. In: Mavronicolas, M., Papadopoulou, V. (Eds.), Proceedings of the 2nd International Symposium on Algorithmic Game Theory. SAGT'09. In: Lecture Notes in Computer Science, vol. 5814. Springer, pp. 109–121.

Balmaceda, F., Balseiro, S., Correa, J.R., Stier-Moses, N.E., 2012. The cost of moral hazard and limited liability in the principal–agent problem. Working Paper, Columbia DRO-2010-06.

Brown, K.C., Harlow, W.V., Starks, L.T., 1996. Of tournaments and temptations: an analysis of managerial incentives in the mutual fund industry. J. Finance 51 (1), 85–110.

Chevalier, J., Ellison, G., 1999. Career concerns of mutual fund managers. Quart. J. Econ. 114 (2), 389–432.

Conlon, J.R., 2009. Two new conditions supporting the first-order approach to multisignal principal–agent problems. Econometrica 77 (1), 249–278.

Correa, J.R., Schulz, A.S., Stier-Moses, N.E., 2008. A geometric approach to the price of anarchy in nonatomic congestion games. Games Econ. Behav. 64, 457–469.

Demougin, D., Fluet, C., 1998. Mechanism sufficient statistic in the risk-neutral agency problem. J. Inst. Theor. Econ. 154 (4), 622–638.

Ghurye, S., Wallace, D., 1959. A convolutive class of monotone likelihood ratio families. Ann. Math. Stat. 30 (4), 1158–1164.

Grossman, S.J., Hart, O.D., 1983. An analysis of the principal–agent problem. Econometrica 51 (1), 7–45.

Healy, P.M., 1985. The effect of bonus schemes on accounting decisions. J. Acc. Econ. 7 (1–3), 85–107.

Holmström, B., Milgrom, P., 1987. Aggregation and linearity in the provision of intertemporal incentives. Econometrica 55, 303–328.

Holmström, B., Milgrom, P., 1991. Multi-task principal–agent analyses: incentive contracts, asset ownership and job design. J. Law, Econ., Organ. 7, 24–52.

Holmström, B., Tirole, J., 1997. Financial intermediation, loanable funds and growth. Quart. J. Econ. 112 (3), 663–691.

Huynh, H., 1994. A new proof for monotone likelihood ratio for the sum of independent Bernoulli random variables. Psychometrika 59 (1), 77–79.

Innes, R.D., 1990. Limited liability and incentive contracting with ex-ante action choices. J. Econ. Theory 52 (1), 45–67.

Jewitt, I., 1988. Justifying the first-order approach to principal–agent problems. Econometrica 56 (1), 1177–1190.

Jensen, M.C., Meckling, W.H., 1979. Rights and production functions: an application to labor-managed firms and codetermination. J. Bus. 52 (4), 469–506.

Johari, R., Tsitsiklis, J.N., 2004. Efficiency loss in a network resource allocation game. Math. Oper. Res. 29, 407–435.

Kelly, F., 1997. Charging and rate control for elastic traffic. Eur. Trans. Telecommun. 8, 33–37.
Kim, S.K., 1997. Limited liability and bonus contracts. J. Econ. Manag. Strategy 6 (4), 899–913.
Koutsoupias, E., Papadimitriou, C.H., 1999. Worst-case equilibria. In: 16th Annual Symposium on Theoretical Aspects of Computer Science, pp. 404–413. Comput. Sci. Rev. 3 (2) (2009) 65–69.
Laffont, J., Martimort, D., 2001. The Theory of Incentives: The Principal–Agent Model. Princeton University Press.
Laux, C., 2001. Limited liability and incentive contracting with multiple projects. RAND J. Econ. 32, 514–526.
Matthews, Steven A., 2001. Renegotiating moral hazard contracts under limited liability and monotonicity. J. Econ. Theory 97, 1–29.
Mirrlees, J., 1974. Notes on welfare economics, information and uncertainty. In: Balch, D.M.M., Wu, S. (Eds.), Essays in Economic Behavior in Uncertainty. North-Holland, Amsterdam.
Mirrlees, J., 1976. The optimal structure of incentives and authority within an organization. Bell J. Econ. 7, 105–131.
Mirrlees, J., 1999. The theory of moral hazard and unobservable behaviour: part I. Rev. Econ. Stud. 66, 3–21.
Moulin, H., 2008. The price of anarchy of serial, average and incremental cost sharing. Econ. Theory 36 (3), 379–405.
Moulin, H., 2009. Almost budget-balanced vcg mechanisms to assign multiple objects. J. Econ. Theory 144 (1), 96–119.
Nisan, N., Roughgarden, T., Tardos, É., Vazirani, V.V., 2007. Algorithmic Game Theory. Cambridge University Press, Cambridge, UK.
Oyer, P., 1998. Fiscal year ends and nonlinear incentive contracts: the effect on business seasonality. Quart. J. Econ. 113 (1), 149–185.
Poblete, J.J., Spulber, D.F., 2012. The form of incentive contracts: agency with moral hazard, risk neutrality and limited liability. RAND J. Econ. 53 (6), 1357–1367.
Rogerson, W.P., 1985. The first-order approach to principal–agent problems. Econometrica 53 (6), 1357–1367.
Rothschild, M., Stiglitz, J.E., 1970. Increasing risk: I. A definition. J. Econ. Theory 2 (3), 225–243.
Roughgarden, T., Tardos, É., 2004. Bounding the inefficiency of equilibria in nonatomic congestion games. Games Econ. Behav. 47 (2), 389–403.
Salanié, B., 2003. Testing contract theory. CESifo Econ. Stud. 49 (3), 461–477.
Shapiro, C., Stiglitz, J.E., 1984. Equilibrium unemployment as a worker discipline device. Amer. Econ. Rev. 74 (3), 433–444.
Spence, M., Zeckhauser, R., 1971. Insurance, information, and individual action. Amer. Econ. Rev. 61 (2), 380–387.
Stiglitz, J.E., 1974. Incentives and risk sharing in sharecropping. Rev. Econ. Stud. 41 (2), 219–255.