Contracting with Moral Hazard: A Review of Theory & Empirics*

George Georgiadis†

August 26, 2022

Abstract

This article surveys the literature on principal-agent problems with moral hazard that gained popularity following the seminal works of Mirrlees (1976), Holmström (1979), and others. This literature is concerned with designing incentives to motivate one or more workers—typically by paying for performance—in settings where his actions cannot be directly contracted on. I begin with the canonical framework, and then in Section 2 I categorize the subsequent literature that relaxes various assumptions of this framework. Section 3 focuses on the empirical research, and Section 4 discusses other strands of the literature. My goal is to provide a broad and concise overview—sometimes at the expense of depth, and focus on the economic insights (rather than methodological contributions).

1 Foundations

This article provides a broad but concise overview of the agency literature under moral hazard, which is concerned with designing incentives—typically performance pay—to motivate one or more workers.

In the canonical framework of Holmström (1979) the game begins with the principal offering a *contract* w that specifies the agent's pay as a function of output produced. The agent then decides whether to accept the contract, in which case he chooses a costly action (effort), or he rejects the contract and collects his outside option, which corresponds to seeking employment elsewhere. Finally, output is generated according to a distribution function that depends on the agent's action, payoffs are realized, and the game ends. The principal's objective is to maximize her expected profit, while the agent chooses an action that maximizes his expected payoff. Mathematically, this problem can be formulated as follows:

$$\max_{w(\cdot),a\geq 0} \mathbb{E}[x - w(x)|a] \tag{P}$$

$$\text{s.t. } a \in \arg\max_{a'\geq 0} \mathbb{E}[u(w(x))|a'] - c(a') \tag{IC}$$

$$\mathbb{E}[u(w(x))|a] - c(a) \geq \underline{u}, \tag{IR}$$

*I am grateful to Enrique Ide, Niko Matouschek, Michael Powell, as well as to Peter Traber for excellent research assistance.

†Kellogg School of Management, Northwestern University. Email: g-georgiadis@kellogg.northwestern.edu.

where x denotes output, a the agent's action, u his strictly increasing and concave utility function, c his strictly increasing and convex cost of choosing each action, and $\underline{u}$ the value of his outside option. The expectation is taken with respect to a family of probability density functions f that map each action into output; i.e., $x \sim f(\cdot|a)$. All parameters are common knowledge. The principal here chooses a contract and a recommended action to maximize her profits subject to two constraints: (IC) stipulates that the recommended action is incentive compatible given w, and (IR) requires that the agent prefers the offered contract to his outside option.

As a benchmark, consider the *first-best outcome*; i.e., the one that maximizes total welfare. This corresponds to the solution of (P) when (IC) is omitted. Because the agent dislikes risk while the principal is risk-neutral, the first-best contract features constant pay.[1]

To solve this program, it is standard to replace (IC) with its first-order condition, and decompose the problem into two steps (Grossman and Hart, 1983).[2] In the first, one finds for each action the contract that minimizes the principal's expected cost. This has been shown to satisfy

$$\frac{1}{u'(w(x))} = \lambda + \mu \frac{f_a(x|a)}{f(x|a)} \text{ for all } x,$$

where f_a is the derivative of f with respect to a, and λ and $\mu > 0$ are the Lagrange multipliers associated with (IR) and the first-order condition of (IC), respectively. In the second step, one solves for the profit-maximizing action.[3]

Intuitively, the cost-minimizing contract attaches large payments to *informative-but-rare* outputs; i.e., outputs for which $f_a(x|a)$ is large, signifying that it is highly likely that the agent indeed chose the target action a rather than $a - \epsilon$ (which is the relevant deviation given the first-order approach and $\mu > 0$), and outputs for which $f(x|a)$ is small so the principal must make these large payments only rarely. Moreover, the variation in payments is decreasing in the curvature of u'; that is, the optimal contract is shaped by a trade-off between providing incentives to the agent and insuring him from volatility in his pay which, being risk-averse, he dislikes.[4]

A key insight from this model is the *informativeness principle*, which asserts that any signal that is informative of the agent's action should be incorporated into the contract. This has several implications. If, for example, a CEO's current efforts impact future profits, then incentive pay should also depend on lagged profits. In similar vein, if multiple executives are subject to common market-wide shocks, then it is valuable to evaluate them relative to one another. See also Holmström (2017) for a more detailed analysis of this framework and discussion.

[1] This outcome can also be achieved in the presence of (IC) if (i) output is deterministic, or equivalently, effort is contractible, (ii) the distribution that maps a to x has shifting support, or (iii) the agent is risk-neutral. See Holmström (2017) for details.

[2] Replacing (IC) with its first-order condition, $\mathbb{E}[u(w(x))f_a(x|a)/f(x|a)|a] = c'(a)$, is without loss of generality if and only if the agent's expected payoff is differentiable and single-peaked in a. Rogerson (1985a) and Jewitt (1988) provide sufficient conditions for this to be the case.

[3] The second-step program, barring stringent conditions (Jewitt, Kadan and Swinkels, 2008), is not convex but can be solved via line-search. Luckily most of the insights can be gleaned from the solution of the first step.

[4] Note that ex post, the principal has no uncertainty about the agent's action. Yet to provide him with ex ante incentives to choose a particular action, the contract must reward him generously for some output realizations and punish him for others.

2 Towards More Realistic Models

We now turn to the literature that has sought to relax various assumptions of the above framework.

2.1 Flexible Actions

In the canonical model, the agent chooses a one-dimensional action which determines the entire distribution of output. In reality agents have at least some flexibility in determining distributions over outcomes in settings including professional investing (Chevalier and Ellison, 1997), management (Shue and Townsend, 2017 and Rahmandad, Henderson and Repenning, 2018), and sales (Oyer, 1998 and Larkin, 2014).

Holmström and Milgrom (1987) were one of the first to model flexible actions. Take a salesperson who is paid according to quarterly sales. Each morning, he checks his sales in the current quarter and chooses how much effort to spend. Under certain conditions, the optimal contract is linear in quarterly sales and induces constant effort over time. Intuitively, if the contract were nonlinear, the agent would be able to exploit it at the principal's expense. Suppose for example that it pays a fixed wage plus a lump-sum bonus if sales exceed a quota. If the agent hits the quota prior to the end of the quarter, he has no incentive to keep working; if he deems it unlikely to meet the quota, he likewise gives up. A linear contract is robust to this type of behavior because each sale is rewarded the same regardless of *when* it was realized. This result is often used to motivate restricting attention to linear contracts in static models, which allows one to solve the entire problem analytically and obtain testable comparative statics; see for example Prendergast (1999).

Others have studied static models where the agent can influence the entire distribution of output.[5] In Palomino and Prat (2003) where the agent controls the first two moments of his output distribution, a bonus contract is optimal; i.e., one with binary payout. In Diamond (1998) the agent can choose any distribution at a cost that depends on its mean, whereas in Barron, Georgiadis and Swinkels (2020) the agent can costlessly add mean-preserving noise to his output. In both cases the optimal contract is linear. Hebert (2017) microfounds a model where the agent can choose any output distribution at a cost that is proportional to its Kullback-Leibler divergence relative to some default distribution and shows that a piece-wise linear contract akin to a call option is optimal, while Georgiadis, Ravid and Szentes (2022) studies a model without functional form assumptions on the agent's cost of choosing a distribution.[6]

2.2 Multiple Tasks

Most jobs comprise many tasks; e.g., salespeople selling multiple products, employees controlling both quantity and quality of output, managers splitting their time between their own work and training

[5] Formally, this can be modeled by letting the agent choose the output distribution F to maximize $\int u(w(x))dF(x) - C(F)$, where $C(F)$ is the agent's private cost of choosing F. Researchers have explored different restrictions on C to suit particular applications while maintaining tractability.

[6] Another way to enrich the agent's action space is to let the resolution of uncertainty be observed by the agent *before* he chooses his action; e.g., he might learn how difficult the task is (Edmans and Gabaix, 2011). As in other models, this leads to *simple* contracts being optimal.

others, and so on. Moreover, sometimes employees can improve their performance measures with legitimate effort as well as with undesirable or illegal tactics.

Holmström and Milgrom (1991) proposes the following framework to model such situations. The agent now chooses a vector of n actions, denoted by $\mathbf{a}$, by bearing cost $c(\mathbf{a})$.[7] The principal has a vector of m performance measures denoted $\mathbf{x} = M\mathbf{a} + \epsilon$, where M is an $m \times n$ matrix and ϵ is m-dimensional Gaussian, zero-mean noise. Assume the principal's profit is $\boldsymbol{\omega} \cdot \mathbf{a}$ net of the agent's pay, where $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_n)$ represents how much the principal values each action, and following Holmström and Milgrom (1987), she restricts attention to linear contracts of the form $w(\mathbf{x}) = b + \boldsymbol{\alpha} \cdot \mathbf{x}$, where b is a wage and $\boldsymbol{\alpha}$ is the commission rate for each dimension of performance.

This model can be solved analytically with some additional assumptions, and generates several insights. First, if some actions are measured with more noise than others, then it is optimal to provide lower-powered incentives on all measures; i.e., to choose smaller α_i's. As an example, let the actions represent the quantity and quality of output, where the latter can only be measured with noise. Given the risk-incentives trade-off, other things equal, the principal would like to offer higher-powered incentives for quantity than for quality. However, because the tasks are substitutes, doing so will lead the agent to substitute effort away from quality and into quantity at the detriment of the principal. In such cases, it is optimal to mute incentives for both quality and quantity.

The second insight applies to settings where none of the performance measures are perfectly aligned with what the principal truly cares about; i.e., none of the rows of M are collinear with $\boldsymbol{\omega}$ (Baker, 1992). In such cases, the optimal power of incentives (i.e., α) increases in the alignment of the performance measures with the principal's preferences. Take for example Volkswagen, whose leadership set out a plan in 2008 to drastically increase sales in the United States, and promised executives and engineers lucrative incentives for meeting certain targets. Those targets could be achieved with product improvements, or alternatively, by cheating. As it turned out however, product improvements were much harder than the (perceived) disutility associated with cheating, and the high-powered incentives led employees to mask the emissions of their diesel engines.[8] Other examples include doctors cherry-picking patients (Dranove et al., 2003), the Wells Fargo fake-accounts scandal, the BP Texas City refinery explosion, and the V.A scandal involving manipulated medical waiting lists. See also Holmström (2017) for further details and Bonham and Riggs-Cragun (2021) for a multitasking model with flexible actions in the spirit of Section 2.1.

2.3 Dynamics

An extension which has received a lot of attention is one where the players interact over multiple periods. Dynamic models come broadly in two flavors: ones in which the principal can commit to a contract at the beginning of the relationship that specifies payments to the agent in each period as a function of current and past output, and models in which she cannot—typically referred to as

[7]The interesting case arises when c is a supermodular function. For example, if a_1 and a_2 represent quantity and quality, respectively, this assumption implies that the tasks are substitutes; i.e., the marginal cost of providing quality increases in quantity, and vice versa.

[8]Such a situation can be modeled by letting a_1 and a_2 represent effort towards product improvements and cheating, respectively, $x = a_1 + a_2 + \epsilon$ denote sales, and $\omega_1 > 0 > \omega_2$ represent the firm's weights on the two dimensions of effort.

relational contracting models.

With commitment. Early contributions are due to Rogerson (1985*b*) and Spear and Srivastava (1987), who study a two-period and infinite-horizon model, respectively. An important insight is that the optimal contract has memory and moreover, it *backloads* incentives: good performance in the current period is rewarded with a larger immediate payment, *and* larger payments in future periods. DeMarzo and Sannikov (2006) and Sannikov (2008) reconsider this problem with a risk-neutral and a risk-averse agent, respectively, whose effort controls the drift of a Brownian motion. Their continuous-time formulation permits a sharp characterization of the optimal contract and several insights—see Athey and Skrzypacz (2017) for an overview. Others including Myerson (2015) and Green and Taylor (2016), analyze this problem with Poisson production technologies, which leads to optimal contracts with somewhat different properties.

A strand of this literature studies models in which the agent can manipulate output intertemporally. Edmans et al. (2012) considers a tractable model in which the agent can engage in short-termism by inefficiently increasing current stock price returns at the expense of future returns. To deter this behavior, the optimal contract has the property that the return in any given period affects log pay in future periods at an increasing rate. In DeMarzo et al. (2014) and Wong (2019), the agent can increase current output at the risk of a low-probability but large-in-magnitude loss the next instant. A takeaway is that deterring risk-taking following a period of poor performance may be unavoidable. See also Makarov and Plantin (2015) for a model with risk-taking where the principal learns about the agent's ability over time.

Without commitment. Performance evaluations are often based on "soft" information such as a subjective evaluation by a supervisor. In these cases, performance-contingent contracts cannot be enforced by a court, so firms must provide incentives on a relational basis. To capture such a situation, MacLeod and Malcomson (1989) considers a discrete-time model in which the principal can voluntarily pay the agent every period after observing his effort. As long as players are sufficiently patient, there exists a stationary equilibrium where in each period, the principal "recommends" a particular effort, the agent follows the recommendation, and receives a "bonus" in return. If the principal ever reneges on paying the bonus—off the equilibrium path—the agent will choose the cheapest effort in every subsequent period; i.e., he plays a grim-trigger strategy. Intuitively, it is in the principal's interest to pay the bonus as long as its cost is smaller than her value of maintaining the relationship.[9] This framework has been extended to incorporate multiple agents (Levin, 2002), unobservable efforts and noisy output (Levin, 2003), endogenous feedback (Fuchs, 2007), multiple agents and relational incentives coupled with formal contracts (Rayo, 2007), asymmetric information (Halac, 2012 and Li and Matouschek, 2013), multiple agents with bilateral monitoring (Andrews and Barron, 2016 and Barron and Powell, 2019), and limited liability (Fong and Li, 2017).

[9] If the game comprises a single period, then the only way to motivate the agent is with a scheme that involves money-burning (MacLeod, 2003).

2.4 Unknown Production Environment

Another strand relaxes the assumption that the production environment is perfectly known by all parties. These models fall in three categories: those in which the principal designs a mechanism to elicit the agent's private information, those where the principal chooses a contract to maximize her worst-case payoff, and career concern models where both players learn about the production environment over time.

Moral Hazard & Adverse Selection. The first features models where the principal is oblivious of one parameter; e.g., the agent's cost function is $c(a, \theta)$ where θ is the agent's private information. Guesnerie and Laffont (1984) studies this problem assuming output is deterministic so that conditional on θ, the agent's action is observable. Chade and Swinkels (2022) incorporates noisy output and gives conditions such that the principal's problem can be disentangled into a pure adverse selection and a pure moral hazard problem. The optimal mechanism involves a menu of output-contingent contracts with each type self-selecting into his designated one; see also Foarta and Sugaya (2021). When the agent is cash constrained, all types are offered the same contract under general conditions if he is risk neutral (Gottlieb and Moreira, 2022), and under more stringent conditions if he is risk averse (Castro-Pires and Moreira, 2021).[10]

Robust Contracts. When multiple parameters are unknown, the mechanism design approach employed above becomes intractable, and even forming a prior might be impractical. Researchers have therefore studied "robust contracts" that provide the principal with the largest payoff guarantee—this is as if an adversarial third party can choose the unknown aspects of the environment after observing the offered contract to minimize the principal's payoff; so the principal seeks a contract to maximize this worst-case payoff. If the principal knows only a subset of the actions available to the agent and their costs, then a linear contract is optimal (Carroll, 2015); under monotone-likelihood-ratio restrictions on the adversarial party's action space, the optimal contract is a mixture of equity and debt (Antic, 2022). Walton and Carroll (2019) provides general conditions for linear contracts to be optimal, Chassang (2013) and Miao and Rivera (2016) embed these ideas in dynamic settings, and Dai and Toikka (2022) incorporates team production. See also Carroll (2019) for an overview.

Career Concerns. A key source of one's motivation is often tied to their reputation: what others think of their abilities, where a favorable perception can lead to promotions, outside offers, and pay raises. Holmström (1999) considers a model with two-sided uncertainty where both the principal and the agent are learning about the agent's ability over time from past performance, and the agent's wage increases in the principal's perception of his ability. Absent incentive contracts, career concerns induce an inefficient pattern of effort over time: the agent works inefficiently hard early on while the principal is learning about their ability, and not hard enough later on when the principal's belief is hard to influence. See also Gibbons and Murphy (1992) for a contracting model with career concerns,

[10] Also related is Garrett et al. (2021) where the agent designs his production technology—by choosing the cost of each output distribution—before contracting with the principal.

Dewatripont, Jewitt and Tirole (1999) for a generalized model, and Cisternas (2018) for a model with endogenous skills acquisition.

2.5 Endogenous Performance Measures

To design an incentive plan, one must first identify appropriate performance measures, and then specify the relationship between pay and those measures. Dye (1986) considers a model a-lá Holmström (1979) where after observing a signal of the agent's effort, the principal can acquire another costly signal. In Georgiadis and Szentes (2020), the principal chooses a strategy for sequentially acquiring signals and a wage scheme as a function of the acquired signals, and shows that a bonus contract is optimal. Hoffmann, Inderst and Opp (2020) studies deferred compensation: signals about the agent's one-shot effort arrive over time and the principal decides when to stop and pay the agent according to the observed signals. The central trade-off in these models is that acquiring more signals is costly, but provides the principal with more accurate information, reducing agency costs.

Khalil and Lawarree (1995) considers a pure adverse selection model where the principal can monitor either the agent's input (e.g., hours worked) or his output. When these monitoring technologies are equally informative, the former is preferable if and only if the principal is the residual claimant of output; see also Baker (2002). Datar, Kulp and Lambert (2001) shows that if one restricts attention to linear contracts, then it may be optimal to ignore some informative signals. In Li and Yang (2020) the principal chooses a partition of the signal-space at a cost that increases in the fineness of the partition and a wage scheme as a function of the cell of the partition in which the signal lies.

2.6 Many Agents

Another significant strand studies models with multiple agents and externalities. These can be positive, if agents work in teams for example, or negative, if they are evaluated and paid relative to one another, as in a tournament or a stack ranking system.

Teams. When people work in teams, disentangling each person's contributions is often difficult (Alchian and Demsetz, 1972). This may then necessitate incentivizing the group as a whole. In a seminal contribution, Holmström (1982) shows that when only joint output is contractible, there exists no sharing rule that implements the efficient output. The intuition is that because agents don't internalize the marginal benefit of their efforts on others, they work too little from the perspective of joint surplus; i.e., they freeride. When efforts are complementary, Winter (2004) shows that the cheapest wage scheme that induces identical agents to work in the unique equilibrium necessarily treats each agent differently; Halac, Lipnowski and Rappoport (2021) further shows that this design hinges on wage schemes being public. Admati and Perry (1991) and Fershtman and Nitzan (1991) extend these ideas to dynamic settings and show that freeriding is reduced if but only if players' efforts are strategic complements across time.[11]

[11]See also Georgiadis (2014) for a continuous-time formulation which provides testable comparative statics, and Cvitanic and Georgiadis (2016) for a mechanism to eliminate freeriding.

Relative Performance Evaluations & Contests. The productivities of workers are often subject to common shocks, that is, the external factors that affect one worker's performance often affect others' performance in similar ways. For example, the productivities of executives across firms or of salespeople within a firm may be jointly affected by the broader economic conditions. In such cases, relative performance evaluation and pay can serve to filter out these shocks hence reducing the riskiness of pay (Holmström, 1982).[12] This observation initiated the literature on tournaments starting with Lazear and Rosen (1981), who characterizes optimal incentive schemes that condition pay on each agent's ordinal rank rather than absolute performance. Green and Stokey (1983) shows that rank-order tournaments dominate individual contracting if the common shock is sufficiently diffuse. In general however, by Holmström's Informativeness Principle, conditioning pay on both absolute and relative performance is optimal (Mookherjee, 1984). Rosen (1986) studies sequential pairwise elimination tournaments commonly encountered in athletic events, while Gradstein and Konrad (1999) shows that depending on the relative importance of effort to random factors, the optimal structure either comprises a single stage, or it is a pairwise elimination tournament.[13] Another benefit of tournaments is that they may facilitate commitment when performance evaluations are subjective: it is easier for the principal to commit to award a fixed prize to the best-performing agent no matter what than to commit to an evaluation-contingent contract (Malcomson, 1984).

3 Empirical Literature

The empirical literature on contract theory can be broadly classified into three buckets. The first comprises studies that evaluate the degree to which workers respond to incentives. The second features papers that examine whether observed contracts have the properties predicted by theory. The third bucket includes papers that *operationalize* different theories, for example by building structural models, recovering the unknown parameters, and computing an optimal contract targeted at real applications.

Responsiveness to Incentives. Several papers use quasi-experimental or experimental variation to show that higher-powered incentives cause individuals to work harder, at least on the dimensions that are highly rewarded. These effects have been found in a variety of settings, including windshield repairers (Lazear, 2000), tree planters (Shearer, 2004), bicycle messengers (Fehr and Goette, 2007), gig workers (Chen and Sheldon, 2016), salespeople (Chung and Narayandas, 2017), workers on Amazon's Mechanical Turk platform (DellaVigna and Pope 2018; 2022), creative tasks (Englmaier et al., 2018), and day laborers (Guiteras and Jack, 2018) in developing countries. An important finding is that incentives affect not only motivation, but also selection. In particular, when a windshield installation firm switched from hourly wages to piece rates, productivity increased by 44%, half of which was attributable to increased motivation and the other half to selection; that is, more productive workers

[12] Another way to filter out common shocks is by evaluating performance relative to a benchmark; e.g., portfolio managers' returns are often evaluated against an index such as the S&P 500.

[13] See also Harris and Vickers (1987) for a dynamic tug-of-war, Taylor (1995) for an optimal dynamic contest, and more recently, Halac, Kartik and Liu (2017) and Ely et al. (2022) for dynamic contests where the principal can curate the agents' beliefs about their own and their rivals' progress.

joining the firm and less productive ones leaving (Lazear, 2000). See also Paarsch and Shearer (2000) and Lo, Ghosh and Lafontaine (2011), as well as Lazear (1986) for a theoretical argument.

Another strand studies multitask settings and how imbalanced incentives can lead to adverse outcomes. Chevalier and Ellison (1997) and Shue and Townsend (2017) presents evidence of mutual fund managers and CEOs taking excessive risks, respectively. Oyer (1998) and Larkin (2014) find that nonlinear incentives led salespeople to manipulate the timing of their sales, and being paid according to sales revenue, they offered excessive discounts. See also Jensen (2001) for additional examples. Tzioumis and Gee (2013) presents evidence that mortgage officers approved marginal applications with higher delinquency rates when close to making their quota at the end of the month. Dranove et al. (2003) finds that when the states of New York and Pennsylvania began publishing health care report cards with surgeon-specific mortality rates, doctors stopped doing certain procedures for very ill patients. In a manufacturing setting, Hong et al. (2018) finds that higher piece rates had a positive effect on quantity at the expense of a large increase in defect rates.[14]

A manifestation of the multitask problem concerns settings where a worker's pay depends both on his and others' productivity; e.g., under relative performance or team pay. Bandiera, Barankay and Rasul (2005) presents evidence that relative performance pay for agricultural workers led to lower productivity vis-a-vis individual piece rates if (but only if) workers could easily monitor each other. Bandiera, Barankay and Rasul (2007) on the other hand, finds that paying managers a performance bonus if aggregate productivity was sufficiently high led to them working harder to increase the productivity of their subordinates; see Itoh (1991) for a theoretical treatment. Berger, Harbring and Sliwka (2013) finds positive effects of a forced distribution system in a real effort experiment.

Are Observed Contracts Optimal? The risk-incentives trade-off asserts that, other things equal, pay-for-performance sensitivity should be lower if the performance measure is more volatile or the worker is more risk averse. However, evidence of this trade-off coming from various occupations is, at best, weak (Prendergast, 2002). Prendergast posits three reasons why pay-for-performance sensitivity may be positively correlated with risk. The first has to do with performance pay being more prominent in more uncertain environments where the choice of what tasks to work on is delegated to the employee—as opposed to a manager telling them what to do and paying fixed wages. The second is that performance reviews are often subjective and are more likely to be biased when stakes are higher. If the environment is noisier, this bias is less costly for the firm, and so higher-powered incentives are optimal. The third reason is that performance investigations are often only triggered if poor decision-making is suspected; e.g., following an accident. If the initial impression that triggers an investigating is noisier, then the worker suffers less from the investigation and so their ex ante incentives are muted. In response, the firm optimally provides stronger incentives.

A second prediction has to do with the shape of the pay-for-performance relationship. In practice, this is typically *simple*—for example efficiency wages, bonus contracts, linear and piecewise linear

[14]Another example comes from the FBI: In the years after 9/11, leadership discovered that the field offices that did the best intelligence work also wrote the most reports on intelligence matters. So in 2007 the FBI started evaluating field offices based on pages of intelligence produced. By 2008 the positive relationship between reports produced and quality had disappeared—agents were simply filling their word-count quotas (Fisman and Sullivan, 2015).

contracts are most common (Prendergast, 1999). In contrast, the shape of the optimal contract in canonical models depends on the primitives and is simple only under specific assumptions (Grossman and Hart, 1983). However, linear contracts have been shown to be *close* to optimal in the canonical Holmström model under relatively mild assumptions (Bose, Pal and Sappington, 2011).[15] Moreover, researchers have proposed richer models that can explain the optimality of simple contracts. In particular, the optimal contracts in Holmström and Milgrom (1987), Carroll (2015) and Barron, Georgiadis and Swinkels (2020) are linear, in Antic (2022) and Hebert (2017) they are piecewise-linear, whereas in MacLeod (2003) and Georgiadis and Szentes (2020) they comprise a fixed wage and a single bonus.

Operationalizing the Theory. Misra and Nair (2011) estimates a dynamic structural contracting model and uses it to evaluate counterfactual contracts for the salesforce of a large contact lens manufacturer; see also Chung, Steenburgh and Sudhir (2014). Chung, Kim and Park (2021) estimates a structural model that includes performance pay, as well as training, recruiting and termination decisions, and applies it to a sales dataset to compute counterfactual contracts.[16] Georgiadis and Powell (2022) establishes conditions under which the information provided by an A/B test of incentive contracts is sufficient for determining how best to improve a status quo incentive contract in the classic Holmström model, and assesses the empirical relevance of these conditions using data from a real effort experiment.[17]

4 Discussion

We conclude with a discussion of additional related literatures and an outline of previous surveys.

CEO Pay. Several researchers have sought to develop and test models tailored to CEO pay. Edmans, Gabaix and Landier (2009) proposes a model with multiplicative preferences, and derives an empirical measure of incentives that is calculated as the dollar change in CEO wealth for a percentage change in firm value divided by annual flow compensation. This measure is independent of firm size, and thus comparable across firms and over time.[18] See also Murphy (1999), Murphy (2013), and Edmans, Gabaix and Jenter (2017) for surveys of theoretical and empirical findings on CEO pay, and Edmans, Gosling and Jenter (2021) for a survey of firm directors and investors about the practical considerations in designing and evaluating pay packages.

Efficiency Wages. Efficiency wage theories argue that firms may raise wages above competitive levels to motivate their workforce, to attract and retain more productive workers, or to foster workers'

[15]See also D'Haultfoeuille and Fevrier, 2020, who uses variation in contracts to estimate the losses associated with using linear contracts when workers are risk neutral.

[16]For details on common estimation methodologies, see Hotz and Miller (1993), Bajari, Benkard and Levin (2007), and Arcidiacono and Miller (2011).

[17]See also Antic and Georgiadis (2022) who reconsiders this question in a setting where the principal is oblivious to the agent's action set and their costs, and evaluates contracts by their worst-case payoff. In this case, the optimal contract is a convex combination of one of the original contracts and a linear one.

[18]This measure is available for a large number of firms at https://alexedmans.com/data/. The authors argue that multiplicative preferences are necessary to deliver empirically consistent predictions for the scaling of various incentive measures with firm size.

loyalty toward the firm (Raff and Summers, 1987 and Krueger and Summers, 1988). In particular, paying supra-competitive wages increases one's cost of losing their job which, coupled with appropriate monitoring and firing policies, can motivate higher effort (Shapiro and Stiglitz, 1984). A critique however is that efficiency wages are typically dominated by other contractual forms (Murphy and Topel, 1990).

Behavioral Frameworks. A substantial literature investigates behavioral responses in contracting settings. Psychologists and sociologists argue that performance pay is counterproductive because it undermines "intrinsic motivation"; see for example Gerhart (2017) and the references therein. To rationalize this view, Benabou and Tirole (2003) shows how performance incentives offered by an informed principal can adversely impact an agent's perception of the task or of his own ability, and ultimately undermine his motivation. For example, the agent might rationally infer that he is being paid a bonus because the task is unenjoyable, and this inference can crowd out his intrinsic motivation. See also Benabou and Tirole (2006) for an alternative model. Another strand studies agency problems with non-standard preferences. For example, in Englmaier and Wambach (2010) there are many inequity-averse agents, whereas in Englmaier and Leider (2012) the agent has reciprocal motivations. See also Fehr, Goette and Zehnder (2009) for a survey of the literature with fairness concerns, and Kamenica (2012) for a review of empirical findings on the anomalous impacts of incentives.

Payments in Kind. Pecuniary incentives such as bonuses, commissions, piece rates, and promotions are only one way to motivate workers. Firms often complement them with various non-monetary incentives. The "delegation literature" explores the motivating effects of granting subordinates decision-making power; see, for example, Aghion and Tirole (1997), Bester (2009), Li, Matouschek and Powell (2017), and Lipnowski and Ramos (2020) for theoretical analyses, and Fehr, Herz and Wilkening (2013) for an experiment. Others have studied the motivating effects of paying (partially) with knowledge such as apprenticeships (Fudenberg, Georgiadis and Rayo, 2021). A more recent strand studies how feedback and information can complement financial incentives to reduce agency costs (Fuchs, 2007 and Ely, Georgiadis and Rayo, 2022). Another strand explores the motivating effects of symbolic awards; see for example Gallus (2017) for an experiment at Wikipedia, and Frey and Gallus (2017) for a survey.

Earlier Surveys and Reviews. Several scholars have compiled excellent surveys. Bolton and Dewatripont (2004) provides the reference textbook. Gibbons (1998) focuses on multitask problems and subjective performance evaluations. In Prendergast (1999) the emphasis is on whether the empirical evidence is consistent with theoretical predictions. Chiappori and Salanie (2003) reviews the empirical literature, Oyer, Schaefer et al. (2011) focuses on how firms (can) attract and hire the right types of workers, while Lazear (2018) surveys the evidence on how workers respond to different kinds of incentives. Garicano and Rayo (2016) describes reasons why organizations fail, many of which are rooted in incentive conflicts. Finally, Gibbs (2016) features a panel discussion among prominent scholars on the state of the contracting literature and important open questions.

References

Admati, Anat R., and Motty Perry. 1991. "Joint Projects without Commitment." *The Review of Economic Studies*, 58(2): 259–276.

Aghion, Philippe, and Jean Tirole. 1997. "Formal and real authority in organizations." *Journal of Political Economy*, 105(1): 1–29.

Alchian, Armen A, and Harold Demsetz. 1972. "Production, information costs, and economic organization." *American Economic Review*, 62(5): 777–795.

Andrews, Isaiah, and Daniel Barron. 2016. "The Allocation of Future Business: Dynamic Relational Contracts with Multiple Agents." *American Economic Review*, 106(9): 2742–59.

Antic, Nemanja. 2022. "Contracting with unknown technologies." *Working Paper.*

Antic, Nemanja, and George Georgiadis. 2022. "Robust Contracts: A Revealed Preference Approach." *Working Paper.*

Arcidiacono, Peter, and Robert A Miller. 2011. "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity." *Econometrica*, 79(6): 1823–1867.

Athey, Susan, and Andrzej Skrzypacz. 2017. "Yuliy Sannikov: Winner of the 2016 Clark Medal." *Journal of Economic Perspectives*, 31(2): 237–56.

Bajari, Patrick, C Lanier Benkard, and Jonathan Levin. 2007. "Estimating dynamic models of imperfect competition." *Econometrica*, 75(5): 1331–1370.

Baker, George. 2002. "Distortion and risk in optimal incentive contracts." *Journal of Human Resources*, 728–751.

Baker, George P. 1992. "Incentive contracts and performance measurement." *Journal of Political Economy*, 100(3): 598–614.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2005. "Social Preferences and the Response to Incentives: Evidence from Personnel Data*." *The Quarterly Journal of Economics*, 120(3): 917–962.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2007. "Incentives for managers and inequality among workers: Evidence from a firm-level experiment." *The Quarterly Journal of Economics*, 122(2): 729–773.

Barron, Daniel, and Michael Powell. 2019. "Policies in relational contracts." *American Economic Journal: Microeconomics*, 11(2): 228–49.

Barron, Daniel, George Georgiadis, and Jeroen Swinkels. 2020. "Optimal contracts with a risk-taking agent." *Theoretical Economics*, 15(2): 715–761.

Benabou, Roland, and Jean Tirole. 2003. "Intrinsic and extrinsic motivation." *The Review of Economic Studies*, 70(3): 489–520.

Benabou, Roland, and Jean Tirole. 2006. "Incentives and Prosocial Behavior." *American Economic Review*, 96(5): 1652–1678.

Berger, Johannes, Christine Harbring, and Dirk Sliwka. 2013. "Performance appraisals and the impact of forced distribution: an experimental investigation." *Management Science*, 59(1): 54–68.

Bester, Helmut. 2009. "Externalities, communication and the allocation of decision rights." *Economic Theory*, 41(2): 269–296.

Bolton, Patrick, and Mathias Dewatripont. 2004. *Contract theory.* MIT press.

Bonham, Jonathan, and Amoray Riggs-Cragun. 2021. "Contracting on What Firm Owners Value." *Working Paper.*

Bose, Arup, Debashis Pal, and David EM Sappington. 2011. "On the performance of linear contracts." *Journal of Economics & Management Strategy*, 20(1): 159–193.

Carroll, Gabriel. 2015. "Robustness and Linear Contracts." *American Economic Review*, 105(2): 536–63.

Carroll, Gabriel. 2019. "Robustness in mechanism design and contracting." *Annual Review of Economics*, 11(1): 139–166.

Castro-Pires, Henrique, and Humberto Moreira. 2021. "Limited liability and non-responsiveness in agency models." *Games and Economic Behavior*, 128: 73–103.

Chade, Hector, and Jeroen Swinkels. 2022. "Disentangling moral hazard and adverse selection." *Working Paper.*

Chassang, Sylvain. 2013. "Calibrated Incentive Contracts." *Econometrica*, 81(5): 1935–1971.

Chen, M Keith, and Michael Sheldon. 2016. "Dynamic Pricing in a Labor Market: Surge Pricing and Flexible Work on the Uber Platform." 455–455.

Chevalier, Judith, and Glenn Ellison. 1997. "Risk Taking by Mutual Funds as a Response to Incentives." *Journal of Political Economy*, 105(6): 1167–1200.

Chiappori, Pierre Andre, and Bernard Salanie. 2003. "Testing Contract Theory: a survey of some recent work." *Advances in Economics and Econometrics*, 1.

Chung, Doug J, and Das Narayandas. 2017. "Incentives versus reciprocity: Insights from a field experiment." *Journal of Marketing Research*, 54(4): 511–524.

Chung, Doug J., Byungyeon Kim, and Byoung G. Park. 2021. "The Comprehensive Effects of Sales Force Management: A Dynamic Structural Analysis of Selection, Compensation, and Training." *Management Science*, 67(11): 7046–7074.

Chung, Doug J., Thomas Steenburgh, and K. Sudhir. 2014. "Do Bonuses Enhance Sales Productivity? A Dynamic Structural Analysis of Bonus-Based Compensation Plans." *Marketing Science*, 33(2): 165–187.

Cisternas, Gonzalo. 2018. "Career concerns and the nature of skills." *American Economic Journal: Microeconomics*, 10(2): 152–89.

Cvitanic, Jakša, and George Georgiadis. 2016. "Achieving efficiency in dynamic contribution games." *American Economic Journal: Microeconomics*, 8(4): 309–42.

Dai, Tianjiao, and Juuso Toikka. 2022. "Robust incentives for teams." *Econometrica*, 90(4): 1583–1613.

Datar, Srikant, Susan Cohen Kulp, and Richard A. Lambert. 2001. "Balancing Performance Measures." *Journal of Accounting Research*, 39(1): 75–92.

DellaVigna, Stefano, and Devin Pope. 2018. "What Motivates Effort? Evidence and Expert Forecasts." *The Review of Economic Studies*, 85(2): 1029–1069.

DellaVigna, Stefano, and Devin Pope. 2022. "Stability of Experimental Results: Forecasts and Evidence." *American Economic Journal: Microeconomics*, 14(3): 889–925.

DeMarzo, Peter M., and Yuliy Sannikov. 2006. "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model." *The Journal of Finance*, 61(6): 2681–2724.

DeMarzo, Peter M, Dmitry Livdan, Alexei Tchistyi, et al. 2014. "Risking other people's money: Gambling, limited liability, and optimal incentives." *Working Paper.*

Dewatripont, Mathias, Ian Jewitt, and Jean Tirole. 1999. "The economics of career concerns, part II: Application to missions and accountability of government agencies." *The Review of Economic Studies*, 66(1): 199–217.

D'Haultfoeuille, Xavier, and Philippe Fevrier. 2020. "The provision of wage incentives: A structural estimation using contracts variation." *Quantitative Economics*, 11(1): 349–397.

Diamond, Peter. 1998. "Managerial Incentives: On the Near Linearity of Optimal Compensation." *Journal of Political Economy*, 106(5): 931–957.

Dranove, David, Daniel Kessler, Mark McClellan, and Mark Satterthwaite. 2003. "Is more information better? The effects of "report cards" on health care providers." *Journal of Political Economy*, 111(3): 555–588.

Dye, Ronald A. 1986. "Proprietary and Nonproprietary Disclosures." *The Journal of Business*, 59(2): 331–366.

Edmans, Alex, and Xavier Gabaix. 2011. "Tractability in incentive contracting." *The Review of Financial Studies*, 24(9): 2865–2894.

Edmans, Alex, Tom Gosling, and Dirk Jenter. 2021. "CEO compensation: Evidence from the field." *CEPR Discussion Paper No. DP16315.*

Edmans, Alex, Xavier Gabaix, and Augustin Landier. 2009. "A multiplicative model of optimal CEO incentives in market equilibrium." *The Review of Financial Studies*, 22(12): 4881–4917.

Edmans, Alex, Xavier Gabaix, and Dirk Jenter. 2017. "Executive compensation: A survey of theory and evidence." *The Handbook of the Economics of Corporate Governance*, 1: 383–539.

Edmans, Alex, Xavier Gabaix, Tomasz Sadzik, and Yuliy Sannikov. 2012. "Dynamic CEO Compensation." *The Journal of Finance*, 67(5): 1603–1647.

Ely, Jeffrey, George Georgiadis, and Luis Rayo. 2022. "Incentive Compatibility in Dynamic Information Design." *Working Paper.*

Ely, Jeffrey, George Georgiadis, Sina Moghadas Khorasani, and Luis Rayo. 2022. "Optimal Feedback in Contests." *Working Paper*.

Englmaier, Florian, and Achim Wambach. 2010. "Optimal incentive contracts under inequity aversion." *Games and Economic Behavior*, 69(2): 312–328.

Englmaier, Florian, and Stephen Leider. 2012. "Contractual and Organizational Structure with Reciprocal Agents." *American Economic Journal: Microeconomics*, 4(2): 146–83.

Englmaier, Florian, Stefan Grimm, David Schindler, and Simeon Schudy. 2018. "The Effect of Incentives in Non-Routine Analytical Teams Tasks-Evidence from a Field Experiment." *CESifo Working Paper Series*, , (6903).

Fehr, Ernst, and Lorenz Goette. 2007. "Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment." *American Economic Review*, 97(1): 298–317.

Fehr, Ernst, Holger Herz, and Tom Wilkening. 2013. "The lure of authority: Motivation and incentive effects of power." *American Economic Review*, 103(4): 1325–59.

Fehr, Ernst, Lorenz Goette, and Christian Zehnder. 2009. "A behavioral account of the labor market: The role of fairness concerns." *Annual Review of Economics*, 1(1): 355–384.

Fershtman, Chaim, and Shmuel Nitzan. 1991. "Dynamic voluntary provision of public goods." *European Economic Review*, 35(5): 1057–1067.

Fisman, Ray, and Tim Sullivan. 2015. *The Org: The Underlying Logic of the Office-Updated Edition.* Princeton University Press.

Foarta, Dana, and Takuo Sugaya. 2021. "The management of talent: Optimal contracting for selection and incentives." *The RAND Journal of Economics*, 52(1): 49–77.

Fong, Yuk-fai, and Jin Li. 2017. "Relational contracts, limited liability, and employment dynamics." *Journal of Economic Theory*, 169: 270–293.

Frey, Bruno S., and Jana Gallus. 2017. "Towards an Economics of Awards." *Journal of Economic Surveys*, 31(1): 190–200.

Fuchs, William. 2007. "Contracting with Repeated Moral Hazard and Private Evaluations." *American Economic Review*, 97(4): 1432–1448.

Fudenberg, Drew, George Georgiadis, and Luis Rayo. 2021. "Working to learn." *Journal of Economic Theory*, 197: 105347.

Gallus, Jana. 2017. "Fostering Public Good Contributions with Symbolic Awards: A Large-Scale Natural Field Experiment at Wikipedia." *Management Science*, 63(12): 3999–4015.

Garicano, Luis, and Luis Rayo. 2016. "Why organizations fail: Models and cases." *Journal of Economic Literature*, 54(1): 137–92.

Garrett, Daniel, George Georgiadis, Alex Smolin, and Balazs Szentes. 2021. "Optimal Technology Design."

Georgiadis, George. 2014. "Projects and Team Dynamics." *The Review of Economic Studies*, 82(1): 187–218.

Georgiadis, George, and Balazs Szentes. 2020. "Optimal Monitoring Design." *Econometrica*, 88(5): 2075–2107.

Georgiadis, George, and Michael Powell. 2022. "A/B Contracts." *American Economic Review*, 112(1): 267–303.

Georgiadis, George, Doron Ravid, and Balazs Szentes. 2022. "Flexible Moral Hazard Problems." *Working Paper.*

Gerhart, Barry. 2017. "Incentives and pay for performance in the workplace." In *Advances in Motivation Science*. Vol. 4, 91–140. Elsevier.

Gibbons, Robert. 1998. "Incentives in organizations." *Journal of Economic Perspectives*, 12(4): 115–132.

Gibbons, Robert, and Kevin J Murphy. 1992. "Optimal incentive contracts in the presence of career concerns: Theory and evidence." *Journal of Political Economy*, 100(3): 468–505.

Gibbs, Michael. 2016. "Past, present and future compensation research: Economist perspectives." *Compensation & Benefits Review*, 48(1-2): 3–16.

Gottlieb, Daniel, and Humberto Moreira. 2022. "Simple contracts with adverse selection and moral hazard." *Theoretical Economics*, 17(3): 1357–1401.

Gradstein, Mark, and Kai A Konrad. 1999. "Orchestrating rent seeking contests." *The Economic Journal*, 109(458): 536–545.

Green, Brett, and Curtis R Taylor. 2016. "Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects." *American Economic Review*, 106(12): 3660–99.

Green, Jerry R., and Nancy L. Stokey. 1983. "A Comparison of Tournaments and Contracts." *Journal of Political Economy*, 91(3): 349–364.

Grossman, Sanford J., and Oliver D. Hart. 1983. "Implicit Contracts Under Asymmetric Information." *The Quarterly Journal of Economics*, 98: 123–156.

Guesnerie, Roger, and Jean-Jacques Laffont. 1984. "A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm." *Journal of Public Economics*, 25(3): 329–369.

Guiteras, Raymond P., and B. Kelsey Jack. 2018. "Productivity in piece-rate labor markets: Evidence from rural Malawi." *Journal of Development Economics*, 131: 42–61.

Halac, Marina. 2012. "Relational contracts and the value of relationships." *American Economic Review*, 102(2): 750–79.

Halac, Marina, Elliot Lipnowski, and Daniel Rappoport. 2021. "Rank uncertainty in organizations." *American Economic Review*, 111(3): 757–86.

Halac, Marina, Navin Kartik, and Qingmin Liu. 2017. "Contests for Experimentation." *Journal of Political Economy*, 125(5): 1523–1569.

Harris, Christopher, and John Vickers. 1987. "Racing with Uncertainty." *The Review of Economic Studies*, 54(1): 1–21.

Hebert, Benjamin. 2017. "Moral Hazard and the Optimality of Debt." *The Review of Economic Studies*, 85(4): 2214–2252.

Hoffmann, Florian, Roman Inderst, and Marcus Opp. 2020. "Only Time Will Tell: A Theory of Deferred Compensation." *The Review of Economic Studies*, 88(3): 1253–1278.

Holmström, Bengt. 1979. "Moral Hazard and Observability." *The Bell Journal of Economics*, 10(1): 74–91.

Holmström, Bengt. 1982. "Moral Hazard in Teams." *The Bell Journal of Economics*, 13(2): 324–340.

Holmström, Bengt. 1999. "Managerial incentive problems: A dynamic perspective." *The Review of Economic Studies*, 66(1): 169–182.

Holmström, Bengt. 2017. "Pay for Performance and Beyond." *American Economic Review*, 107(7): 1753–77.

Holmström, Bengt, and Paul Milgrom. 1987. "Aggregation and Linearity in the Provision of Intertemporal Incentives." *Econometrica*, 55(2): 303–328.

Holmström, Bengt, and Paul Milgrom. 1991. "Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design." *Journal of Law, Economics & Organization*, 7: 24.

Hong, Fuhai, Tanjim Hossain, John A. List, and Migiwa Tanaka. 2018. "Testing the Theory of Multitasking: Evidence from a Natural Field Experiment in Chinese Factories." *International Economic Review*, 59(2): 511–536.

Hotz, V Joseph, and Robert A Miller. 1993. "Conditional choice probabilities and the estimation of dynamic models." *The Review of Economic Studies*, 60(3): 497–529.

Itoh, Hideshi. 1991. "Incentives to help in multi-agent situations." *Econometrica*, 611–636.

Jensen, Michael C. 2001. "Corporate Budgeting Is Broken-Let's Fix It." *Harvard Business Review*, 79(10): 94–101.

Jewitt, Ian. 1988. "Justifying the First-Order Approach to Principal-Agent Problems." *Econometrica*, 56(5): 1177–1190.

Jewitt, Ian, Ohad Kadan, and Jeroen M. Swinkels. 2008. "Moral hazard with bounded payments." *Journal of Economic Theory*, 143(1): 59–82.

Kamenica, Emir. 2012. "Behavioral economics and psychology of incentives." *Annual Review of Economics*, 4(1): 427–452.

Khalil, Fahad, and Jacques Lawarree. 1995. "Input versus output monitoring: who is the residual claimant?" *Journal of Economic Theory*, 66(1): 139–157.

Krueger, Alan B, and Lawrence H Summers. 1988. "Efficiency wages and the inter-industry wage structure." *Econometrica*, 259–293.

Larkin, Ian. 2014. "The Cost of High-Powered Incentives: Employee Gaming in Enterprise Software Sales." *Journal of Labor Economics*, 32(2): 199–227.

Lazear, Edward P. 1986. "Salaries and piece rates." *Journal of Business*, 405–431.

Lazear, Edward P. 2000. "Performance Pay and Productivity." *American Economic Review*, 90(5): 1346–1361.

Lazear, Edward P. 2018. "Compensation and incentives in the workplace." *Journal of Economic Perspectives*, 32(3): 195–214.

Lazear, Edward P., and Sherwin Rosen. 1981. "Rank-Order Tournaments as Optimum Labor Contracts." *Journal of Political Economy*, 89(5): 841–864.

Levin, Jonathan. 2002. "Multilateral contracting and the employment relationship." *Quarterly Journal of Economics*, 117(3): 1075–1103.

Levin, Jonathan. 2003. "Relational Incentive Contracts." *American Economic Review*, 93(3): 835–857.

Li, Anqi, and Ming Yang. 2020. "Optimal incentive contract with endogenous monitoring technology." *Theoretical Economics*, 15(3): 1135–1173.

Li, Jin, and Niko Matouschek. 2013. "Managing conflicts in relational contracts." *American Economic Review*, 103(6): 2328–51.

Li, Jin, Niko Matouschek, and Michael Powell. 2017. "Power dynamics in organizations." *American Economic Journal: Microeconomics*, 9(1): 217–41.

Lipnowski, Elliot, and Joao Ramos. 2020. "Repeated delegation." *Journal of Economic Theory*, 188: 105040.

Lo, Desmond, Mrinal Ghosh, and Francine Lafontaine. 2011. "The Incentive and Selection Roles of Sales Force Compensation Contracts." *Journal of Marketing Research*, 48(4): 781–798.

MacLeod, Bentley W. 2003. "Optimal Contracting with Subjective Evaluation." *American Economic Review*, 93(1): 216–240.

MacLeod, W Bentley, and James M Malcomson. 1989. "Implicit contracts, incentive compatibility, and involuntary unemployment." *Econometrica*, 447–480.

Makarov, Igor, and Guillaume Plantin. 2015. "Rewarding Trading Skills without Inducing Gambling." *The Journal of Finance*, 70(3): 925–962.

Malcomson, James M. 1984. "Work incentives, hierarchy, and internal labor markets." *Journal of Political Economy*, 92(3): 486–507.

Miao, Jianjun, and Alejandro Rivera. 2016. "Robust Contracts in Continuous Time." *Econometrica*, 84(4): 1405–1440.

Mirrlees, James A. 1976. "The Optimal Structure of Incentives and Authority within an Organization." *The Bell Journal of Economics*, 7(1): 105–131.

Misra, Sanjog, and Harikesh S. Nair. 2011. "A structural model of sales-force compensation dynamics: Estimation and field implementation." *Quantitative Marketing and Economics*, 9(3): 211–257.

Mookherjee, Dilip. 1984. "Optimal Incentive Schemes with Many Agents." *The Review of Economic Studies*, 51(3): 433–446.

Murphy, Kevin J. 1999. "Executive compensation." *Handbook of Labor Economics*, 3: 2485–2563.

Murphy, Kevin J. 2013. "Executive compensation: Where we are, and how we got there." In *Handbook of the Economics of Finance*. Vol. 2, 211–356. Elsevier.

Murphy, Kevin M, and Robert H Topel. 1990. "Efficiency wages reconsidered: Theory and evidence." In *Advances in the Theory and Measurement of Unemployment*. 204–240. Springer.

Myerson, Roger B. 2015. "Moral hazard in high office and the dynamics of aristocracy." *Econometrica*, 83(6): 2083–2126.

Oyer, Paul. 1998. "Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality*." *The Quarterly Journal of Economics*, 113(1): 149–185.

Oyer, Paul, Scott Schaefer, et al. 2011. "Personnel Economics: Hiring and Incentives." *Handbook of Labor Economics*, 4: 1769–1823.

Paarsch, Harry J, and Bruce Shearer. 2000. "Piece rates, fixed wages, and incentive effects: Statistical evidence from payroll records." *International Economic Review*, 41(1): 59–92.

Palomino, Frederic, and Andrea Prat. 2003. "Risk Taking and Optimal Contracts for Money Managers." *The RAND Journal of Economics*, 34(1): 113–137.

Prendergast, Canice. 1999. "The provision of incentives in firms." *Journal of Economic Literature*, 37(1): 7–63.

Prendergast, Canice. 2002. "The tenuous trade-off between risk and incentives." *Journal of Political Economy*, 110(5): 1071–1102.

Raff, Daniel MG, and Lawrence H Summers. 1987. "Did Henry Ford pay efficiency wages?" *Journal of Labor Economics*, 5(4, Part 2): S57–S86.

Rahmandad, Hazhir, Rebecca Henderson, and Nelson P Repenning. 2018. "Making the numbers? "Short termism" and the puzzle of only occasional disaster." *Management Science*, 64(3): 1328–1347.

Rayo, Luis. 2007. "Relational Incentives and Moral Hazard in Teams." *The Review of Economic Studies*, 74(3): 937–963.

Rogerson, William P. 1985*a*. "The first-order approach to principal-agent problems." *Econometrica*, 1357–1367.

Rogerson, William P. 1985*b*. "Repeated Moral Hazard." *Econometrica*, 53(1): 69–76.

Rosen, Sherwin. 1986. "Chapter 12 The theory of equalizing differences." In . Vol. 1 of *Handbook of Labor Economics*, 641–692. Elsevier.

Sannikov, Yuliy. 2008. "A Continuous- Time Version of the Principal: Agent Problem." *The Review of Economic Studies*, 75(3): 957–984.

Shapiro, Carl, and Joseph E Stiglitz. 1984. "Equilibrium unemployment as a worker discipline device." *The American Economic Review*, 74(3): 433–444.

Shearer, Bruce. 2004. "Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment." *The Review of Economic Studies*, 71(2): 513–534.

Shue, Kelly, and Richard R Townsend. 2017. "Growth through rigidity: An explanation for the rise in CEO pay." *Journal of Financial Economics*, 123(1): 1–21.

Spear, Stephen E., and Sanjay Srivastava. 1987. "On Repeated Moral Hazard with Discounting." *The Review of Economic Studies*, 54(4): 599–617.

Taylor, Curtis R. 1995. "Digging for Golden Carrots: An Analysis of Research Tournaments." *American Economic Review*, 85(4): 872–890.

Tzioumis, Konstantinos, and Matthew Gee. 2013. "Nonlinear incentives and mortgage officers' decisions." *Journal of Financial Economics*, 107(2): 436–453.

Walton, Daniel, and Gabriel Carroll. 2019. "When are Robust Contracts Linear?" *Working Paper.*

Winter, Eyal. 2004. "Incentives and Discrimination." *American Economic Review*, 94(3): 764–773.

Wong, Tak-Yuen. 2019. "Dynamic agency and endogenous risk-taking." *Management Science*, 65(9): 4032–4048.