Statement of Research

George Georgiadis

Kellogg School of Management, Northwestern University

October 2022

I am an applied microeconomic theorist with a focus on organizational economics and industrial organization. At a broad level, my research studies how incentives—predominately financial ones—affect the behaviors of individuals and organizations. Most of my research belongs to one of two agendas. The first studies how firms can design effective incentive schemes for their workforce. The second analyzes public good provision problems, and specifically explores ways to mitigate freeriding and other inefficiencies. I analyze these problems using primarily game theoretic tools, and in some cases, also simulations and empirical methods.

1 Contract Theory

My recent research contributes to the agency literature under moral hazard that was spearheaded in the 1970's by Mirrlees (1976), Holmström (1979), and others. In the canonical model, a principal offers a wage scheme to an agent that specifies the latter's pay as a function of output. If the agent accepts the contract (wage scheme), then they privately choose a costly action (effort), which determines the probability distribution of output. Finally, output and payoffs are realized, and the game ends. Both players are expected payoff maximizers, and importantly, the agent's action is not contractible. This model is very elegant, and sheds light on the key trade-offs. Perhaps the most celebrated takeaway is that optimal wage schemes are shaped by a trade-off between incentives and insurance, and reward the agent for those output realizations which are informative about his effort.

Much of my work contributes to the literature that seeks to relax various assumptions of the above framework. In particular, Georgiadis and Szentes (2020) endogenizes the performance measure. Barron, Georgiadis and Swinkels (2020), as well as Georgiadis, Ravid and Szentes (2022) enrich the agent's action space. Garrett et al. (2022) varies the timing of the contracting game. In Ely et al. (2022, WIP) we explore the role of feedback in incentive design. Georgiadis and Powell (2022) and Antic and Georgiadis (WIP) are interested in operationalizing contract theoretic model. That is, they ask and answer, in the context of two different frameworks, what data does a firm need and how it should use it to improve an existing incentive contract. Ely et. al. (2022a, 2022b) study the combined use of monetary and informational incentives in a dynamic moral hazard setting. Finally, Fudenberg et al. (2021) studies apprenticeships: how to optimally compensate an agent using not only cash, but also knowledge that raises their productivity.

1.1 Endogenous Performance Measures

Designing an incentive plan requires first to identify appropriate measures of performance, and then to tie these measures to outcomes for the agent; e.g., via performance pay. Indeed, firms devote substantial resources to identifying and designing effective performance measures. In ***Optimal Monitoring Design*** (*Econometrica*, 2020, with B. Szentes) we study a principal-agent model in which both the performance measure and the pay-for-performance relationship are endogenous.

In our framework, the principal commits to a strategy for sequentially acquiring costly signals that are informative of the agent's effort, and a wage scheme that specifies a payment to the agent as a function of the acquired signals. The agent then privately chooses how much effort to exert. Finally, the principal acquires signals according to the chosen strategy and payoffs are realized. The principal's objective is to minimize the cost of inducing a risk-averse and liquidity-constrained agent to exert a particular level of effort.[1]

To conceptualize this model, you can think of a periodic employee evaluation: The principal collects and reviews evaluation-relevant information, and remunerates the employee according to the acquired information, but this process takes time and effort. For concreteness, suppose that each piece of information is a Gaussian signal with mean equal to the employee's effort and some fixed variance, and it costs a fixed amount to acquire and review. A signal acquisition strategy is a *stopping time*: it specifies at every moment and as a function of the already acquired signals whether to stop, possibly probabilistically.

Attacking this problem head-on is challenging, because the space of signal acquisition strategies is vast. Instead, we establish a series of lemmas that allow us to reduce the problem to one that is equivalent yet tractable. For any given signal acquisition strategy, the problem boils down to a canonical principal-agent problem. Our first lemma shows that the cost-minimizing wage is a function of a Lagrange multiplier and the so-called (Fisher) *score*, which is a zero-mean random variable and can loosely be interpreted as a log-likelihood ratio.[2]

[1] The problem of finding a profit-maximizing contract is typically decomposed into two steps. As is common in the literature, we focus on the first, which entails finding for each effort, the contract that minimizes the principal's expected cost. The second step, which solves for the profit-maximizing effort is generally omitted.

[2] A score close to $+\infty$ implies that the agent almost certainly chose the target effort level, whereas a score close to $-\infty$ implies that he almost certainly did not.

Next, we show that instead of choosing a stopping time, equivalently, the principal might choose a probability distribution over scores at a cost that is proportional to the variance of that distribution, and a wage scheme as a function of the realized score. Dualizing the agent's incentive compatibility constraint, this problem reduces to a min-max program, in which the principal first maximizes a Lagrangian function with respect to the dual multiplier (to obtain the optimal wage scheme given an arbitrary distribution over scores), and then minimizes it with respect to the distribution over scores. Solving this problem is hard, but the corresponding max-min problem is *not*.

Our main result establishes that under a condition on the agent's utility function, the two problems are equivalent, and moreover, the optimal contract features a *binary* distribution over scores.[3] The upshot is that the optimal wage scheme is also binary: it pays a wage plus a fixed bonus if the acquired signals are sufficiently favorable. It can also be interpreted as an efficiency wage, whereby the agent is paid a wage and is dismissed if the acquired signals are not sufficiently favorable.

What signal acquisition strategy corresponds to the optimal distribution over scores? This strategy is characterized by two parallel upward-sloping lines, one with a positive intercept and one with a negative one. The principal continues to acquire information as long as the sum of the acquired signals lies in-between the two lines. Intuitively, the principal stops upon becoming sufficiently certain that the agent did or did *not* exert the target effort. This strategy is reminiscent of (s, S) policies, familiar from inventory management and other dynamic models of lumpy decision making.

Turning to empirical implications, our model predicts that if the agent is only moderately risk-averse, then the bonus is large and is paid only if the acquired information is overwhelmingly positive. As the agent becomes more risk-averse, the principal acquires more information (in expected terms), the bonus decreases, but it is paid with higher probability—all in line with insuring the agent more.

1.2 Flexible Actions

The optimal contract in the canonical model depends on the minutiae of the distributions that map each action to output. This is a consequence of the agent choosing a typically one-dimensional action that determines the entire distribution of output (Holmström and Milgrom, 1987). Below I describe three articles that explore the consequences of enriching the agent's action space.

[3] This condition is always met if the agent has CARA utility, or CRRA utility with coefficient of relative risk aversion greater than 1/2. If the coefficient is less than 1/2, then the first-best outcome can be approximated by a single-bonus contract. More generally, if a solution exists, then the optimal distribution over scores is either binary or trinary.

In ***Optimal Contracts with a Risk-Taking Agent*** (*Theoretical Economics*, 2020, with D. Barron and J. Swinkels) we study optimal contracts in settings where the agent can game their incentives by engaging in risk-taking. Excessive risk-taking has been linked to poorly designed incentives (Rajan, 2011), and has been documented in various settings, including portfolio managers (Chevalier and Ellison, 1997), executives (Shue and Townsend, 2017), and entrepreneurs (Vereshchagina and Hopenhayn, 2009).

In our principal-agent model, the agent privately chooses a costly effort that produces a non-contractible *intermediate output*. The agent privately observes this output and can manipulate it by costlessly adding mean-preserving noise to it, which determines the final, contractible output. (In the context of a portfolio manager, you can think of intermediate output as a portfolio's mid-year performance. The manager can add risk to the portfolio to determine end-of-year performance, which is the final output.)

The key twist of this model is that the agent can take on risk by adding noise to their output. In equilibrium, the principal optimally designs the contract to deter such risk-taking. Doing so adds a new constraint, which requires the agent's utility to be concave in output. We show that the problem of characterizing the optimal contract reduces to imposing this *no-gaming constraint* on a classic moral hazard problem.

If the agent is risk-neutral, then the no-gaming constraint implies that linear contracts are optimal. If we ignored the no-gaming constraint, the principal would offer a convex contract to concentrate high pay on high output. The no-gaming constraint therefore binds everywhere, so the principal offers the most convex contract that is also concave, which is a linear contract.

If the agent is risk-averse, then we provide a characterization of the unique profit-maximizing contract that implements any given effort. Analogous to the case with a risk- neutral agent, the optimal contract "irons" the agent's utility so that it is linear in output whenever the no-gaming constraint binds. However, the no-gaming constraint does not necessarily bind everywhere, leading to optimal incentives that are linear (in utility) over some outputs and concave (in utility) over others.

Our framework is elemental enough to easily incorporate into richer models of optimal incentives. We illustrate this point with three extensions: the first incorporates a cost of engaging in risk-taking, the second alters the timing of the game so that risk-taking occurs before the agent observes output, and the third reinterprets risk-taking as intertemporal gaming, as documented in Oyer (1998). In each of these extensions, we derive a natural analogue of the no-gaming constraint and characterize optimal incentives.

Earlier models in which the agent's action space is rich imposed specific functional form assumptions on the agent's cost of choosing an output distribution. In ***Flexible Moral***

Hazard Problems (Working paper, with D. Ravid and B. Szentes) we consider a principal-agent model, in which the principal first offers a wage scheme, and then the agent can choose *any* output distribution with support in a given compact set. We assume only that the agent's cost function is *smooth* and increasing in first-order stochastic dominance. Then output is drawn according to the chosen distribution, and the agent is paid according to the wage scheme.

Our first result characterizes, for each output distribution, the lowest-cost wage scheme that implements it. The key takeaway is that the cost-minimizing wage scheme reflects the marginal cost of choosing a particular output distribution but *not* the information content of the output, as is the case in one-dimensional effort models (Holmström, 2017). This wage scheme is constructed so that the target distribution satisfies a generalized first-order condition: the agent's marginal cost of choosing a nearby distribution is approximately his marginal benefit from doing so.

Our second result show that the wage scheme is always increasing in output as long as the agent's cost of choosing a distribution is monotone in first-order stochastic dominance.[4] To see why this result is true, suppose that the wage is larger at a small output level than at other higher outputs. Then the agent could profitably modify that distribution by moving probability mass from those higher outputs to the low output. This modification increases the expected wage, and moreover, because the modified distribution is first-order stochastically dominated by the original one, it is less costly.

Finally, we consider the principal's problem of finding the profit-maximizing distribution and the corresponding optimal contract. Under a smoothness assumption (which is necessary for tractability), we derive properties of the principal-optimal distribution, and provide sufficient conditions such that it has one or two outputs in its support.

When analyzing agency problems, the determinants of agency frictions are typically taken as given. In hidden-information models, for example, the distribution of types is typically treated as exogenous, and in hidden-action models, the cost of each action is usually part of the model description. However, if an agent's payoff depends on agency frictions, then he is likely to pursue to shape them in a way that enhances their payoff; for example, a worker might make human capital investments that affect their subsequent productivity or influence which tasks they work on. In ***Optimal Technology Design*** (Conditionally accepted at the *Journal of Economic Theory*, with D. Garrett, A. Smolin, and B. Szentes), we reconsider the standard limited-liability moral hazard problem and examine how an agent might maximize

[4] In one-dimensional models, because larger outputs are not necessarily more informative about the agent's effort than smaller ones, optimal wage schemes are monotone in output only under strong conditions on actions available to the agent, such as the monotone likelihood ratio property.

rents by optimally designing the production technology.

In our model, before interacting with the principal, the agent chooses a *production technology*, which specifies their cost of *producing* each output distribution with support on a compact set. After observing the agent's production technology, the principal offers a wage scheme, which is a mapping from output realizations to nonnegative payments. Finally, the agent privately chooses an output distribution at a cost determined by their first-stage choice.

Our model is one of a hold-up problem. The agent could choose a production technology that makes it possible to costlessly produce the largest output with certainty. While doing so would be socially efficient, the agent would then be unable to extract any rents. Instead, the agent designs the production technology so that generating large outputs is artificially costly. We show that there is an optimal production technology involving only binary distributions (i.e., the cost of any other distribution is prohibitively high), and we characterize the equilibrium technology defined on the binary distributions. In equilibrium, the principal offers a bonus which induces the agent to complete the project with probability one.

The assumption that the agent can choose *any* production technology—even ones that generate large outputs with little to no cost—serves to deliver sharp results and highlight the key trade-off. To shed light on the case in which not every production technology is feasible, we characterize the payoff combinations that can arise in principal-agent models with limited liability for some given production technology. The equilibrium payoff profile when the agent instead chooses the technology corresponds to the point in this set where the agent's payoff is maximized.

1.4 Prescriptive Contract Theory

Agency models under moral hazard have generated invaluable insights into fundament trade-offs. However, their use as prescriptive theories has been limited, as the optimal contracts they prescribe in a given environment often depend in complicated and subtle ways on unobservable characteristics of that environment.

With ***A/B Contracts*** (*American Economic Review*, 2022, with M. Powell) we aim to improve the practical applicability of the classic theory of incentive contracts under moral hazard. In particular, we establish modeling assumptions and a computationally tractable estimation method that enable an experimental test between two incentive contracts to be sufficient for determining the optimal contract and predicting performance of any counterfactual contract. We then test our approach using data from DellaVigna and Pope's 2018 experimental study of six different incentive treatments.

To introduce our main ideas, let us consider an example. Suppose you hire teenagers each summer to sell kitchen knife sets door to door, and you pay them a simple piece rate for doing so. You have access to sales data for your workforce, and you are interested in knowing whether, and how, you should change the piece rate. Suppose your gross profit margin for selling a knife set is m, the piece rate is α, and your worker's average sales are a. Your expected profits are therefore $\Pi = (m - \alpha)\, a$. If you were to marginally increase your piece rate, the effect on your profits would be

$$\frac{d\Pi}{d\alpha} = (m - \alpha)\,\frac{da}{d\alpha} - a.$$

You know your gross profit margin, the current piece rate, and the current average sales, but you don't know your workers' *behavioral response*, $da/d\alpha$, to an increase in the piece rate. Given observational data alone, figuring out this behavioral response requires knowing your workers' preferences with respect to money and effort, as well as the effect of effort on the distribution of their sales. These are questions you likely do not know the answer to, but importantly, they are questions you do not need to know the answer to if you are willing to run an experiment. Suppose you decide to run an A/B test on your workforce. You randomly divide it into a treatment and a control group, you increase the piece rate by a small amount in the treatment group, and you have access to the data on the distribution of output for both the status quo contract and the test contract. You can use this data to estimate $da/d\alpha$, and you can use the above expression to determine whether you should marginally increase or decrease your piece rate.

This example sidesteps two important issues. First, it restricts attention to linear contracts. Second, it asks a local question—how best to marginally improve upon the status quo contract—and for practical applications, we are interested in non-local adjustments. We address each of these issues in turn.

Our principal–agent framework is as in Holmström (1979), but with a non-standard informational assumption. We assume that the principal knows neither the agent's effort costs, nor the transition probability function that maps each effort into output. Instead, she has output data corresponding to two contracts—a status quo and a test contract.

Our main conceptual contribution shows that given knowledge of the agent's preferences for money, a *single* A/B test of incentive contracts suffices to estimate the agent's behavioral response to *any* marginal adjustment to the status quo contract. We then show that the problem of how best to *locally* adjust a status quo contract is equivalent to figuring out the direction of steepest ascent in the principal's objective, which can be determined by solving a convex program.

The second important issue that the above example sidestepped was the question of how to predict the effects of *non-local* adjustments to the status quo contract. We establish assumptions under which a single A/B test provides all the information needed to predict how the principal's profits will respond to *any* adjustment to the status quo contract. In doing so, we provide a procedure for using this information to optimally adjust the status quo contract.

We then explore the quantitative implications of our results using data from DellaVigna and Pope's (2018) large-scale experimental study of how a variety of incentive schemes motivate subjects in a real-effort task. Our first exercise asks whether subjects' average performance varies in the way our model predicts with our measure of the subjects' marginal incentives. We take the data from two treatments supposing that in one of the treatments, the subjects were on the status quo contract, and in the other, they were on the test contract. For each such pair, we predict the mean performance in each of the remaining four treatments and compare it to the actual average performance. The mean absolute percentage error across all such pairs is approximately 2%. Moreover, our predictions for a given treatment are similar no matter which pairs we use to make our predictions. Taken together, the correlation between our predictions and actual performance is 0.94.

Our second empirical exercise assesses the performance of the contract generated by our procedure. We use the data from seven treatments to fit the parameters of the production environment using nonlinear least squares estimation. Given those parameters, we compute, as a benchmark, the optimal contract and the principal's corresponding expected profit. Then, we take data from each pair of treatments and use our procedure to construct the optimally adjusted contract. Averaging across all A/B tests, the principal can attain just over two-thirds of the profit gains that she could attain if she knew the entire production environment and put the optimal contract in place.

To characterize an optimal contract given information contained in an A/B test of incentive contracts, we imposed arguably strong assumptions about agent's cost function and the transition probability function that maps each effort into output. In ***Robust Contracts: A Revealed Preference Approach*** (Work in progress, with N. Antic) we revisit this question under minimal assumptions about the production environment. In this model, the agent's "action" is a probability distribution over output. The principal is oblivious to the full set of actions available to the agent and their costs, but observes the agent-optimal action in response to each of K "known" contracts, and aims to choose a contract with the maximum possible profit guarantee. This is *as if* an adversarial third party picks the agent's action set and the cost of each action to minimize the principal's profit subject to a set of revealed preference constraints, and the principal maximizes this worst-case profit.[5]

[5]Conceptually, you can think that the principal has a large number of output of observations under each of the K contracts, and is unwilling to make assumptions about the agent's action set or the cost of each

Our main result shows that the optimal contract is either the most profitable of the known contracts, or a mixture of *one* of the known contracts and a linear contract; specifically the one that makes the agent full residual claimant. If moreover there are two known contracts both of which are linear, then a mixture contract is never optimal. Applying our methodology to the data from DellaVigna and Pope's 2018 real-effort experiment, we find that in virtually every case, a mixture contract does not provide a larger profit guarantee than the most profitable of the known contracts. This finding suggests that unless one is willing to make additional assumptions about the production environment, it is typically best to stick to one of the known contracts. It also provides a rationale for why firms are reluctant to experiment with different incentive schemes when they have one that works reasonably well.

1.5 Paying with Information (and Money)

Another strand of my research studies the combined use of monetary and informational incentives to motivate one or more workers. Conceptually, a firm can shape its workers' informational incentives by curating the information their receive about their own or others' performance.

In ***Optimal Feedback in Contests*** (Forthcoming at the *Review of Economic Studies*, with J. Ely, S. Khorasani, and L. Rayo) we consider a game between a principal and N agents. Each agent continuously chooses to work or shirk, and working generates a "Poisson success"—a signal that arrives at a random date and is observed by the principal.[6] The principal has a fixed prize to award and designs a mechanism—in particular, a contest—to maximize total effort. In particular, she chooses rules dictating *when* the mechanism will end, how the prize is to be allocated, and a feedback policy that specifies a message to be sent to each agent as a function of her past observations.

This framework can be applied, for instance, to a professional partnership seeking to promote one of their associates to partner. Here a "success" represents an associate exceeding an exogenous threshold for promotion and an associate's effort is presumably valuable even after they, or any of their peers, have cleared that bar for promotion.[7]

The key challenge when searching for an optimal contest is the vast range of potential designs. We attack this problem by providing a sufficient condition for a contest to maximize effort—namely, that it maximizes the probability that the prize is awarded while giving zero

action beyond those observations. Absent a prior over the various possibilities, we posit that this is a natural objective.

[6] It turns out to be immaterial whether agents observe their own signal, but it is important that they don't observe others' successes. In the baseline model, each agent can generate at most one such signal.

[7] In other applications, the designer's objective might be to maximize, instead of effort, the number of Poisson successes. As it turns out, the two objectives are equivalent.

rents to the contestants—and then displaying contests that meet these criteria. One such contest, which we term "cyclical egalitarian", features a cyclical structure whereby the contest is terminated at the end of each fixed-length cycle if at least one agent has succeeded by then, and it is otherwise reset. The prize is shared equally among all successful agents irrespective of when they succeeded, and the feedback policy keeps agents fully apprised of their own success, but only periodically informs them about their rivals' successes—at the end of each cycle—so as to not discourage further effort.

Only a much smaller set of contests, however, is capable of converting all prize money into effort in the shortest expected duration—a property that would be valuable to the designer if for instance running the contest entailed a flow cost. These contests, which we term "2^{nd} chance", all have in common that the contest continues until some pre-specified number of successes arrive and, once that occurs, the contest enters a countdown phase where contestants are given a final (potentially random) deadline to succeed (a "2^{nd} chance"), with the contest ending before that deadline if one more contestant succeeds. These contests minimize duration because they guarantee that the number of agents working at a given time is as similar as possible across different histories—which in turn prevents inefficient scenarios where the contest continues with only a small number of agents are still working.

We also consider several extensions that suggest a degree of robustness to our results: the possibility of multiple successes, heterogeneity of success rates across agents, and an increasing hazard rate that captures a notion of progress or knowledge accumulation over time. In all these cases, a 2^{nd} chance contest attains maximum effort; and in several cases, it does so in the minimum possible duration. We also relax the designer's commitment power and show that an egalitarian prize structure, which is implicitly present in all our other designs, remains optimal.

This problem was made tractable by characterizing an upper bound for the principal's payoff, and then characterizing contests that achieve it. To do so, we assumed that the number of agents is sufficiently large. In ***Incentive Compatibility in Dynamic Information Design*** (Work in progress, with J. Ely and L. Rayo) we consider a similar dynamic agency model between a principal and *one* agent. The principal commits to a terminal date, a payment schedule as a function of the time that a random binary signal of effort is (privately) observed, and a feedback policy that specifies, for each point in time, messages as a function of past messages and the arrival date of the random signal, if any. The principal's objective is to maximize total effort net of monetary payments.

Because extracting all rents is not possible in this setting, we must optimize directly with respect to the choice variables. By a revelation-principle-like argument, it suffices to consider direct policies in which the principal recommends to the agent whether to work at every moment. We identify a key necessary condition for incentive compatibility and use it to pin down, for any recommendation policy, the cost-minimizing payment schedule. Conceptually, paying the agent with information (i.e., by being transparent about whether the random signal has arrived) is a substitute for monetary payments, and the optimal contract must balance the two means of motivation. Our main result shows that under the optimal policy, the agent works continuously up to come time-cutoff. If the random signal has arrived by that time, the agent is paid a fixed lump-sum and stops. Otherwise, the agent continues to work right until the signal arrives, and is paid a smaller lump-sum if the signal arrives prior to the terminal date.[8]

1.6 Paying with Knowledge (and Money)

Careers in a wide range of industries, such as medicine, academia, professional services, and culinary arts, frequently begin with a lengthy "apprenticeship" stage where novices gain knowledge from their masters while working hard and receiving relatively low wages. ***Working to Learn*** (*Journal of Economic Theory*, 2021, with D. Fudenberg and L. Rayo) studies such a work-for-knowledge exchange. In our model, a principal offers to an agent an "apprenticeship" consisting of time paths of knowledge transfers, wages, and effort, subject to the constraint that the agent can walk away at any time, and subject to a learning constraint that bounds how quickly they can learn.

Our analysis solves for the whole family of Pareto-optimal contracts as parameterized by the agent's outside option. Every such contract has two phases. In the first one, the agent learns as fast as their learning-by-doing constraint allows while earning rents in the sense that they are more than compensated for the economic cost of working for the principal. Then in the second phase, the principal only allows the agent to learn as quickly as is consistent with the agent being willing to remain in the apprenticeship; here the principal keeps all rents.[9]

Throughout the apprenticeship, effort is distorted above the static first-best. This serves two purposes: first, higher effort allows the agent to learn faster, and second, increased effort transfers rents to the principal. To smooth the agent's consumption, wages are constant in the first phase. During phase 2, the principal offers an increasing wage path (i.e., backloads wages) in order to relax the binding participation constraint.

Our model has novel implications for (optimal) regulation, of which apprenticeships are a frequent target. In particular, a planner who is able to set upper bounds on effort and lower

[8]This recommendation policy can be implemented by staying silent until the cutoff, and thereafter, informing the agent as soon as the signal arrives.

[9]Perhaps surprisingly, the second phase never disappears completely, even in the agent-optimal contract. This is because when the agent's outside option is high, phase 1 pays them more than they produce, thus placing them in "debt." Phase 2 then allows the principal to gradually collect on this debt.

bounds on wages is able to implement any contract on the Pareto frontier, without needing to regulate the path of knowledge transfers (which is presumably much harder to do).

2 Provision of Public Goods

A second research agenda studies public good provision problems; see Admati and Perry (1991) and Marx and Matthews (2000) for early contributions. In the workhorse model, time is continuous and at every moment, each of N agents makes a costly contribution (effort) to a joint project. The project progresses—possibly stochastically—at a rate that depends on the sum of the agents' instantaneous efforts. As soon as the *state* of the project hits a specific threshold, the project is completed and generates a lump-sum payoff for each agent. Thus, this is a game with positive externalities, and will be prone to the freerider problem (Holmström, 1982). For example, you might think of individuals collaborating to develop a new product, or different nations contributing towards a *discrete* public good such as the International Space Station. Agents are impatient, and choose their efforts as a function of the current state of the project to maximize their payoff; that is, we focus on Markov Perfect equilibria (hereafter MPE).

This agenda comprises five articles, which use variations of this model to explore how different incentive structures and different allocations of decision rights can improve equilibrium outcomes.[10]

Projects and Team Dynamics (*Review of Economic Studies*, 2015) proposes the aforementioned model and comprises two parts. The first focuses on the agents' incentives and establishes comparative statics. The second part introduces a profit-maximizing principal who chooses the number of agents and an incentive contract for each agent.

This game has a unique MPE, wherein each agent's effort is strictly lower that the first-best level, and increases as the project progresses. This is because agents are impatient and are rewarded only upon completion. This implies that efforts are strategic complements across time: by frontloading his effort, an agent brings the project closer to completion, thus inducing others to raise their future efforts.[11]

A key comparative static shows that given a fixed total reward, larger teams work harder than smaller ones—both individually and on aggregate—if, but only if, the project is sufficiently far from completion. Increasing the team size has two effects. First, agents have stronger incentives to frontload their effort since they can influence the future efforts of more agents. Second, the (positive) externality, and hence the incentives to freeride, becomes

[10] A lecture note covering this literature is available here.

[11] Thus equilibrium effort paths are shifted downwards and flatter (i.e., tilted clockwise) relative to the first-best effort path.

stronger. During the early stages of the project, the frontloading motive is strong since a lot of progress remains, whereas the freeriding motive is muted as agents exert relatively little effort (and hence their marginal cost of effort is low). Thus the first effect dominates the latter and larger teams work harder than smaller ones. The opposite holds when the project is close to completion.

Turning to the principal's problem, if agents must be treated symmetrically, then it is optimal to backload incentives so that each agent is rewarded only upon completion. In this case, the problem reduces to choosing the number of agents and a budget for rewarding them. The restriction to symmetric contracts, however, is not without loss of generality. With two agents, the principal may prefer to have both working in the beginning, and *retire* one agent once the project gets sufficiently close to completion.[12]

A key decision in any project involves its *scope*: Which features to include in a new product before it is marketed? What results to pursue in a research project? In ***Project Design with Limited Commitment and Teams*** (*RAND Journal of Economics*, 2014, with S. Lippman and C. Tang) we embed the dynamic public good provision game described above in an agency model, where the principal chooses the *size* of the project. A bigger project requires more effort to complete but generates a larger payoff, to be split according to pre-specified proportions among the principal and the agents. A central feature of the model is that the principal is unable to commit to the size of the project until it is sufficiently close to completion. In developing an innovative new product for example, it may be difficult to specify its final characteristics until sufficient progress is made. The inability to commit can also be due to an asymmetry in the bargaining power between the principal and the agents. Formally, when the current state of the project is q, she can commit to any project size $Q \leq q + y$ where y captures her *commitment power.*

Our main result shows that the principal has incentives to extend the project as it progresses, for example by introducing additional requirements. Intuitively, the principal trades off the marginal benefit of a bigger project and the marginal cost associated with a longer wait until the bigger project is completed, but not the effort cost needed to do so. So as progress is made and past efforts are sunk, there is a motive to extend the project.

With sufficient commitment power, the principal commits to the optimal project size at time zero. Otherwise, they can commit to a smaller than ideal project at time zero or wait. However, once progress has been made, the optimal project size is larger than it was originally, and so faces the same dilemma. It turns out that the principal always prefers to

[12]The intuition stems from the team-size comparative static discussed above. One way to implement this scheme is by rewarding one agent upon reaching an intermediate milestone at which point he stops, and rewarding the other agent (only) upon completion.

wait, and eventually chooses a bigger project the smaller is their commitment power.

The agents of course anticipate this and decrease their effort to the principal's detriment. With too little commitment power, the principal is actually better off delegating the decision rights over the project size to the agents. In that case, the agents will choose a smaller project than is optimal for the principal but their preferences are time-consistent. Intuitively this is because unlike the principal, they do internalize the effort cost needed to complete the project.

Freeriding is a key impediment to collaboration. Since some of the benefits of my collaborative efforts accrue to you, I have a natural tendency to underprovide effort and so do you. In a dynamic setting such as the one described above, there is a second, less well understood inefficiency: collaborative efforts are frontloaded relative to the social optimum. I provide too much effort early on to bring the completion date forward, and thereby motivate you to provide more effort in the future. In ***Achieving Efficiency in Dynamic Contribution Games*** (*American Economic Journal: Microeconomics*, 2016, with J. Cvitanic) we propose a mechanism that induces first-best efforts as the outcome of an MPE in a dynamic contribution game.

The mechanism specifies for each agent flow payments that are a function of the progress made to date, and a reward that is disbursed upon completion. Intuitively, this mechanism must neutralize the two inefficiencies described above. To eradicate freeriding, each agent is effectively made the residual claimant. To neutralize frontloading, the flow payments increase with progress at a rate such that each agent's benefit from frontloading effort is exactly offset by the cost associated with having to make larger flow payments in the future. Assuming that flow payments are placed in a savings account that accumulates interest, we show that there exists a budget balanced mechanism that achieves efficiency—*ex post* if the project progresses deterministically, and *ex ante* if it progresses stochastically.[13]

We also adapt our mechanism to a dynamic common-resource extraction problem such as the one studied by Levhari and Mirman (1980), and a strategic experimentation problem as in Keller, Rady and Cripps (2005). For the resource extraction problem, the efficiency-restoring mechanism specifies that each agent receives a flow subsidy that decreases as the resource becomes more scarce to neutralize his incentives to overharvest it. To ensure budget balance, each agent pays a combination of an entry fee in exchange for access to the resource and a penalty that becomes due as soon as it is depleted. For the experimentation problem, the mechanism specifies that each agent pays an entry fee and receives a subsidy that depends

[13]Some features of incentive structures in startups resemble this mechanism. In particular, the flow payments in our model can be interpreted as the investments that entrepreneurs make until they raise capital, plus the difference between their income and the (higher) market salary rate. Finally, equity ownership makes them residual claimants.

on the (public) belief about the risky project being lucrative.

A ubiquitous feature of projects is a *deadline*—a date by which the project must be concluded (or abandoned). The article ***Deadlines and Infrequent Monitoring in the Dynamic Provision of Public Goods*** (*Journal of Public Economics*, 2017) is one of few in this literature to incorporate a deadline. The main challenge is that there are two state variables to keep track of: the progress made to date on the project, and the time remaining until the deadline. I overcome this challenge by exploiting that *on* the equilibrium path (though *not* off it), there is a one-to-one mapping between the two state variables, so the problem can be analyzed using standard optimal control techniques. The game has a unique and symmetric MPE that can be characterized analytically. When the deadline is sufficiently long, equilibrium efforts are inefficient (due to freeriding) and frontloaded. Shortening the deadline forces the agents to work harder, but efforts continue to be frontloaded.

The source of frontloading is that agents observe the state of the project continuously, so they have a motive to influence peers' future efforts by modifying their present effort. In team projects however, progress is often tallied only periodically; e.g., during group meetings. So I assume that agents observe the state of the project only at discrete dates, and in-between, they choose their efforts according to their beliefs about others' strategies. If agents were to observe progress only at the deadline, then frontloading would be eliminated, and with an appropriate choice of deadline, equilibrium efforts would be efficient.

A problematic feature of deadlines, however, is that they are not renegotiation proof: if the project is not completed by the deadline, agents will have an incentive to defer it *ex post*, undermining its intended purpose. I show that given an exogenous deadline, agents can maximize their ex ante payoffs by monitoring progress at some date prior to the deadline, and then not monitoring it again until the deadline. If that date is chosen appropriately, it will be in each agent's interest to complete the project by then, and so it acts as a self-enforcing deadline.

In many settings, agents with different preferences must collectively decide the scope of a project. For example, entrepreneurs collaborating on a joint business venture must choose whether to seek a blockbuster product or one with a quicker, if smaller, payoff. In ***Collective Choice in Dynamic Public Good Provision*** (*American Economic Journal: Microeconomics*, 2019, with R. Bowen and N. Lambert) we investigate the effect of the collective choice institution (*i.e.*, dictatorship and unanimity in two-member teams) on the equilibrium size of the project and on total welfare. We use the workhorse model described above with two agents who differ in their marginal cost of effort and their stake in the project.[14]

[14] To simplify the exposition, here I focus on the case in which both agents receive the same reward upon

We first characterize the project size that maximizes each agent's payoff. In particular, we show that the efficient agent prefers a smaller project than the inefficient agent, and moreover, his ideal project size *decreases* with progress. This is because he increases his effort at a faster rate than the inefficient agent, so his *share* of the remaining project cost *increases* as the project gets closer to completion. In contrast, the inefficient agent's ideal size *increases* as the project progresses. The agents' preferences over the project size are thus time-inconsistent and divergent.

Next, we analyze the equilibrium project size under dictatorship and unanimity. For brevity, I only discuss the case in which the parties cannot commit to an *ex ante* decision to implement a particular project size. If the efficient agent is dictator, then the ideal project size is implemented in the unique MPE. If instead the inefficient agent is dictator, then there is a continuum of equilibrium project sizes, all of which are smaller than the inefficient agent's ideal but larger than the efficient agent's. Last, because the inefficient agent prefers a larger project scope than the efficient agent, the set of equilibria under unanimity are the same as when the inefficient agent is dictator.

From a welfare perspective, it may be desirable to allocate decision rights to the inefficient agent, because the efficient agent obtains *real* control since he carries out the majority of the progress in equilibrium. This provides a rationale for unanimity as the collective choice institution in various settings, and it resonates with Galbraith (1952), who argues that when one party is strong and the other weak, it is preferable to give formal authority to the latter.

3 Other Research

In this section I briefly discuss other work beyond the above agendas.

In ***The Absence of Attrition in a War of Attrition under Complete Information*** (*Games & Economic Behavior*, 2022, with Y. Kim and D. Kwon) we study a canonical model of war of attrition under complete information. We show that if the players' payoffs whilst in "war" vary stochastically *and* their exit payoffs are heterogeneous, then the game admits Markov Perfect equilibria in pure strategies only. This is true irrespective of the degree of randomness and heterogeneity, thus highlighting the fragility of mixed-strategy equilibria to a natural perturbation of the canonical model.

Optimal Reservation Policies and Market Segmentation (*International Journal of Production Economics*, 2014, with C. Tang) studies optimal reservation policies for a capacity-constrained firm facing customers who are privately informed about their valuation

completion. I will refer to the agent with the smaller marginal cost as "efficient" and to the other agent as "inefficient".

for the service on offer and their show-up probability, such as an airline or an airport. The main takeaway is that when the firm's capacity is small relative to demand, it optimally screens customers mostly based on their show-up probability by requiring non-refundable deposits. As its capacity grows however, it turns to screening customers predominantly based on their valuation by overbooking in lieu of demanding a deposit.

In ***The Retail Planning Problem Under Demand Uncertainty*** (*Production and Operations Management*, 2013, with K. Rajaram) we consider a supply chain planning problem where a firm selects suppliers, and makes decisions on production, distribution and inventory at the retail locations for each of their (private label) products. Given that this problem is strongly NP-hard, we develop heuristics to obtain feasible solutions. Computational analysis indicates that our best heuristic yields feasible solutions that are close to optimal with an average suboptimality gap of 3.4%.

References

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

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

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

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. 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.

Keller, Godfrey, Sven Rady, and Martin Cripps. 2005. "Strategic experimentation with exponential bandits." *Econometrica*, 73(1): 39–68.

Levhari, David, and Leonard J Mirman. 1980. "The great fish war: an example using a dynamic Cournot-Nash solution." *The Bell Journal of Economics*, 322–334.

Marx, Leslie M., and Steven A. Matthews. 2000. "Dynamic Voluntary Contribution to a Public Project." *The Review of Economic Studies*, 67(2): 327–358.

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

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

Rajan, Raghuram G. 2011. "Fault lines." In *Fault Lines.* princeton University press.

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.

Vereshchagina, Galina, and Hugo A Hopenhayn. 2009. "Risk taking by entrepreneurs." *American Economic Review*, 99(5): 1808–30.