Optimal Scheduling under Adverse Selection & Hidden Actions[*]

Murali Agastya and Oleksii Birulin
School of Economics
University of Sydney
Sydney, NSW 2006
AUSTRALIA

Email: murali.agastya@sydney.edu.au
Email: oleksii.birulin@sydney.edu.au

April 4, 2021

Abstract

A Principal owns a project consisting of several tasks. Tasks differ, both in their innate success probabilities and their incremental benefits. Moreover, only specialists can perform these tasks. Subject to moral hazard and adverse selection, in what order should the Principal commission the tasks and when should she terminate the project? What investments into changing tasks' characteristics yield the highest marginal profit? These issues arise in diverse areas such as drug development, sequencing proxy wars or job assignment. We show that, despite informational constraints, a simple index – a task's effective marginal contribution – determines the optimal schedule/mechanism.

Keywords: Task Scheduling, Hidden Action, Adverse Selection, Effective Marginal Contributions, Gittins Index

[*]The paper has benefited from helpful conversations with Parimal Bag, JC Carbahal, Indranil Chakraborty, Nick de Roos, Moshe Haviv, George Mailath, Suraj Prasad, Vladimir Smirnov, Andrew Wait and seminar participants at AETW, NUS, UNSW, IIM Bangalore and the Econometric Society Meetings.

1 Introduction

Many projects can be viewed as collections of risky tasks that must be performed by specialists. A Principal determines the order in which the tasks will be attempted and whether to continue on to the next stage contingent on the tasks completed successfully. The tasks themselves differ in their characteristics, some are more likely to succeed if attempted, others offer greater benefit to the project *if* they succeed. The specialists undertaking their respective tasks largely care about their monetary transfers and their privately known costs of effort. They do not necessarily share the Principal's interest in the project overall.

A typical example is drug development. It requires computer modeling, multiple toxicology studies on animals, multiple clinical trials on humans, environmental risk assessments, etc. An important trend is the outsourcing of these R&D tasks to academic and private contract research organizations (CROs), largely driven by the use of increasingly sophisticated technologies. Visiongain (2016), a prominent business intelligence forecaster, estimates the drug discovery outsourcing market to be worth $68 billion by 2028, as compared to about 20 billion in 2016.[1] As in most R&D projects, there is some flexibility on the order in which these tasks are commissioned. (See Reyck and Leus (2008) for an elaborate case study.) Moreover, a drug development project will be terminated at any stage following the discovery of harmful side effects or the inefficacy of the approach/product. These features are also characteristic to software development.[2]

Scheduling proxy wars is arguably another notable area with similar issues. Unlike in conventional wars, a Principal, usually a foreign state, does not involve itself in active combat against an enemy. Instead it equips its proxies with arms and ammunition, provides technical knowhow, monetary incentives etc. The proxies are then expected to act.[3] Although proxy wars are relatively inexpensive compared to having "boots

[1]Similar or even higher estimates are presented in a range of online blogs on the Pharma industry, see also for example Buvalio (2018) or Lynch (2019).

[2]Gassmann et al. (2004) (on pg. 119) explain that "If a drug candidate fails during the development phase it is withdrawn entirely from further testing. Unlike in the automobile industry, drugs are not modular products where a faulty stick shift can be replaced without throwing the entire car design away. In pharmaceutical R&D, drug design cannot be changed." Similarly, (Ready, 2011, p. 28) points out that in software development practice "...it was all or nothing with our work. On numerous projects, we would only make money if we got everything lined up in perfect fashion. If we got 95 percent done ... we would still be in the same state as if we had just sat and watched *Matlock* reruns on the USA Network for the last ten months".

[3]Notable recent examples of such intervention include the US support of the Contras in Nicaragua and of the Mujaheddin in Afghanistan in the 1980s, the continuing Iran's support of several movements (from the Hezbollah to Houthi) especially since 2003, the involvement of the Russian backed "private military companies" first in Donbass, then in Syria and now in Libya. In fact, Mumford (2013) forcefully

on the ground", they are not cheap, not in the least due to the changing loyalties and commitments of the proxies. Hence when multiple conflicts are of concern to a Principal, the corresponding proxy wars have to be scheduled in an optimal manner while providing the right incentives to the proxies.[4]

Once the optimal schedule is settled, other questions of interest arise. Looking ahead, the Principal may consider making ex-ante investments that change the characteristics of the tasks. In the drug R&D example, the Principal may engage in costly lobbying to redefine the meaning of a success for a task. Alternatively, she could invest into "drug re-purposing" of the compounds resulting from successful tasks. This would increase the incremental benefit of a successful task, without affecting the chances of success. What is the return on a marginal investment in either of these parameters for a given task, and between tasks? For example, is it better to improve the success probability of a certain task or invest in another task with higher incremental benefit but lower success probability?

The common elements in all of the above examples are agency problems in a production process where the sequence of tasks is chosen by a Principal. It is difficult to monitor the CROs (or the proxies) actions, hence the problem involves moral hazard. It is also difficult to assess the CROs' (or the proxies') costs, hence the problem involves adverse selection. We are the first to construct a model that includes all these elements.[5] A task in our model also has the two features mentioned above – its success probability and a measure of its essenatiality (based on its marginal contribution to the project).

One of our main findings is that if a task is more essential than another but the latter is less risky, then in the optimal mechanism priority (on average) should be given to the former. Moreover, if the Principal could make ex-ante investments to improve the marginal contributions or the success probabilities of tasks, she should invest into the success probability of the most essential task.

Our results may also be applied to optimal job assignment. (See Gibbons and Wald-

argues that proxy warfare has all but replaced conventional wars since the Cold War, echoing a prominent sentiment on the obsolescence of major wars expressed earlier in Mueller (1989).

[4]The importance of sequencing the support to the proxies is evident. The debilitating effect of the Portuguese Colonial Wars (effectively three *simultaneous* proxy wars in Angola, Guinea-Bissau and Mozambique) on Portugal itself is well documented. Although an unprecedented $20 billion of US aid was provided to the Mujaheddin in Afghanistan between 1979 and 1989, the aid was increased dramatically in 1987, the same year the US aid to the Contras was severely limited after the Iran-Contra scandal. (See Lynch (2013), Prevost (1987) and (Riedel, 2014, pp. ix–xi, 21–22, 93, 98–99, 105. Chapter 4)).

[5]Optimal scheduling has been extensively studied in the operations research literature as a one person decision problem. Our model can be seen as introducing agency concerns to a particular version of the model explored in Reyck and Leus (2008). More closely related literature is surveyed in the latter part of the Introduction.

man (1999) for a survey.) That is, suppose that the tasks must be performed in a given sequence or even that, as in Kremer (1993), the output at each stage is an intermediate good to be used as the input to the next stage. The Principal can assign workers with varying abilities to these tasks. Kremer finds that workers with higher success probability must be assigned to later tasks in the general equilibrium of a complete information economy. Reinterpreting our results confirms this to be the case even with adverse selection and a profit maximizing Principal. (See Remark 2.)

More specifically, in our model, the Principal owns a project consisting of several tasks. Each of these can only be attempted by an agent with specific skills. When an agent is commissioned to undertake his task, he must take a costly action for the task to succeed with a (fixed) positive probability. The success of a task is verifiable. However, whether the agent has taken the costly action and the cost of that action to the agent is her private information. Each successful task brings an incremental benefit to the Principal, but the tasks also complement each other to deliver a sizeable boost to the project's value if *all* the tasks are completed successfully. Tasks, however, can be undertaken in any order and the choice of the later tasks can be contingent on the success or failure of the earlier tasks. Motivated by the drug development application above, we focus on the environments where it is inefficient to continue with the remaining scheduled tasks upon the failure of any task.

In Section 2.1 we introduce the notion of the *effective marginal contribution* (EMC) of a task. Much like the famed Gittins Index or the reserve price in Weitzman (1979), the EMC of a task is the index that depends on that task's parameters alone. As it turns out, virtually all of our analysis on optimal scheduling, regardless of whether an agent's cost of exerting effort is common-knowledge or his private information, involves commissioning the tasks in decreasing order of their EMCs for the *appropriately* chosen transfers.

If costs were common-knowledge, the expected transfer is equal to the agent's cost. The resulting optimal schedule has the following implications when the costs of all the tasks are equal. In choosing between a pair of tasks with equal marginal benefits (equal success probabilities), the task with the lower success probability (the higher marginal benefit) must be commissioned earlier. More generally, we say that there is a trade off between tasks i and j if the former has a higher success probability but the latter has a greater marginal benefit. Then task j precedes task i in the optimal schedule.

When an agent's cost of effort is his private information, the Principal must design a mechanism that provides incentives *both* to reveal the information required for scheduling *and* to take the costly actions. Mixed models of moral hazard and adverse selection are rarely tractable.[6] Proposition 1, however, provides a complete characterization of the optimal (Expected Profit Maximizing) mechanism here.

First, the agents simultaneously report their costs. Then, at each reported costs profile, the tasks are scheduled in decreasing order of their EMCs, computed *as though* the expected transfer to each agent is his *virtual cost*. The transfers required for incentive compatibility are, of course, more involved than simply covering the agents' virtual costs. In fact, we design the transfers such that the Perfect Bayesian Equilibrium that delivers the highest expected profit is also an ex-post equilibrium.[7]

Remarkably, each of our earlier statements on optimal scheduling with equal and commonly known costs has a direct counterpart in the asymmetric information setting when the agents' costs are i.i.d. When comparing a pair of tasks with equal marginal benefits (equal success probabilities), the task with the lower success probability (the higher marginal benefit) has a *higher ex-ante probability* of being commissioned earlier. Similarly, when there is a trade off between a pair of tasks, the task with a greater marginal benefit has a *higher ex-ante probability* of being commissioned earlier, see Proposition 2 in Section 4.3.

Proofs of all our results rely on the expression for the optimal expected profit in Proposition 1, which resembles the optimal revenue in typical auction-like IPV environments. The resemblance is however superficial. Due to moral hazard, the transfers depend not only on the reported but also on the true costs profile. The usual payoff equivalence then does not apply. Nevertheless, the Envelope Theorem methods can be used to derive the upper bound on optimal profit. The proof of Proposition 1 constructs a mechanism that delivers that upper bound in ex-post equilibrium. See Section 4.2 for more details.

Section 4.4 examines the incentives of the Principal to make ex-ante investments that improve the success probabilities or the marginal benefits of the tasks. We show that whenever a pair of tasks exhibits a trade off and the boost to the project's value that results from completing both tasks is large enough, the Principal benefits the most from improving the success probability of the task with the higher marginal benefit. It is worth emphasizing that all of our results under asymmetric information hold without any special assumptions on the prior distribution of the costs, beyond the usual i.i.d.

Throughout we assume that the agent makes a binary effort choice. It is a typical modeling choice in most of the cited literature even under pure moral hazard (and no

[6]Some notable exceptions are reviewed in Chapter 7 of Laffont and Martimort (2002).

[7]An agent prefers to reveal his cost and take the action when instructed regardless of the reports of the others as long as the other agents take actions on their turn.

adverse selection). Section 5 extends our results to the setting with continuous choice of effort under pure moral hazard. Furthermore, incorporating adverse selection may be possible along the lines of McAfee and McMillan (1986) and Laffont and Tirole (1987). However, this extension has proved to be significantly more involved and warrants a separate analysis.

Related Work

Optimal scheduling and stopping as a one person decision problem has been extensively studied in Operations Research.[8] In Economics, the literature on scheduling and stopping decisions under uncertainty but without agency considerations originates from Weitzman (1979). See Olszewski and Weber (2015) and Armstrong (2017) and the references therein for recent progress in such settings. In contrast to that literature, the tasks here are delegated to the agents, hence the Principal also has to optimize on the agency costs. There are three strands in the Economics literature with asymmetric information where an "allocation" is an ordering of agents as it is here.

Ordering of agents is also a key variable in the "queuing" problem, see the literature that follows Mitra (2001) and Suijs (1996). In contrast to the Operations Research literature referred to earlier, these models incorporate adverse selection. A Principal allows access to the resource to n agents, one at a time. An agent's type is an n-dimensional vector representing his cost of being served in each of the n possible positions. Queuing literature is mainly concerned with identifying cost structures where first-best mechanisms exist and with axiomatic characterization of particular mechanisms. The objective of the Principal is different from ours, she needs to order the agents in a queue efficiently, i.e., to minimize the aggregate waiting time. Moreover, the moral hazard constraints in our model induce ex-post individual rationality considerations typically absent in the queuing literature.

Optimal allocation of the slots on the list is a key issue in the search positions auctions literature. (See e.g., Varian (2007) or Edelman et al. (2007).) The advertised positions in the search results can be purchased at auctions run by the internet search engines. Each position has a fixed "click through rate" while each advertiser is privately

[8] The literature on scheduling in environments where a job once started necessarily succeeds, i.e., deterministic scheduling, is vast. It is mostly algorithmic and seeks to minimize the project's duration. (See Li and Ierapetritou (2008)) Scheduling with a view of maximizing a discounted profit is surveyed in Herroelen et al. (1997). The literature that incorporates uncertainty into a job's success is also vast. Herroelen and Leus (2005) surveys this literature, which again mostly focuses on minimizing the expected duration of the project. Reyck and Leus (2008) is among the more recent studies where the objective is to maximize the expected profit when the tasks are uncertain to succeed.

informed of its "value per click". The product of these two is the value of a position to an advertiser. Assigning advertisers to positions is akin to choosing the stages at which to commission different tasks. This similarity is, however, superficial, since even setting aside the moral hazard considerations, the payoff structures here and in the positions auctions problem are very different. Section 5 explains the differences between positions auctions and task scheduling.

Task scheduling with only moral hazard is explored in Winter (2006).[9] His focus is on scheduling of R&D activities performed *within* an organization. Success or failure of an individual task is observable but not verifiable. We address the complementary case of outsourcing the tasks to the *external* firms. Thus, in our setting the success of an individual task is verifiable.

Finally, there is a literature on procurement of a large project in pieces modelled as a contracting problem with one agent under moral hazard or as sequential auctions. In that literature, each piece (task) can be supplied by *all* the suppliers and the focus is on optimal use information about the suppliers generated over time. In our model, as each task requires job specific skills of a single agent, the primary sequencing is the information it reveals on the project and not of the agents.

The rest of this paper is organized as follows. Section 2 sets up the model and derives the optimal schedule when costs are common knowledge. Section 3 deals with the pure moral hazard model. Section 4 contains our main results for the model with moral hazard and adverse selection. Section 5 discusses the comparative statics and alternative applications of our methods. Section 6 concludes and provides directions for further research.

2 The Model

A Principal owns a project that consists of a set of tasks $N = \{1, \dots, n\}$. In the R&D interpretation of our model, these tasks can be performed in any order by specialized agents. Each *task* i has two parameters q_i and r_i. q_i is the success probability of i's task provided i takes a costly hidden action. Without that action the task fails. Success of

[9]In Mylovanov and Schmitz (2008) three identical tasks are to be accomplished within two stages. Unlike here or in Winter (2006) any agent can perform any task, but no agent can be allocated more than two tasks in one stage. Schmitz (2005) studies benefits of integration vs. separation under moral hazard with exogenously given order of the tasks. These are pure moral hazard models with complete contracting. Gilbert and Riordan (1995) and Da Rocha and Angeles de Frutos (1999) study the benefits of integration under adverse selection and perfect complementarity of the tasks. Severinov (2008) also allows the tasks to be substitutes. In the last three papers both tasks are always attempted and which one goes first is irrelevant.

task i adds $r_i \geq 0$ to the value of the project. There is also an additional benefit of $\alpha > 0$ if all the tasks are completed successfully (and the project itself is deemed a success).

In the alternative job-assignment interpretation of our model, the order of the tasks is fixed but any agent can be assigned to any one of the tasks. q_i and r_i are then interpreted as the productivity parameters of *agent* i. In particular, in the Kremer's O-ring theory of economic development, the production process consists of several stages and the output of stage k is an input for stage $k+1$. This is equivalent to setting $r_i = 0$ for all i, since the output of each stage is being used up in the next stage. Of course, the benefit of α is achieved if all stages succeed. For concreteness, we proceed with the R&D interpretation. However, the reader could easily change the emphasis to job assignment by swapping the "optimal ordering of the tasks" with the "optimal ordering of the agents".

We assume throughout that the parameters $\mathbf{q} = (q_1, \ldots, q_n)$, $\mathbf{r} = (r_1, \ldots, r_n)$ and α are common-knowledge.[10] In the event that a subset S of tasks are successful, the value of the project is

$$r(S) \quad = \quad \begin{cases} \sum_{i \in S} r_i & \text{if } S \neq N \\ \sum_{i \in N} r_i + \alpha & \text{otherwise.} \end{cases}$$

Throughout, in comparing a pair of tasks we call the task with the higher success probability the *easier* task. At the same time, we call the task with the higher expected benefit, $q_i r_i$, the *more essential* task.[11]

Successes and failures of the tasks are verifiable but the Principal incurs costs to induce agents to exert costly effort (to take the hidden action). The Principal can commission the tasks in any order and her later decisions can be contingent on the outcomes of earlier tasks. She may also terminate the project mid-way and realize the partial profits accumulated so far. A question of interest is whether the Principal gives priority to the easier or to the more essential tasks.

[10] Here the outcome of a task is uncertain. Instead, we may assume that r_i can be obtained for sure but after a duration t_i, and a discount factor $0 < \delta < 1$ captures the unit time cost. By setting $q_i = \delta^{t_i}$ such a model is equivalent to ours. Also, see Agastya et al. (2016) for a version of our model with complete information and two tasks only, but allowing for correlation between tasks' success probabilities.

[11] The rationale for describing a task with higher expected benefit as being more essential is as follows. A natural *a priori* measure of a task's essentiality is the marginal increment it brings to the project's expected value. That is, let Δ denote the project's expected value conditional on attempting *all* the tasks and Δ_{-i} the corresponding expected value conditional on attempting all the tasks but i. The marginal benefit of task i is then defined as $\Delta - \Delta_{-i}$. Accordingly, we could say that task i *more essential* than task j if $\Delta_{-j} > \Delta_{-i}$. With $n = 2$, $\Delta_{-i} = q_j r_j$, since only task j is being attempted. Then task i is relatively more essential than j if $q_i r_i > q_j r_j$. Importantly, this characterization of essentiality generalizes to $n > 2$. See Lemma 0 in the Appendix.

Admittedly assuming any sequence of tasks to be feasible is extreme. FDA, for instance, requires four main stages in drug development – basic research, pre-clinical, clinical and FDA review. Although these major phases cannot be swapped, there are several tasks within and between these phases that are interchangeable. Hence, more generally, we can consider an exogenously given "precedence network" that describes the set of feasible schedules. Our results should apply to any given sequence of interchangeable tasks within this precedence network.

Agents are risk neutral, and their reservation utilities are normalized to zero. If agent i receives a transfer τ_i his payoff is $\tau_i - c_i$ if he exerts effort and τ_i if he does not. The Principal observes neither the effort choice of i nor his cost of exerting effort, denoted by c_i. This c_i is i's private information, distributed on an interval $[\underline{c}, \bar{c}]$ according to a continuous and regular cumulative distribution function F.[12] An agent is motivated solely by his transfer net of his cost. He does not directly care about the project's benefits, as those benefits accrue to the Principal.

As the agents' actions are hidden, the Principal has to provide the right incentives through transfers and schedule the tasks optimally to achieve the highest expected profit.[13] A careful description of his strategy as the choice of a mechanism is taken up in later sections. Essentially, the agents are offered to play an extensive form game that is contingent on the information elicited at the preliminary stage. The game form comprises of a *schedule* and a profile of transfers $\boldsymbol{\tau} = (\tau_1, \dots, \tau_n)$. A schedule itself consists of i) the order in which the Principal makes a decision on the various tasks, described by a mapping $\pi : N \longrightarrow N$ and ii) a decision rule, denoted by λ, that specifies whether to commission a given task as a function of the histories.

To be precise, given a schedule (π, λ) and a profile of transfers, the agents are offered to play the following game denoted by $g = (\pi, \lambda, \boldsymbol{\tau})$. At stage k,

1. The Principal commissions a subset of task(s) with $\pi(i) = k$ as per her decision rule λ.

2. The agents whose tasks are commissioned at this stage choose whether to exert

[12]The cumulative probability F corresponding to f is said to be regular if $c + \frac{F(c)}{f(c)}$ is non-decreasing. Regularity is a standard assumption and plays the same role here as elsewhere in mechanism design, importantly, it ensures that local incentive compatibility implies global incentive compatibility.

[13]The formulation of an agent's payoff from a task does not allow for any direct benefit either from the completion of the task nor that of the project. One could easily include a benefit s_i to the agent upon the success of *his* task. This would imply repeating the analysis by redefining the cost as $c_i := c_i - q_i s_i$. "Types" are then no longer distributed identically but this assumption, unlike independence, is invoked only for Proposition 2. Including a benefit to an agent from the completion of the whole project would, on the other hand, significantly alter the analysis. For our motivational examples, however, a preference for a shared objective such as the project's success appears unreasonable.

effort.

3. The tasks' outcomes are realized. Agent i receives a transfer τ_i/q_i if and only if his task succeeds.

The decision rule λ can depend on the outcomes of tasks from previous stages as well as on the information elicited in the preliminary stage. However, as we explain in the next section, this potential complexity notwithstanding, the Principal's expected profit maximizing mechanism uses only two simple schedules. Here, we briefly motivate the two key assumptions that make this model tractable.

First, the outcomes of *individual tasks* are verifiable. This assumption seems reasonable, for example, in the case of outsourcing drug R&D. Each outsourced task implies strict contractual obligations both ways between the Principal and the labs/contractors. This allows for the later phases in the R&D to be contingent on the earlier outcomes. Verifiability of task outcomes takes us outside the scope of the literature on moral hazard in teams, where only the project's *overall* success is verifiable.

Second, given our applications it is also natural to assume that a task does not succeed "just by luck" when the agent exerts no effort. In a pure moral hazard model with binary effort choices and verifiable outcomes the assumption that the task fails for sure when the agent shirks is, in fact, without loss of generality. One could instead assume that this task succeeds with a positive albeit lower probability even without effort. The resulting incentive constraints would be equivalent to ours. The model in this paper, however, combines moral hazard with adverse selection. In such mixed models the incentive constraints limit the Principal's choice of transfers in a non-trivial manner that depends on the specification. We shall discuss this further in Section 5 following Proposition 4.[14]

2.1 Profits and the Effective Marginal Contribution Schedule

Now suppose $g = (\pi, \lambda, \boldsymbol{\tau})$ (for an arbitrary λ) is played out when the true cost profile is $\mathbf{c} = (c_1, \ldots, c_n)$. The sequentially rational behavior of each agent i is to exert effort if and only if

$$\tau_i \quad \geq \quad c_i \tag{1}$$

[14] Laffont and Martimort (2002) point out that mixed models with "false" moral hazard are tractable. These are models in which there is a deterministic relation between an agent's action and the variable that the Principal observes. This is not the case here. Upon observing that a task has failed, the Principal cannot conclude that the agent has shirked.

We refer to these as the *moral hazard constraints*. Letting 1 and 0 denote an agent's effort choice, the sequentially rational choice of agent i is $\delta(\tau_i, c_i)$, where $\delta = 1$ if $\tau_i \geq c_i$ and $\delta = 0$ otherwise. The schedule (π, λ) implies an ex-ante probability that task i is commissioned under g. With a slight abuse of notation, let $\lambda_i(g)$ denote this probability. Then, if the true cost profile is $\mathbf{c}$, the expected profit of the Principal from g is

$$P(g \mid \mathbf{c}) \;=\; \sum_{i=1}^{n} \lambda_i(g)\delta(\tau_i, c_i)(q_i r_i - \tau_i) + \alpha\lambda_0(g)\prod_{i=1}^{n} \delta(\tau_i, c_i) q_i, \tag{2}$$

where $\lambda_0(g)$ denotes the probability that all the tasks are commissioned.

Under adverse selection, the Principal designs a mechanism that specifies g as a function of the reported types. The Principal's objective is to maximize the expectation of (2) over $\mathbf{c}$. Details are in the next two sections. Here, we shall propose two important schedules that provide an upper bound on (2) for a given profile of transfers. The optimal mechanism then reduces to a choice between these two schedules at each realized $\mathbf{c}$.

Fix a profile of transfers $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_n)$ and define $q_i r_i - \tau_i$ to be the (potential) *profit flow* from task i. Under the first schedule, the Principal simultaneously commissions all the tasks with non-negative profit flows in the first stage and stops the project. We denote this schedule by $\pi^+_{\boldsymbol{\tau}}$. In particular, if every task has a negative profit flow, then $\pi^+_{\boldsymbol{\tau}}$ is simply the *status-quo* and will instead be denoted by s.

To describe the second schedule, we first introduce the notion of the *effective marginal contribution* (EMC). The EMC of task i at $\boldsymbol{\tau}$ is

$$e_i(\tau_i) \;=\; \frac{q_i r_i - \tau_i}{1 - q_i}.$$

The EMC of a given task can be interpreted as follows. Consider task i in isolation and the transfers contingent on success and failure of task i that induce agent i to exert effort. If we denote by τ_s the transfer paid upon success, the EMC of task i is the value of the transfer in case of failure, say τ_f, that makes the Principal indifferent between commissioning and not commissioning this task, that is, τ_f solves the equation $q_i r_i - q_i \tau_s - (1 - q_i)\tau_f = 0$.

Definition 1 ($\boldsymbol{\tau}$-EMC schedule) *Given $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_n)$, the $\boldsymbol{\tau}$-EMC schedule is as follows:*

1. *First, commission all the tasks with non-negative profit flows.*

2. *Then, commission exactly one task with a negative profit flow in each stage in the*

non-ascending order of their effective marginal contributions, provided <u>*all*</u> *the tasks from previous stages were successful.*

3. *Otherwise, stop the project.*

Let $\pi^*_{\boldsymbol{\tau}}$ *denote the* $\boldsymbol{\tau}$*-EMC schedule.*

Lemma 1 *Given* $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_n)$*,*

$$P(\pi, \lambda, \boldsymbol{\tau} \mid \mathbf{c}) \quad \leq \quad \max\{P(\pi^*_{\boldsymbol{\tau}}, \boldsymbol{\tau} \mid \mathbf{c}),\, P(\pi^+_{\boldsymbol{\tau}}, \boldsymbol{\tau} \mid \mathbf{c})\}.$$

for all schedules (π, λ) *and for* <u>*all*</u> $\mathbf{c}$*.*

Some parts of the proof of the Lemma are immediate. For instance, observe that the expression in (2) is a weighted sum of the profit flows of various tasks with non-negative coefficients. It is therefore clear that the value of $P(\pi, \lambda, \boldsymbol{\tau} \mid \mathbf{c})$ weakly increases by unconditionally commissioning every task with a non-negative profit flow. Now, the only reason for commissioning a task with a *negative* profit flow is to retain a positive probability of achieving α. Note that if the moral hazard constraint is not met for even one agent, α is unattainable as that agent's task necessarily fails. Therefore, if $\tau_i < c_i$ for some i, it is (weakly) beneficial not to commission any of the tasks with a negative profit flow, and $P(\pi^+_{\boldsymbol{\tau}}, \boldsymbol{\tau} \mid \mathbf{c}) \geq P(\pi, \lambda, \boldsymbol{\tau} \mid \mathbf{c})$.

Next, let us consider a $\mathbf{c}$ such that $\tau_i \geq c_i$ for all i. Every commissioned agent exerts effort. Nonetheless, if a task fails, α is still unattainable and, as before, there is no reason to commission a task with a negative profit flow. Therefore all tasks with a non-negative profit flow *must* precede tasks with a negative profit flow. Moreover, even if no task has failed, it is not optimal to commission a task i with a negative profit flow simultaneously with another task j. Instead, the Principal should save on the negative profit flow of task i by commissioning task j first and proceeding with task i only if the former succeeds.

In view of the above, the Principal only needs to choose between $\pi^+_{\boldsymbol{\tau}}$ and the schedules in which tasks with negative profit flows are commissioned only one at a time and only if all prior tasks are successful. (Tasks with non-negative profit flows are of course commissioned unconditionally). The $\boldsymbol{\tau}$-EMC schedule has these features. That tasks should be commissioned in the non-ascending order of their EMCs (as in the $\boldsymbol{\tau}$-EMC schedule) follows from a standard interchange argument. That is, pick any pair of tasks i and j that occur at consecutive positions $k = \pi(i)$ and $k + 1 = \pi(j)$, and switch their positions to obtain the schedule π'. The proof of Lemma 3 essentially verifies that

$\mathrm{sgn}(P(g \mid \mathbf{c}) - P(g' \mid \mathbf{c})) = \mathrm{sgn}(e_i(\tau_i) - e_j(\tau_j))$. If at a given cost profile the prospect of getting α outweighs the expected loss from completing the negative profit flow tasks in the $\boldsymbol{\tau}$-EMC schedule, the optimal schedule is $\pi^*_{\boldsymbol{\tau}}$. Otherwise, the optimal schedule is $\pi^+_{\boldsymbol{\tau}}$.

The EMC of a task is therefore an index of "precedence desirability" that depends *only* on the characteristics of the task. This mirrors the famed Gittins index for multi-armed bandit problems and the scheduling rule in Weitzman (1979).[15] That a simple index, the EMC, can be used to schedule the tasks optimally is noteworthy. In most formulations of task-scheduling, even when posed as one person decision problem with deterministic payoffs, the solution is NP-complete and requires heuristic arguments to obtain approximate solutions. The simplicity of the characterization in Lemma 1 for the optimal schedule is critical for studying scheduling in the agency framework.

3 Optimal Scheduling without Adverse Selection

This section considers the case where the agents' costs profile $\mathbf{c}$ is also common-knowledge along with $\mathbf{q}$, $\mathbf{r}$ and α. Now, setting transfers $\boldsymbol{\tau} = \mathbf{c}$ is enough to ensure that moral hazard constraints (1) hold for every i. Optimal schedule is then given by a simple Corollary to Lemma 1.

Corollary 1 *Suppose* $\mathbf{c}$ *is common-knowledge. The profit maximizing strategy for the Principal is to offer the game form* $(\pi^+_{\mathbf{c}}, \mathbf{c})$ *if* $P(\pi^+_{\mathbf{c}}, \mathbf{c} \mid \mathbf{c}) \geq P(\pi^*_{\mathbf{c}}, \mathbf{c} \mid \mathbf{c})$ *and* $(\pi^*_{\mathbf{c}}, \mathbf{c})$ *otherwise.*

Remark 1 (Priority of the essential tasks over the easier tasks) A pair of tasks are said to exhibit a *tradeoff* if one is the easier but the other is the more essential task. That is, tasks i and j exhibit a tradeoff if and only if

$$(q_i - q_j)(q_i r_i - q_j r_j) \quad < \quad 0 \tag{3}$$

In fact, when there is a tradeoff between i and j and say task j is the more essential task, its higher mean benefit, i.e. $q_j r_j > q_i r_i$ is accompanied by a greater variance of

[15] In Weitzman (1979), tasks can be executed in any order and the process can stop at any time. The Principal's payoff is $\max_i\{x_i\}$ over the tasks executed net of their combined cost. The reserve price of task i is w_i that solves $E_i[\max\{x_i - w_i, 0\}] = c_i$ with c_i being the cost and x_i the benefit that results from task i. It is optimal to execute the tasks in the descending order of their reserve prices. Weitzman's reserve price increases in the "marginal benefit" and decreases in the cost of the task. Since his Decision Maker expects a positive profit from any task, the reserve prices are positive. In our setting the profit flow from a task can be negative and hence the EMC of a task may a negative number.

its benefits (and profit flows), i.e. $(1-q_j)q_jr_j > (1-q_i)q_ir_i$. In that sense the more essential task is also the riskier task. A natural question is whether the more essential or the easier tasks are given priority in the optimal schedule.

Of course, the answer also depends on the costs. However note that when a tradeoff exists and both tasks receive an equal transfer, the EMC of the more essential task is higher, i.e. $e_j(\tau) > e_i(\tau)$. Since the $\mathbf{c}$-EMC schedule is optimal when $\mathbf{c}$ is common-knowledge, we then conclude:

> *Whenever there is a trade off between the tasks, and the more essential one also costs less, the more essential task is commissioned earlier.*

Moreover, suppose a task's cost is drawn at random from $[\underline{c}, \bar{c}]$ as per the original assumption, and the Principal chooses the optimal schedule point-wise, subject to the moral hazard constraints, at each $\mathbf{c}$. In view of the above observation, the following answers the question of priority in probabilistic terms:

Claim. *Whenever a pair of tasks exhibits a tradeoff, the event where the more essential task precedes the easier task has a higher ex-ante probability than the complementary event where the easier task is given a priority.*

Remarkably, as we show later in Proposition 2, the above Claim holds even under adverse selection. A formal proof of the above claim is a direct analogue of the proof of that proposition. We present below a simple geometric argument for the case of two tasks.

Suppose there are only two tasks, $n = 2$. Let task 2 (task 1) be the more essential (the easier) task. First consider the case where $q_ir_i < \underline{c}$ for $i = 1, 2$. In this case the profit flows of both tasks are negative when transfers are set equal to $\mathbf{c}$ at each realization. Therefore, only three schedules can occur under the profit maximizing strategy as $\mathbf{c}$ varies: the status-quo s, π_1 and π_2, where π_1 means that the Principal first commissions task 1 and then task 2 if and only if the earlier task is successful. Schedule π_2 is defined analogously. The regions in which each of these are optimal are depicted on the Panel 1, Figure 1. Noting that the profit from the status-quo is zero, the regions where π_1, π_2 and s are optimal are determined by the three lines ℓ_1, ℓ_2 and ℓ_e. ℓ_i is the locus of all tuples (c_1, c_2) such that $P_i(\pi_i, c_1, c_2 \mid \mathbf{c}) = 0$ for $i = 1, 2$, while $\ell_e(c_1)$ is the locus of all the tuples with equal EMC, i.e., $e_1(c_1) = e_2(c_1)$. $P(\pi_1, \mathbf{c} \mid \mathbf{c}) = P(\pi_2, \mathbf{c} \mid \mathbf{c})$ on this line.

Given the tradeoff, the following may be readily verified: i) $\ell_e(c_1)$ has a slope $(1 - q_2)/(1 - q_1) > 1$ with a positive y-intercept, ii) the line $\ell_1(c_1)$ is flatter than the line

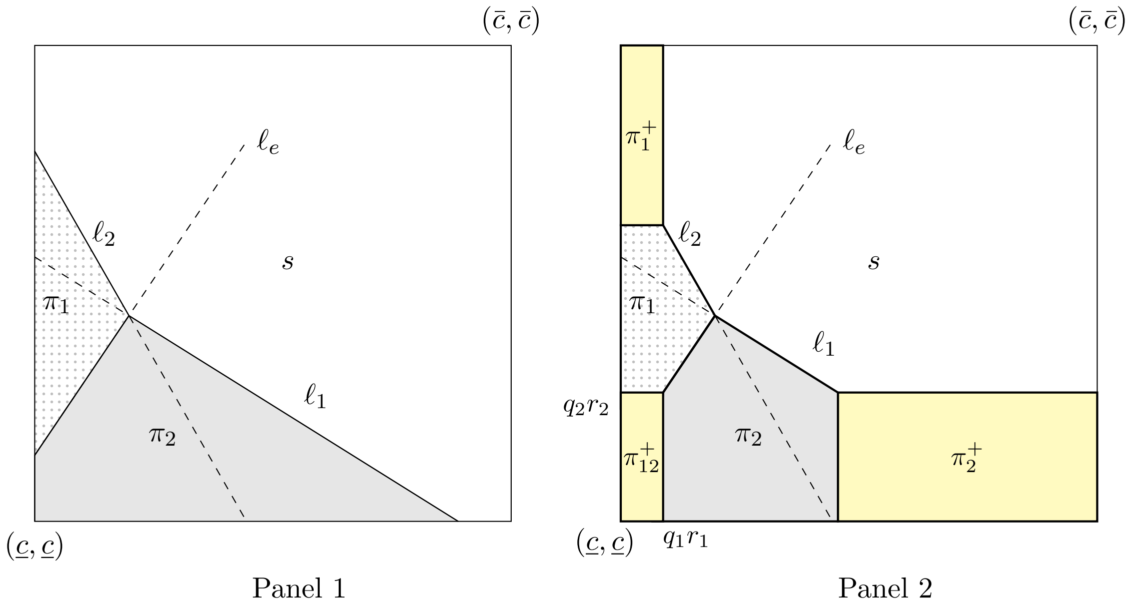

Figure 1: Reports of agent 2 in $\{v_a, v_b, v_c\}$ and the corresponding thresholds in agent 1's types that determine his transfers.

$\ell_2(c_1)$ and iii) all three lines have a common intersection. Therefore, the grey region (where π_2 is optimal) has greater area than the dotted region (where π_1 is optimal).

A similar construction arises when $q_i r_i > \underline{c}$ for some i. The case where this inequality holds for both i is depicted on Panel 2, Figure 1. In comparison to Panel 1, three new regions arise. Both tasks are commissioned unconditionally in region $\pi_{1,2}^+$. Only task 1 (only task 2) is commissioned in the region (π_1^+) (correspondingly π_2^+).

4 Optimal Scheduling with Adverse Selection

4.1 Incomplete Information & Direct Revelation Games

We now return to the case where costs are private information. The Principal is now subject to both moral hazard and adverse selection problems. We take a strategy for the Principal to be a *direct revelation game*. Formally, a direct revelation game is denoted by a pair of functions $\mu := \langle \psi, \mathbf{t} \rangle$ where $\psi : [\underline{c}, \bar{c}]^n \longrightarrow \mathcal{P} \cup \{s\}$ and $\mathbf{t} : [\underline{c}, \bar{c}]^n \longrightarrow \mathbb{R}^n_+$. In stage 0, the Principal commits to a μ. In stage one, agents simultaneously report their types. Given the reported $\hat{\mathbf{c}} = (\hat{c}_1, \ldots, \hat{c}_n)$, the game ends with the status-quo as the outcome if $\psi(\hat{\mathbf{c}}) = s$, otherwise the game proceeds to stage two. In this stage $\hat{\mathbf{c}}$ is made public and the agents play $g = (\pi, \boldsymbol{\tau})$ with $\pi = \psi(\hat{\mathbf{c}})$ and $\boldsymbol{\tau} = (t_1(\hat{\mathbf{c}}), \ldots, t_n(\hat{\mathbf{c}}))$.

Most generally, to implement the desired outcome one could allow arbitrary extensive forms in the Principal's choice set. Restricting attention to direct revelation games is, however, without loss of generality despite the multi-stage nature of the interaction here. Conceivably, in addition to the communication with the Principal the agents may exchange messages between themselves or obtain the information about the types of the other agents by observing their actions. Such information, however, is not payoff relevant as the allocation and the transfers depend only on the types reported in stage one and the costs of the agents are known to them. The (conditional) message exchanges can be incorporated into the agents strategies as in Myerson (1982) so that his Revelation Principle fully applies here.[16]

Every μ induces a multi-stage game of incomplete information among the agents. A strategy for agent i consists of a pair of functions (R_i, δ_i), where $R_i(c_i)$ is type c_i's report in stage one and $\delta_i(\hat{\mathbf{c}} \mid c_i) \in \{0, 1\}$ denotes his effort choice in stage two. To ensure sequential rationality, at each $\mathbf{c}$, the effort choices must obey the moral hazard constraints with respect to the transfers $\mathbf{t}(\hat{\mathbf{c}})$ if $\hat{\mathbf{c}} = (R_1(c_1), \ldots, R_n(c_n))$ were reported in stage one. The expected profit of the Principal at $\mathbf{c}$ would then be $P(g \mid \mathbf{c})$, where $g = (\psi(\hat{\mathbf{c}}), \mathbf{t}(\hat{\mathbf{c}}))$. Let $P(s \mid \mathbf{c}) = 0$ denote her profit when the game does not progress to stage two. Then the *equilibrium* expected profit can be expressed as

$$\Pi(\mu) \quad = \quad E[P(\mu(\hat{\mathbf{c}}) \mid \mathbf{c})].$$

The Principal's objective is to maximize $\Pi(\mu)$ in the class of all direct revelation games μ. A mechanism μ^* is said to be *optimal* if it maximizes $\Pi(\cdot)$ across all direct revelation games, and $\Pi^* = \Pi(\mu^*)$ is the *optimal profit.*

In view of Lemma 1, at every cost profile in optimal mechanism, the schedule is either $\pi^+_{\boldsymbol{\tau}}$ or the $\boldsymbol{\tau}$-EMC schedule for some transfer $\boldsymbol{\tau}$. To simplify the exposition, through the rest of the paper we assume

$$q_i r_i \quad < \quad \underline{c} \quad \forall\, i = 1, \ldots, n. \tag{4}$$

This implies that the transfer required to induce effort on a task necessarily results in a negative profit flow. Then π^+ is just the status-quo. We emphasize that except for inclusion of cumbersome expressions for incentive compatible transfers to the types with positive profit flows, violation of (4) does not alter our analysis.

[16]Even though the Principal faces a dynamic problem, each agent has only one of piece information and acts only once, therefore the Revelation Principle for models with adverse selection & moral hazard in static environments is applicable here.

4.2 Optimal Mechanism and Optimal Profit

For a given cost c_i, define its corresponding *virtual cost* $v_i = c_i + \frac{F(c_i)}{f(c_i)}$. Let $\mathbf{v} = (v_1, \ldots, v_n)$ denote the profile of the virtual costs for a given cost profile $\mathbf{c} = (c_1, \ldots, c_n)$. The following is the main result on optimal scheduling.

Proposition 1 *In an optimal mechanism, for any cost realization* $\mathbf{c}$*, either the* $\mathbf{v}$*-EMC schedule is implemented or the status-quo is retained. Moreover the optimal profit is*

$$\Pi^* \quad = \quad E[\max\{P(\pi^*_{\mathbf{v}}, \mathbf{v} \mid \mathbf{c}), 0\}]. \tag{5}$$

Proposition 1 says that to achieve Π^* the Principal behaves at each $\mathbf{c}$ as if she faces the agents with costs $\mathbf{v}$ and chooses the ex-post optimal $\mathbf{v}$-EMC schedule or the status-quo. This resembles the optimal allocation in typical auction-like IPV environments. This simplicity is noteworthy, given the usual complexity that moral hazard brings to models of adverse selection.[17] In IPV environments with only adverse selection, the allocation and the transfers depend *only* on the *reported* profile of types. This important feature, in particular, leads to the well known payoff equivalence theorem.

Here, an allocation refers to both the task schedule and the required actions along that schedule. These actions are subject to moral hazard. Therefore, whether the desired allocation (which itself is contingent on the reported profile of types) is realized in stage two also depends on the *actual* profile of types. In particular, the expected probability that an agent exerts effort (and hence incurs the cost) depends both on his true type *and* his report. A key step in the proof of Proposition 1 is to show that despite this, one can bring the familiar Envelope Theorem methods to obtain a canonical representation for the transfers. This enables us to obtain the RHS of (5) as an upper bound for the optimal profit. To complete the proof, we construct an explicit direct revelation game, $\mathcal{G}^*$ which achieves the upper bound.

4.2.1 Optimal Direct Revelation Game for Two Tasks

In this section we illustrate the construction of the above game $\mathcal{G}^*$ on the special case of $n = 2$. The general case is presented in the Appendix. We can reprise the discussion

[17] The reason is subtle. In mixed models, the incentive constraints at the reporting stage must take account of the agent's incentives to exert effort when commissioned. Usually, the complication arises from the fact that a high cost agent must be prevented from pretending to be a low cost agent at the reporting stage while planning to shirk later, see e.g., Chakraborty et al. (2017). In our setting since the task fails for sure without effort, such an incentive constraint is equivalent to an ex-post participation constraint.

leading to Panel 1, Figure 1, and appealing to Proposition 1, we depict the schedules that are optimal in different states in Figure 2. Except for the fact that under adverse selection the axis measure *the virtual costs* instead of the actual costs this Figure is identical to the earlier Figure 1. The three lines ℓ_1, ℓ_2 and ℓ_e. ℓ_i are now the locus of all tuples (v_1, v_2) such that $P_i(\pi_i, v_1, v_2 \mid \mathbf{c}) = 0$ for $i = 1, 2$, while $\ell_e(v_1)$ is the locus of all the tuples with equal EMC, i.e., $e_1(v_1, v_2) = e_2(v_1, v_2)$.

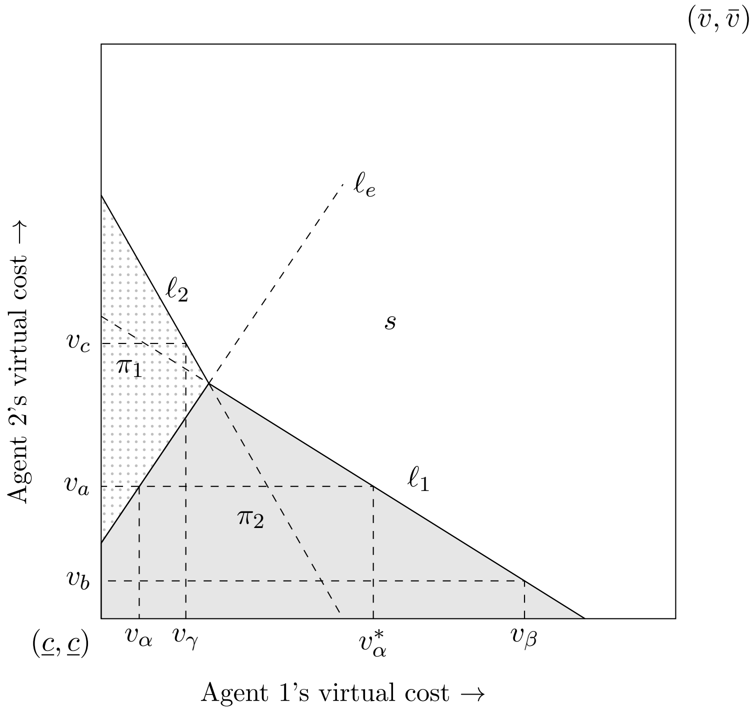

Figure 2: Reports of agent 2 in $\{v_a, v_b, v_c\}$ and the corresponding thresholds in agent 1's types that determine his transfers.

To complete the description of $\mathcal{G}^*$, we need to specify the transfers that elicit the true costs in equilibrium. We shall use Figure 2 as a reference. For each report $\hat{c}_2$ or equivalently $\hat{v}_2$ of agent 2, partition agent 1's reports into at most three regions where the schedules π_1, π_2 and the status-quo are correspondingly optimal. For example, for $\hat{v}_2 = v_a$, let c_α and c_α^* denote the costs corresponding to the thresholds v_α and v_{α^*}. These partition agent 1's types into $\{[\underline{c}, c_\alpha], [c_\alpha, c_\alpha^*], [c_\alpha^*, \bar{c}]\}$.

As c_1 increases, the position at which agent 1 acts in the optimal schedule changes from position 1 to position 2, and then to the status-quo, with the switches occurring at the threshold types c_α and c_α^*. Depending on which segment agent 1's report $\hat{c}_1$ belongs to, his transfer is a constant τ_1, τ_2 or 0. The transfers are chosen to make the threshold types indifferent between their adjacent positions. That is $(\tau_1 - c_\alpha) = q_2(\tau_2 - c_\alpha)$ and $\tau_2 - c^*_\alpha = 0$.

Type c_α is indifferent between acting at position 1 and receiving transfer τ_1 and acting at position 2 and receiving a higher transfer τ_2. Shifting to a later position implies lower probability of being called into play which lowers the expected payoff. The types below c_α prefer the earlier position, the types above c_α prefer the later position. Every type below c^*_α prefers being commissioned to staying idle. Therefore agent 1 benefits from reporting her own cost truthfully regardless of whether the report of agent 2 is truthful or not. The argument in the Appendix generalizes the above construction to the case of n tasks and verifies that moral hazard constraints are also met.

Note that depending on $\hat{v}_2$, the partition may include only two segments. As can be seen in Figure 2 for $\hat{v}_2 = v_b$ or $\hat{v}_2 = v_c$, the corresponding partitions are $\{[\underline{c}, c_\beta], [c_\beta, \bar{c}]\}$ and $\{[\underline{c}, c_\gamma], [c_\gamma, \bar{c}]\}$ respectively where c_β and c_γ denote the costs corresponding to agent 1's virtual costs v_β and v_γ.

4.3 Tradeoffs and Optimal Scheduling

Remark 1 in Section 3 describes how the trade off between the tasks is reflected in their relative positions in the optimal schedule. In this Section we generalize this to the case of asymmetric information. Since the costs are unknown, the question is which task has a higher *expected probability* of precedence? That is, "on average" do we expect an easier or a more essential task to be scheduled earlier?

The answer is fairly transparent if the project has exactly two tasks that exhibit a tradeoff. Figure 2 is in fact drawn for case where task 1 is easier and task 2 is more essential, i.e., $q_1 > q_2$ but $q_1 r_1 \leq q_2 r_2$. Then the area under ℓ_2 is larger than the area under ℓ_1 and ℓ_e is steeper than the 45 degree line.

With identically distributed costs this translates into a higher expected probability of precedence for task 2. Indeed, for any profile (v_1, v_2) where π_1 is the optimal schedule, in a symmetric profile (v_2, v_1) the optimal schedule is π_2. Given that the costs are identically distributed, schedule π_2, where the more essential task is commissioned first, is chosen more often than schedule π_1. This is an important observation that allows to predict the properties of the optimal task schedule ex-ante, relying only on the characteristics of the individual tasks – the more essential task on average precedes the easier task in the case of a tradeoff. This makes formal an intuitive approach to a complex problem where the key issues are attempted first. Their successful resolution makes the entire project viable, and opens the gate to the less essential tasks that are deferred until then.[18]

The following proposition is thus a generalization of Remark 1 to incomplete information.

Proposition 2 *Consider n tasks and suppose that task j is weakly more essential than task i, but task i is easier. In the optimal mechanism, the event that task j precedes task i is more likely than the event that task i precedes task j.*

Note that the above argument involving only two tasks used the fact that in a flip of the virtual costs from (v_1, v_2) to (v_2, v_1), the only contention is a switch from π_1 to π_2. With n tasks there are many possible alternative schedules to consider and, in principle, the probability that a given task is commissioned may increase or decrease. The key observation used in the proof is that since at a cost profile $\mathbf{c} = (c_i, c_j, \mathbf{c}_{-i,j})$ the tasks are ordered by their EMC's, the less essential task i can precede the more essential j only if i is relatively cheaper. Then at a symmetric cost profile $\mathbf{c}' = (c_j, c_i, \mathbf{c}_{-i,j})$ task j will have the cost advantage and will precede task i. The main difficulty in the proof of Proposition 2 lies with establishing that whenever the status-quo is not optimal at profile $\mathbf{c}$ it is also not optimal at $\mathbf{c}'$.

Remark 2 Recall that in the O-ring production model alluded to at the beginning of Section 2, $r_i = 0$ for all i. With complete information, the EMC reduces to the simple ratio $-c_i/(1-q_i)$. Now, as in Kremer (1993) if costs are known and equal across the agents, those with higher q's are assigned to later stages. (See also Sobel (1992)) Proposition 1 and Proposition 2 would show the same to be true in expectation when the costs are private information.

4.4 Ex-ante Investments in Technology

Before embarking on the project, the Principal could make ex-ante investments that will increase q_i or r_i for some tasks to improve the expected value of the project. For example, it is common to have tangible performance standards to determine if a task is a success. In particular, suppose task i is deemed a success if the verifiable realization of a random variable X_i exceeds an exogenously given performance standard, say x_i^*. Assuming that the distribution of X_i is G_i if the agent exerts effort then $q_i = (1 - G_i(x_i^*))$. Assume the standard cannot be met if the agent shirks. The Principal can either engage in costly

[18] For more detailed description of the stage-gate process see Cooper (1990).

lobbying to reduce x_i^*, or alternatively (at a cost) improve production standards leading to an improvement in G_i. Either way, she seeks to increase q_i.[19]

Similarly, in anticipation of the project's failure, the Principal could make early investments into "drug re-purposing or drug re-profiling" of the compounds that resulted from the successful tasks, see Oprea et al. (2011)). She could also develop alternative uses for the outcomes within a successful project through advertisement and public relations campaigns, etc. These amount to the Principal increasing r_i at a cost.

Any ex-ante investments in either q_i or r_i increase the expected profit from the optimal mechanism $\Pi^* = \Pi^*(\mathbf{q}, \mathbf{r}, \alpha)$. The question, which of these instruments has the highest marginal effect?

Proposition 3 *Suppose $n = 2$ and the two tasks exhibit a tradeoff. When α is large enough, i.e. if the tasks are sufficiently complementary, the highest marginal effect on the expected profit is with respect to the success probability of the more essential task.*

In other words, if the Principal can choose whether to spend the marginal dollar on increasing the value of drugs by re-purposing, or engage in lobbying activities to lower the threshold for success (i.e. lower x_i^*), she should do the latter.

5 Discussion

The tractability of our setup hinges on obtaining an index analogous to the EMC of a task that depends *only* on the parameters of that task to create the optimal schedule. Even abstracting from agency concerns altogether, the existence of such an index in the one person decision problem requires that i) the task outcomes must be independently distributed across tasks and ii) the value of the project essentially has one of two specifications: either it is a weighted sum of individual task outcomes (as it is here) or it is the maximum of the task outcomes. The additive specification is related to the existence of the well-known Gittins Index for multi-armed bandit problems, and the latter to the celebrated "Pandora Box" model of Weitzman (1979)'s search for the best alternative. Appendix B shows that our approach can be used to readily incorporate agency concerns into Weitzman (1979) with binary outcomes.

[19] In environmental regulation, the task would be deemed a success if its total pollution does not exceed a certain threshold. For example, production of electrical components requires insulation with PCB materials that are known to have adverse effects on humans. Production may necessarily result in a random fraction of defective items, each of which counts toward pollution. If the Principal is able to influence G_i, then she can execute technological changes that reduce the fraction of the defective items.

Mixed models of moral hazard and adverse selection do not usually yield closed form solutions in contrast to ours. The binary structure of our setup (two outcomes and two actions) contributes to the tractability here. A setting with continuous effort choices and continuous outcomes requires a separate analysis. Retaining the assumption of binary outcomes for any task, the model readily extends to the case of continuous effort choice under pure moral hazard.

Scheduling under pure moral hazard and continuous effort choice. Indeed, let $q_i(e)$ be the probability that task i succeeds if agent i's effort is $e \geq 0$ and c_i is his commonly known marginal cost of exerting effort. Agent i receives a transfer, say τ_i, if and only if his task is a success.[20] For simplicity, and without an essential loss of generality, assume the counterpart of (4), i.e. $q_i(e)r_i < c_i$ for all i so that profit flow of every task is negative. As before, it suffices to think of schedules denoted by a permutation π of N with the interpretation that i is commissioned at stage $\pi(i)$ only if all the prior tasks are a success. Having chosen a schedule π, and a profile of transfers $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_n)$, *if* the effort profile is $\boldsymbol{e} = (e_1, \ldots, e_n)$, then the probability that agent i is commissioned is $\lambda_i(\pi, \boldsymbol{\tau} \mid \boldsymbol{e}) = \Pi_{\ell=1}^{\pi(i)-1} q_i(e_i)$ and the resulting profit is

$$P(\pi, \boldsymbol{\tau} \mid \boldsymbol{e}, \boldsymbol{c}) \quad = \quad \sum_{i=1}^{n} \lambda_i(\pi, \boldsymbol{\tau} \mid \boldsymbol{e}) q_i(e_i)(r_i - \tau_i) + \alpha \prod_{i=1}^{n} q_i(e_i). \tag{6}$$

Effort choices of course depend on the transfers. In fact, given $\boldsymbol{\tau}$, the resulting effort $e_i = e_i(\tau_i)$, where

$$e_i(\tau_i) \quad = \quad \text{argmax}_e \{\tau_i q_i(e) - c_i e\} \qquad \text{for each } i. \tag{7}$$

(Assume that q_i is well-behaved so that this argmax is well-defined.) The Principal maximizes the profit $P(\pi, \boldsymbol{\tau} \mid \boldsymbol{e}, \boldsymbol{c})$ with respect to the schedule and the transfers $(\pi, \boldsymbol{\tau})$ subject to (7). Let $(\pi^*, \boldsymbol{\tau}^*)$ denote the optimal choice and $e_i^* = e_i(\tau_i^*)$ denote the induced effort.

Proposition 4 *The optimal schedule π^* is the $\boldsymbol{\tau}^*$-EMC schedule with $q_i = q_i(e_i^*)$.*

Incorporating adverse selection in the above analysis is not straightforward. In mixed models, the incentive constraints at the reporting stage must take account of the agent's incentives to exert effort when commissioned. Usually, the complication arises from the

[20] Consequently, the expression for profit given in (6) is slightly different from (2), where a transfer of τ_i/q_i was assumed to be the payment *conditional* on the task's success.

fact that a high cost agent must be prevented from pretending to be a low cost agent at the reporting stage while planning to shirk later. (For example see Chakraborty et al. (2017).) With binary choices, assuming that the task fails for sure without effort, such an incentive constraint is equivalent to an ex-post participation constraint. With continuous effort choices, such simplification is not feasible.

Scheduling under pure adverse selection and positions auctions. With pure adverse selection, our scheduling problem resembles but does not reduce to the problem of "positions auctions". Positions auctions have attracted a fair deal of attention recently, especially given their application to the sales of the internet ads by search engines. (See Varian (2007) or Edelman et al. (2007)) In these auctions, there is a given set of positions $j = 1, \ldots, n$ with given "click-through" rates $x_1 > \cdots > x_n$. That is, an ad placed in the internet search position j yields x_j "clicks" per unit of time. If this position is given to advertiser i who values a click at v_i and pays τ_i per click to the search engine, he receives a payoff $x_j(v_i - \tau_i)$.[21]

In our scheduling problem agents also effectively bid for "positions". When the chosen schedule is π, and agent i is placed at position $\pi(i) = j$, his payoff is $\lambda_i(\pi)(\tau_i - c_i)$, which looks remarkably similar to the advertiser's payoff in the positions auction. However, unlike x_j, $\lambda_i(\pi)$ depends not only on the position j. Rather, $\lambda_i(\pi)$, is given by $\prod_{\ell=1}^{\pi(i)-1} q_{\pi(\ell)}$ if the moral hazard constraints are met for every agent i' commissioned prior to agent i (i.e. prior to stage $\pi(i)$), and equals zero otherwise. In other words, the "click-through" rate here is endogenous, its value for position j depends explicitly on which agents are scheduled to act prior to stage j.[22]

In search positions auction the expected revenue is the only concern of the Principal. Her objective is more involved here. By altering the schedule the Principal affects not only the expected transfers she makes to the agents but also the expected value of the project itself. Clearly, the positions auction problem does not nest the scheduling problem studied here. Given that $x_1, \ldots, x_n$ can be an arbitrary non-increasing sequence in the positions auction, our problem also does not nest the positions auctions problem.

[21] A similar preference for the (otherwise identical) object sold earlier also may appear in sequential auctions, see e.g. Jeitschko (1999), where this preference is driven by the exogenous uncertainty over whether the second object is available for sale.

[22] Due to this, the mechanism in Edelman et al. (2007), where the private information is revealed as the allocation unfolds, cannot be used here. The Principal has to elicit the entire profile of the costs before choosing the schedule. Athey and Ellison (2011) embed position auctions in consumer search problems, which makes the value of a position endogenous, and yet independent of the schedule.

6 Conclusion

The ability of the Principal to commit to a schedule, the additive structure of the technology, private, independently drawn values and risk neutrality all contribute to a highly tractable model of optimal scheduling, despite the presence of both moral hazard and adverse selection. It would be of considerable interest to relax any subset of these assumptions. However, the problem then becomes sufficiently different and complex to warrant a separate analysis.

Given the (imperfect) complementarity of the tasks, it is natural to inquire whether it is in the Principal's interest to procure the "bundle" of services of a group of agents instead of dealing with them individually. Such an inquiry would mirror the regulatory issues explored in Gilbert and Riordan (1995) but for the case of sequential production processes. Their work effectively addresses this issue when tasks are perfect complements. Our preliminary results is that bundling is superior for a class of cost distributions.

Appendix

A Proofs

Lemma 0 below follow up on the discussion in Footnote 10.

Lemma 0. *Task i is relatively more essential than task j if and only if $q_i r_i > q_j r_j$.*

Proof of Lemma 0. Pick any i and j. $Q(S) = \prod_{k\in S,\ell\in N\setminus\{i,j\}} q_k(1-q_\ell)$ is the probability that exactly the tasks in S succeed when only the tasks in $N\setminus\{i,j\}$ have been attempted. Also let $r(S) = \sum_{i\in S} r_i$. The expected value from attempting $N \setminus \{i\}$ tasks can be expressed as

$$\begin{aligned}\Delta_{-i} &= \sum_{S\subseteq N\setminus\{i,j\}} Q(S)[q_j r(S\cup j) + (1-q_j)r(S)] \\ &= \sum_{S\subseteq N\setminus\{i,j\}} Q(S)q_j r_j + \sum_{S\subseteq N\setminus\{i,j\}} Q(S)r(S).\end{aligned}$$

Therefore,

$$\Delta_{-j} - \Delta_{-i} = \sum_{S\subseteq N\setminus\{i,j\}} Q(S)[q_i r_i - q_j r_j]$$

and $\Delta_{-j} > \Delta_{-i}$ (hence $\Delta - \Delta_{-i} > \Delta - \Delta_{-j}$) if and only if $q_i r_i > q_j r_j$. $\square$

Proof of Lemma 1. Fix a transfer profile $\boldsymbol{\tau}$ and partition the tasks into the set of those with non-negative profit flows and its complement:

$$N^+ = \{i \in N \mid q_i r_i - \tau_i \geq 0\} \quad \text{and} \quad N^- = N \setminus N^* \tag{8}$$

From the discussion that follows Lemma 1, it suffices to consider the case where $\tau_i \geq c_i$ for all i, and focus on πs such that all the tasks in N^+ precede those in N^- and no task from the latter set is commissioned simultaneously with another task. That is, we need to consider only those π such that i) $\pi(i) < \pi(j)$, for all $i \in N^+$ and all $j \in N^-$ and ii) the restriction of π to N^- is a permutation of N^- . The decision rule λ commissions all the tasks in N^+ unconditionally, and those in N^- only if every one of the tasks in the preceding stages is a success. For any such schedule, the tasks are commissioned in the game $g = (\pi, \lambda, \boldsymbol{\tau})$ with probabilities

$$\lambda_i(g) = 1 \quad \forall\, i \in N^+ \quad \text{and} \quad \lambda_i(g) = \prod_{\ell=1}^{\pi(i)-1} q_{\pi(\ell)} \quad \forall i \in N^-,$$

which in turn allows us to express the profit as

$$P(g \mid \mathbf{c}) \;=\; \sum_{i \in N^+} (q_i r_i - \tau_i) + \sum_{i \in N^-} \lambda_i(g)(q_i r_i - \tau_i) + \alpha \prod_{i=1}^{n} q_i.$$

To complete the proof, we need to show that for all such schedules, $P(g \mid \mathbf{c}) \leq P(\pi^* \mid \mathbf{c})$.

Indeed, fix such a $g = (\pi, \lambda, \boldsymbol{\tau})$ and a pair of tasks $i, j \in N^-$ at adjacent positions, i.e. $k = \pi(i)$ and $k + 1 = \pi(j)$. Let π' be the ordering of tasks obtained from π by interchanging the positions of i and j. Since the initial choice of π and the choice of the position k are arbitrary, it suffices to show that

$$\operatorname{sgn}(P(\pi, \boldsymbol{\tau} \mid \mathbf{c}) - P(\pi', \boldsymbol{\tau} \mid \mathbf{c})) \;=\; \operatorname{sgn}(e_i(\tau_i) - e_j(\tau_j)) \tag{9}$$

to complete the proof. For each agent ℓ who is scheduled to act *prior* to i we have $\lambda_\ell(\pi) = \lambda_\ell(\pi')$ and for each agent m who is scheduled to act *after* j, we have $\lambda_m(\pi) = \lambda_m(\pi')$. Hence,

$$\begin{aligned} P(\pi, \boldsymbol{\tau} \mid \mathbf{c}) - P(\pi', \boldsymbol{\tau} \mid \mathbf{c}) \;=\; & \lambda_i(\pi)(q_i r_i - \tau_i) + \lambda_j(\pi)(q_j r_j - \tau_j) \\ & - \lambda_j(\pi')(q_j r_j - \tau_j) + \lambda_i(\pi')(q_i r_i - \tau_i). \end{aligned}$$

Moreover, $\lambda_i(\pi) = \lambda_j(\pi')$, since this is simply the probability that all the tasks scheduled *prior* to k have succeeded. Further, $\lambda_j(\pi) = q_i \lambda_i(\pi)$ and $\lambda_i(\pi') = q_j \lambda_j(\pi') = q_j \lambda_i(\pi)$. Substituting in the above,

$$\begin{aligned} P(\pi, \boldsymbol{\tau} \mid \mathbf{c}) - P(\pi', \boldsymbol{\tau} \mid \mathbf{c}) &= \lambda_i(\pi)[(q_i r_i - \tau_i) + q_i(q_j r_j - \tau_j) - (q_j r_j - \tau_j) - q_j(q_i r_i - \tau_i)] \\ &= \lambda_i(\pi) \times [e_i(\tau_i) - e_j(\tau_j)] \times (1 - q_i)(1 - q_j), \end{aligned}$$

from which it is evident that (9) holds. □

Proof of Proposition 1 We begin with a few preliminary observations and prove Lemma 2. Using this Lemma, we proceed in two steps. Step 1 shows that the RHS of (5) is, in fact, an upper bound of $\Pi(\mu)$ for any direct revelation game μ. Step 2 then exhibits a mechanism that achieves this upper bound in ex-post equilibrium.

Preliminaries.

Fix a direct revelation game μ. Suppose $\mathbf{c}$ is the true cost profile but the reported profile is $\hat{\mathbf{c}}$. If $\mu(\hat{\mathbf{c}}) = s$ the game ends. Otherwise, agents are called to play according to a given schedule π determined by $\mu(\hat{\mathbf{c}})$ and transfers $(t_1(\hat{\mathbf{c}}), \ldots, t_n(\hat{\mathbf{c}}))$. Recall the interpretation of these transfers – agent i receives a payment of $t_i(\hat{\mathbf{c}})/q_i$ if and only if his task is a success. Now, regardless of i's posterior beliefs or his position in the schedule π, the payoff of type c_i is $t_i(\hat{\mathbf{c}}) - c_i$ if he exerts effort and zero otherwise. Therefore, sequential rationality requires the stage two behavioral strategy profile $(\delta_1, \ldots, \delta_n)$ to satisfy

$$\delta_i(\hat{\mathbf{c}} \mid c_i) = 1 \quad \Leftrightarrow \quad t_i(\hat{\mathbf{c}}) \geq c_i, \qquad \forall \hat{\mathbf{c}}, c_i \text{ and } \forall i \tag{10}$$

in any Perfect Bayesian Equilibrium of μ, which are essentially the moral hazard constraints (1). Correspondingly, the probability that position $k = \pi(i)$ is reached when the reported profile is $\hat{\mathbf{c}}$ and the true profile of the others is $\mathbf{c}_{-i}$ is given by

$$p_i(\hat{\mathbf{c}} \mid \mathbf{c}_{-i}) = \begin{cases} 0 & \text{if } \psi(\hat{\mathbf{c}}) = s \text{ and} \\ \lambda_i(\pi) & \text{if } \psi(\hat{\mathbf{c}}) = \pi, \end{cases} \tag{11}$$

for a non-trivial schedule π.

Now, appealing to the Revelation Principle, we need to only look for a Bayesian

Equilibrium of the direct revelation game which induces truth telling, i.e. $R_i(c_i) = c_i$ for all c_i. Consider the incentives of agent i to report her type truthfully given that all the other agents choose truth-telling. Define

$$\begin{aligned} p_i(\hat{c}_i \mid c_i) &= E_{\mathbf{c}_{-i}}[p_i(\hat{c}_i, \mathbf{c}_{-i} \mid \mathbf{c}_{-i})\delta_i(\hat{\mathbf{c}} \mid c_i)] \\ \text{and} \quad t_i(\hat{c}_i \mid c_i) &= E_{\mathbf{c}_{-i}}[p_i(\hat{c}_i, \mathbf{c}_{-i} \mid \mathbf{c}_{-i})\delta_i(\hat{\mathbf{c}} \mid c_i)t_i(\hat{c}_i, \mathbf{c}_{-i})]. \end{aligned}$$

Assume $\hat{c}_j = c_j$ for all the other agents. Then $p(\hat{c}_i \mid c_i)$ is the expected probability that type c_i of agent i exerts effort (and hence incurs the cost) if she reports her type as $\hat{c}_i$ in stage one. Similarly, $t_i(\hat{c}_i \mid c_i)$ is this type's expected transfer. The expected interim payoff of type c_i from (mis)reporting her type as $\hat{c}_i$ may then be written compactly as

$$U_i(\hat{c}_i \mid c_i) = t_i(\hat{c}_i \mid c_i) - p(\hat{c}_i \mid c_i)c_i. \tag{12}$$

Let $U_i(c_i) = U(c_i \mid c_i)$ be type c_i's expected payoff from telling the truth, given all the others also report truthfully. Mechanism μ is *incentive compatible* if and only if

$$U_i(c_i) \geq U_i(\hat{c}_i \mid c_i) \qquad \forall \hat{c}_i, c_i, \; \forall i. \tag{13}$$

The Principal's objective is then to maximize

$$\begin{aligned} \Pi(\mu) &= E[P(\psi(\mathbf{c}), \mathbf{t}(\mathbf{c})) \mid \mathbf{c})] \\ &= \sum_{i=1}^{n} (E_{c_i}[p_i(c_i \mid c_i)q_i r_i] - E_{c_i}[t_i(c_i \mid c_i)]) \end{aligned} \tag{14}$$

with respect to μ, subject to (13).

Lemma 2 *In a Perfect Bayesian Equilibrium of a direct revelation game μ, the ex-ante expected transfer of agent i is*

$$E_{\mathbf{c}}\left[p_i(\mathbf{c} \mid \mathbf{c}_{-i})\delta_i(\mathbf{c} \mid c_i)\left(c_i + \frac{F(c_i)}{f(c_i)}\right)\right] + U_i(\bar{c}_i).$$

Lemma 2 is our analogue of the payoff equivalence. Incentive compatibility constraints implicit in any equilibrium behavior restrict the transfers once the schedule at $\mathbf{c}$, and hence $p_i(\mathbf{c} \mid \mathbf{c}_{-i})$, are fixed. At any $\mathbf{c}$, the product $p_i\delta_i$ is the expected probability that agent i is called into action, and the transfer to agent i is *as if* he is paid v_i/q_i if and only if his task is a success. There is an important caveat to this analogy with payoff equivalence in the auction-like problems in the IPV setting. Note that the probability

$p_i(\mathbf{c} \mid \mathbf{c}_{-i})\delta_i(\mathbf{c} \mid c_i)$ itself depends on the chosen transfers. So, Lemma 2 is a canonical representation of the expected transfers rather than a claim of payoff equivalence.

Proof of Lemma 2. We first show that if μ is incentive compatible, then

$$U_i'(c_i) \quad = \quad -p_i(c_i \mid c_i).$$

Assuming that all agents other than i report truthfully, for each report $\hat{c}_i$ of i, define the random variables $X = p_i(\hat{c}_i, \mathbf{c}_{-i} \mid \mathbf{c}_{-i})$ and $Y = t_i(\hat{c}_i, \mathbf{c}_{-i})$ and let $G_i(\cdot \mid \hat{c}_i)$ be their joint cumulative distribution function. Then,

$$U_i(\hat{c}_i \mid c_i) \quad = \quad \int_0^1 \int_{c_i}^{\bar{c}} x(y - c_i)\mathrm{d}G(x, y \mid \hat{c}_i).$$

By incentive compatibility, $U_i(c_i)$ is the value function of $U_i(\hat{c}_i \mid c_i)$. By the Envelope Theorem,

$$\begin{aligned} U_i'(c_i) &= \frac{\partial}{\partial c_i} U_i(\hat{c}_i \mid c_i)\Big|_{\hat{c}_i = c_i} \\ &= -\int_0^1 x \int_{c_i}^{\bar{c}} \mathrm{d}G(x, y \mid c_i) \\ &= -E_{\mathbf{c}_{-i}}[p_i(c_i, \mathbf{c}_{-i} \mid \mathbf{c}_{-i})\delta_i(\mathbf{c} \mid c_i)] = -p_i(c_i \mid c_i). \end{aligned}$$

To complete the proof, we proceed along the lines of Myerson (1981), i.e., $U_i(c_i)$ is reconstructed from its derivative and the resulting expression is used to express the transfers using (12). After the usual change of the order of integration the ex-ante expected transfers can be written as

$$E[t_i(c_i \mid c_i)] \quad = \quad E\left[p_i(c_i \mid c_i)\left(c_i + \frac{F(v_i)}{f(v_i)}\right)\right] + U_i(\bar{c}_i), \tag{15}$$

as claimed in the statement of the Lemma. □

Step 1 - The Upper Bound

We may now substitute for $E[t_i(c_i \mid c_i)]$ in (14) and re-express $\Pi(\mu)$ as

$$\Pi(\mu) \quad = \quad E[\Gamma(\mathbf{c})] - \sum_{i=1}^{n} U_i(\bar{c}_i),$$

where

$$\Gamma(\mathbf{c}) \;=\; \sum_{i=1}^{n} p_i(\mathbf{c} \mid \mathbf{c}_{-i})\delta_i(\mathbf{c} \mid c_i)\left(q_i r_i - c_i - \frac{F(c_i)}{f(c_i)}\right) + \left(\prod_{i=1}^{n} \delta_i(\mathbf{c} \mid c_i) q_i\right)\alpha.$$

Now, there are three cases to consider. First, $\psi(\mathbf{c}) = s$, i.e. the status-quo is retained at $\mathbf{c}$, and $\Gamma(\mathbf{c}) = 0$. Second, $\psi(\mathbf{c}) = \pi$ but the transfers are such that some of the moral hazard constraints (10) violated, i.e. $t_i(\mathbf{c}) < c_i$ for some i. Then, the coefficient of α in $\Gamma(\mathbf{c})$ is zero, and by (4), $\Gamma(\mathbf{c}) \leq 0$. Finally, $\psi(\mathbf{c}) = \pi$ and $t_i(\mathbf{c}) \geq c_i$ for all i, in which case we have $\Gamma(\mathbf{c}) = P(\pi, \mathbf{v} \mid \mathbf{c}) \leq P(\pi^*, \mathbf{v} \mid \mathbf{c})$. The equality comes from the fact that $v_i \geq c_i$ and hence (1) holds for all i, and the inequality from appealing to Lemma 1. Therefore,

$$\Pi(\mu) \quad \leq \quad E_{\mathbf{v}}[\max\{0, P(\pi^*_{\mathbf{v}}, \mathbf{v} \mid \mathbf{c})\}] = \Pi^*.$$

Step 2. Achieving the Upper Bound.

The Game $\mathcal{G}^*$. The allocation rule is the $\mathbf{v}$-EMC schedule for each reported $\mathbf{c}$. The transfers are as follows. At a reported cost profile $(\hat{c}_i, \hat{\mathbf{c}}_{-i})$, there are no transfers if the status-quo is retained. Otherwise, fixing $\hat{\mathbf{c}}_{-i}$, transfers for agent i in the corresponding $\hat{\mathbf{v}}$-EMC schedule as $\hat{c}_i$ varies are determined as follows. Let e_k be the kth highest element in the set $\{e_j(\hat{v}_i) \mid j \neq i\}$. Without loss of generality $e_k \neq e_{k'}$ if $k \neq k'$. Introduce k^* so that agent i by reporting his lowest possible cost appears in position k^*, that is $e_i(\underline{v}) \in (e_{k^*-1}, e_{k^*}]$.[23] Introduce also K^* so that agent i by reporting his highest possible cost appears in position K^*, that is $e_i(\bar{v}) \in (e_{K^*-1}, e_{K^*}]$.[24] Let $L = K^* - k^*$.

Now, since F is regular, $e_i(\hat{v}_i)$ is strictly decreasing in $\hat{c}_i$. By varying $\hat{c}_i$, we obtain a sequence of exactly L thresholds $\underline{c} = c_0 < c_i^1 < \cdots < c_i^L \leq \bar{c}$, such that whenever $\hat{c}_i \in (c_i^{\ell-1}, c_i^{\ell}]$, the corresponding $e_i(\hat{v}_i) \in (e_{k^*+\ell-1}, e_{k^*+\ell}]$, hence i will be commissioned at position $k^* + \ell$ of the $\hat{\mathbf{v}}$-EMC schedule, for $\ell = 1, \ldots, L$.[25]

In the following let q_ℓ denote the success probability of task j such that $e_j(\hat{v}_j) = e_{k^*+\ell}$. This is the success probability of the agent whom i "lets to go ahead" of himself in the schedule by (mis)reporting his own type slightly above c_i^ℓ. Define recursively,

$$\begin{aligned} \tau_i^L &= \bar{c}, \\ \tau_i^\ell &= q_\ell \tau_i^{\ell+1} + (1 - q_\ell) c_i^\ell, \quad \text{for } \ell = L-1, \ldots, 1. \end{aligned}$$

[23] Ties are broken such that agent i appears earlier on the schedule.

[24] The case where such report "induces" the status-quo is treated similarly by setting $\tau_i^L = 0$.

[25] Note that k^* and K^* as well as each threshold c_i^ℓ depend on $\hat{\mathbf{c}}_{-i}$, but we keep this implicit.

Now define the transfer

$$t_i(\hat{c}_i, \hat{\mathbf{c}}_{-i}) \quad = \quad \begin{cases} \tau_i^1 & \text{if } \hat{c}_i = \underline{c}, \\ \tau_i^\ell & \text{if } \hat{c}_i \in (c_i^{\ell-1}, c_i^\ell] \quad \text{for } \ \ell = L, \ldots, 1, \ \text{ and} \\ \tau_i^L & \text{if } \hat{c}_i > c_i^L. \end{cases}$$

Truth telling is an Equilibrium in $\mathcal{G}^*$. Here we will verify that $\mathcal{G}^*$ admits a Perfect Bayesian Equilibrium with Π^* as the resulting expected payoff for the Principal. For this, let $\mathcal{G}^*_{\mathbf{c}}$ be the direct revelation game $\mathcal{G}^*$ where it is common-knowledge that the cost profile is $\mathbf{c}$. We will verify that for each agent to truthfully report his c_i in the first stage and obey the moral hazard constraints (10) in the second stage constitutes a SPE of $\mathcal{G}^*_{\mathbf{c}}$. As this will be true for every $\mathbf{c}$, for the agents to report truthfully and obey (10) will constitute a Perfect Bayesian Equilibrium in $\mathcal{G}^*$ that is also an ex-post equilibrium. In addition, at each $\mathbf{c}$, the SPE outcome is the $\mathbf{v}$-EMC schedule if $P(\pi^*_{\mathbf{v}}, \mathbf{v} \mid \mathbf{c}) > 0$ and the status-quo otherwise. This ensures that Π^* is the Principal's payoff.

Observe that each τ_i^ℓ is a convex combination of the threshold types $c_i^\ell, \ldots, c_i^L, \overline{c}$ for all ℓ except for $\tau_i^L{=}\overline{c}$. Therefore, at every $\hat{\mathbf{c}}$ such that $P_i(\pi^*_{\hat{\mathbf{v}}}, \hat{\mathbf{v}} \mid \mathbf{c}) > 0$,

$$t_i(\hat{c}_i, \hat{\mathbf{c}}_{-i}) \quad \geq \quad \hat{c}_i \qquad \forall \hat{c}_i \in \left[\underline{c}, c_i^L\right], \ \forall i. \tag{16}$$

Pick a profile $\mathbf{c}$ such that $P(\pi^*_{\mathbf{v}}, \mathbf{v} \mid \mathbf{c}) > 0$. Indeed, if every agent were to report the truth in the first stage, then, by (16), each agent has an incentive to incur the cost in the second stage and the $\mathbf{v}$-EMC schedule is, in fact, obtained. The payoff of every agent is positive (except perhaps for $c_i = c_i^L$, in which case it is zero), hence none of them prefers to induce the status-quo where the payoff is zero. We will now verify that agents also prefer not to mis-report in the first stage to alter their position in the schedule.

For any (mis)report $\hat{c}_i < c_i^L$, by definition, $P(\pi^*_{(\hat{v}_i, \mathbf{v}_{-i})}, \hat{v}_i, \mathbf{v}_{-i} \mid \mathbf{c}) > 0$. So by (16) $t_j(\hat{c}_i, \mathbf{c}_{-i}) \geq c_j$ for all $j \neq i$ hence all these agents exert effort when asked to act. Given this, each threshold type c_i^ℓ, by construction, is indifferent between acting at position $k^* + \ell$ (of the schedule) with transfer τ_i^ℓ and at position $k^* + \ell + 1$ with transfer $\tau_i^{\ell+1}$. Moreover, since τ_i^ℓ and c_i^ℓ are independent of $\hat{c}_i$, by varying his report $\hat{c}_i$ agent i can change his payoff only by changing his position in the schedule.

Given the linearity of the payoffs in the costs, every type $c_i' < c_i^\ell$ strictly prefers acting at position $k^* + \ell$ to acting at position $k^* + \ell + 1$, whereas the opposite is true for a type $c_i' > c_i^\ell$. Since this holds for any ℓ, agent i with $c_i \in (c_i^{\ell'-1}, c_i^{\ell'}]$, for any $\ell' = 1, \ldots, L$, prefers to act at position $k^* + \ell'$, hence prefers to report truthfully.

Proof of Proposition 2. Recall that in the optimal mechanism, it is as if at each cost profile $\mathbf{c}$, the $\mathbf{v}$-EMC schedule $\pi^*_{\mathbf{v}}$ is chosen, unless, of course, the status-quo is retained. For any i, j let $\Omega(i, j)$ denote the set of all cost profiles in which some EMC schedule (not the status-quo) is chosen, and i appears earlier than j. Now pick a pair i, j that exhibit a tradeoff, i.e.

$$q_i > q_j \quad \text{and} \quad q_i r_i \leq q_j r_j, \tag{17}$$

which means j is more essential than i but i is relatively easier.

Pick any $\mathbf{c} = (c_i, c_j, \mathbf{c}_{-i,j}) \in \Omega(i, j)$ and let $\mathbf{c}'$ be the cost profile obtained from $\mathbf{c}$ by swapping the costs of i and j. That is, $c'_i = c_j, c'_j = c_i$ and $c'_k = c_k$ for all $k \neq i, j$. We now make two claims both of which follow from (17).

First, we claim that if tasks i and j are scheduled for positions ℓ and m in the $\mathbf{v}$-EMC schedule ($\ell < m$), then at the $\mathbf{v}'$-EMC schedule j must be scheduled no later than position ℓ and i must be scheduled no earlier than position m. Hence comparing $\mathbf{c}$ and $\mathbf{c}'$, we could conclude that j is undertaken with as high probability at $\mathbf{c}'$ as i was undertaken at $\mathbf{c}$. The second claim is that $P(\pi^*_{\mathbf{v}'}, \mathbf{v}' \mid \mathbf{c}) > 0$ so that

$$\mathbf{c} \in \Omega(i, j) \implies \mathbf{c}' \in \Omega(j, i). \tag{18}$$

These two claims and the observation that costs are symmetrically distributed complete the proof.

We begin by observing that for any $\mathbf{c} \in \Omega(i, j)$, the following inequalities are satisfied for the corresponding $\mathbf{v}$:

$$e_i(v_j) < e_j(v_j) \leq e_i(v_i) < e_j(v_i) \tag{19}$$

The first and the last inequality come from (17), i.e., when the costs are equal, the more essential task also has the higher effective marginal contribution. The second inequality is from the definition of $\Omega(i, j)$.

From (19) we see that at $\mathbf{c}'$ the effective marginal contribution of task j, namely $e_j(v'_j) = e_j(v_i)$ is above that of task i at $\mathbf{c}$. Similarly, the effective marginal contribution of task i at $\mathbf{c}'$, $e_i(v'_i) = e_i(v_j)$, is below that of task j at $\mathbf{c}$. The effective marginal contributions of the other tasks are the same at $\mathbf{c}$ and $\mathbf{c}'$. This proves the first claim above on the $\mathbf{v}'$-EMC schedule.

To prove the second claim, start with the original $\mathbf{v}$-EMC schedule $\pi^*_{\mathbf{v}}$ where $\pi^*_{\mathbf{v}}(i) = k_1$ and $\pi^*_{\mathbf{v}}(j) = k_2$ where $k_2 > k_1$. Now consider the schedule π' obtained by interchanging the positions at which tasks i and j occur in $\pi^*_{\mathbf{v}}$. Note that $\pi'(j) = k_1$, $\pi'(i) = k_2$,

$v_i' = v_j$, $v_j' = v_i$ and $\pi'(m) = \pi^*_{\mathbf{v}}(m)$, $v_m' = v_m$ for all agents $m \neq i, j$. In words, $(\pi^*_{\mathbf{v}}, \mathbf{v})$ and $(\pi', \mathbf{v}')$ are identical except for the positions at which tasks i and j are commissioned. These tasks are swapped without changing the transfers to the tasks at various positions.

We claim, that

$$P(\pi^*_{\mathbf{v}}, \mathbf{v} \mid \mathbf{c}) \quad < \quad P(\pi', \mathbf{v}' \mid \mathbf{c}).$$

The proof is explicit. First compare the probabilities with which agent m is commissioned in the two schedules $\pi^*_{\mathbf{v}}$ and π'. These schedules are identical except for the interchange of agents i and j between positions k_1 and k_2. Until position k_1 neither i nor j could have acted, and after position k_2 both of them would have acted. Therefore, $\lambda_m(\pi') = \lambda_m(\pi^*_{\mathbf{v}})$ for all agents m such that $\pi^*_{\mathbf{v}}(m) \leq k_1$ or $\pi^*_{\mathbf{v}}(m) \geq k_2 + 1$. Since the agents acting at positions k_1 and k_2 were interchanged, $\lambda_j(\pi') = \lambda_i(\pi^*_{\mathbf{v}})$ and $\lambda_i(\pi') = \frac{q_j}{q_i}\lambda_j(\pi^*_{\mathbf{v}})$. Moreover, in both $P(\pi^*_{\mathbf{v}}, \mathbf{v} \mid \mathbf{c})$ and $P(\pi', \mathbf{v}' \mid \mathbf{c})$, the flow payoffs $(q_m r_m - v_m)$ are the same for all $m \neq i, j$. These payoffs $(q_i r_i - v_j)$ and $(q_j r_j - v_i)$ for i and j since π' is implemented at $\mathbf{v}'$.

Bearing these facts in mind, and letting $S = \{m : k_1 + 1 \leq \pi^*(m) \leq k_2 - 1\}$, we have

$$\begin{aligned}
P(\pi^*_{\mathbf{v}}, \mathbf{v} \mid \mathbf{c}) - P(\pi', \mathbf{v} \mid \mathbf{c}) \quad &= \quad \lambda_i(\pi^*_{\mathbf{v}})(q_i r_i - v_i) - \lambda_j(\pi')(q_j r_j - v_i) \\
&\qquad + \lambda_j(\pi^*_{\mathbf{v}})(q_j r_j - v_j) - \lambda_i(\pi')(q_i r_i - v_j) \\
&\qquad\quad + \sum_{m \in S} (\lambda_m(\pi^*_v) - \lambda_m(\pi'))(q_m r_m - v_m) \\
&= \quad \lambda_i(\pi^*_{\mathbf{v}})(q_i r_i - q_j r_j) \\
&\qquad + \lambda_j(\pi^*_{\mathbf{v}})(q_j r_j - v_j) - \frac{q_j}{q_i}\lambda_j(\pi^*_{\mathbf{v}})(q_i r_i - v_j) \\
&\qquad\quad + \left(1 - \frac{q_j}{q_i}\right) \sum_{m \in S} \lambda_m(\pi^*_{\mathbf{v}})(q_m r_m - v_m)
\end{aligned}$$

Adding and subtracting $q_j r_j$ within the parenthesis of the second last term and factoring appropriately gives us

$$\begin{aligned} P(\pi^*_{\mathbf{v}}, \mathbf{v} \mid \mathbf{c}) - P(\pi', \mathbf{v} \mid \mathbf{c}) &= \left(1 - \frac{q_j}{q_i}\right) \lambda_i(\pi^*_{\mathbf{v}})(q_i r_i - q_j r_j) \\ &\quad + \left(1 - \frac{q_j}{q_i}\right) \lambda_j(\pi^*_{\mathbf{v}})(q_j r_j - v_j) \\ &\quad + \left(1 - \frac{q_j}{q_i}\right) \sum_{m \in S} \lambda_m(\pi^*_{\mathbf{v}})(q_m r_m - v_m) \\ &< 0 \end{aligned}$$

Since $P(\pi^*_{\mathbf{v}'}, \mathbf{v}' \mid \mathbf{c}) \geq P(\pi', \mathbf{v}' \mid \mathbf{c})$ and $P(\pi^*_{\mathbf{v}}, \mathbf{v} \mid \mathbf{c}) > 0$ (since $\mathbf{c} \in \Omega(i, j)$), it follows that $\mathbf{c}' \in \Omega(j, i)$ which completes the proof of the second claim, and the proposition. □

Proof of Proposition 3. In the **v**-EMC schedule the moral hazard constraints are always met since the virtual cost exceeds the actual cost. Therefore, $P(\pi_i, \mathbf{v} \mid \mathbf{c}) = q_i A_i - v_i - q_i v_j$, where

$$A_i = r_i + q_j(r_j + \alpha) \tag{20}$$

is the expected value of the project contingent on the success of the first task, namely task i. While working through the rest of the proof, the reader may find it helpful to refer to Figure 2, drawn for the case where task 2 is relatively more essential while task 1 is more likely to succeed.

Let $G(\mathbf{v})$ denote the cumulative distribution of the valuation profiles $\mathbf{v}$, obtained from the prior distribution $F(c_i)$. Divide the space $[\underline{c}, \bar{c}]^2$ into three regions Ω_s, Ω_1 and Ω_2, corresponding to the realizations of $\mathbf{v}$ where the status-quo, and the schedules beginning with task 1 and correspondingly task 2 are optimal. The Principal's expected payoff is

$$\Pi^* = \int_{\mathbf{v} \in \Omega_1} (q_1(A_1 - v_2) - v_1) \mathrm{d}G(\mathbf{v}) + \int_{\mathbf{v} \in \Omega_2} (q_2(A_2 - v_1) - v_2) \mathrm{d}G(\mathbf{v}),$$

where A_1 and A_2 are the conditional expected values given by (20).

We will show that

$$\frac{\partial \Pi^*}{\partial q_2} = \max\left\{\frac{\partial \Pi^*}{\partial q_1}, \frac{\partial \Pi^*}{\partial q_2}, \frac{\partial \Pi^*}{\partial r_1}, \frac{\partial \Pi^*}{\partial r_2}\right\} > 0 \tag{21}$$

The regions where Ω_1 and Ω_2 are optimal, in fact, depend on (q_1, q_2, r_1, r_2). However, the expected profit from either schedule is the same on the boundary between Ω_1 and

Ω_2. Such expected profit is also 0 along the boundary between either Ω_1 or Ω_2 and Ω_s. Hence, only the direct effects (of changing (q_1, q_2, r_1, r_2) in the integrands) matter and for $i, j = 1, 2$.

$$\frac{\partial \Pi^*}{\partial q_i} = \int_{\mathbf{v}\in\Omega_i} (A_i - v_j)\mathrm{d}G(\mathbf{v}) + q_j(r_i + \alpha)\mathrm{Prob}(\Omega_j).$$

When $\mathbf{v} \in \Omega_i$, the optimal schedule starts with i, so that $q_i(A_i - v_j) - v_i \geq 0$. Hence $A_i - v_j > 0$ and $\partial\Pi^*/\partial q_i > 0$ for $i = 1, 2$.

Further,

$$\begin{aligned}\frac{\partial \Pi^*}{\partial q_2} - \frac{\partial \Pi^*}{\partial q_1} &= \int_{\mathbf{v}\in\Omega_2} (A_2 - v_1 - q_2(r_1 + \alpha))\mathrm{d}G(\mathbf{v}) \\ &\quad - \int_{\mathbf{v}\in\Omega_1} (A_1 - v_2 - q_2(r_2 + \alpha))\mathrm{d}G(\mathbf{v})\end{aligned}$$

As in the proof of Proposition 2, $q_1 > q_2$ and $q_2 r_2 \geq q_1 r_1$ imply $\Omega_1' = \{(v_2, v_1) \mid (v_1, v_2) \in \Omega_1\} \subset \Omega_2$. Therefore, given the i.i.d. of the costs,

$$\begin{aligned}\frac{\partial \Pi^*}{\partial q_2} - \frac{\partial \Pi^*}{\partial q_1} &= \int_{\mathbf{v}\in\Omega_2\setminus\Omega_1'} (A_2 - v_1 - q_2(r_1 + \alpha))\mathrm{d}G(\mathbf{v}) \\ &\quad + \int_{\mathbf{v}\in\Omega_1} [(A_2 - v_2 - q_2(r_1 + \alpha)) - (A_1 - v_2 - q_1(r_2 + \alpha))]\mathrm{d}G(\mathbf{v}) \\ &= \int_{\mathbf{v}\in\Omega_2\setminus\Omega_1'} (r_2 + (q_1 - q_2)(r_1 + \alpha) - v_1)\mathrm{d}G(\mathbf{v}) \\ &\quad + [A_2 - q_2(r_1 + \alpha) - A_1 + q_1(r_2 + \alpha)]P(\Omega_1)\end{aligned}$$

The square bracket on the last line simplifies to $(r_2 - r_1) + (q_1 - q_2)(r_1 + r_2 + 2\alpha) > 0$, since the tradeoff implies $r_2 \geq r_1$. The integrand in the second last line is also positive for sufficiently large α given that $q_1 > q_2$. Hence $\frac{\partial \Pi^*}{\partial q_2} > \frac{\partial \Pi^*}{\partial q_1}$ for sufficiently large α.

Next,

$$\frac{\partial \Pi^*}{\partial r_j} = q_i q_j \mathrm{Prob}(\Omega_i) + q_j \mathrm{Prob}(\Omega_j).$$

Hence,

$$\frac{\partial \Pi^*}{\partial q_i} - \frac{\partial \Pi^*}{\partial r_j} = \int_{\mathbf{v}\in\Omega_i} (A_i - v_j)\mathrm{d}G(\mathbf{v}) + q_j(r_i+\alpha)\mathrm{Prob}(\Omega_j) - q_i q_j \mathrm{Prob}(\Omega_i) - q_j \mathrm{Prob}(\Omega_j)$$
$$= \int_{\mathbf{v}\in\Omega_i} (r_i + q_j(r_j+\alpha-q_i) - v_j)\mathrm{d}G(\mathbf{v}) + q_j(r_i+\alpha-1)\mathrm{Prob}(\Omega_j),$$

which for sufficiently large α also implies $\frac{\partial \Pi^*}{\partial q_i} > \frac{\partial \Pi^*}{\partial r_j}$. Thus,

$$\frac{\partial \Pi^*}{\partial q_2} > \frac{\partial \Pi^*}{\partial q_1} > \frac{\partial \Pi^*}{\partial r_2} > 0$$
$$\text{and} \qquad \frac{\partial \Pi^*}{\partial q_2} > \frac{\partial \Pi^*}{\partial r_1} > 0$$

□

B Agency and the Pandora Box Problem

In our model, the value of a task is a random variable x_i ($\in \{0, r_i\}$) and the Principal's benefit after commissioning a subset of tasks S, is $\sum_{i\in S} x_i$. Perhaps the most celebrated model of scheduling and stopping in Economics is Weitzman (1979). His is a one-person decision problem where the corresponding benefit is $\max_{i\in S}\{x_i\}$. Our analysis of agency concerns in task-scheduling can be easily ported to his setting.

We first recall, that when x_i is either r_i with probability q_i and zero otherwise, for a given $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_n)$ as given, the optimal scheduling and stopping rule is given by the (Weitzman's) index

$$w_i(\tau_i) = \frac{q_i r_i - \tau_i}{q_i}.$$

$w_i(\tau_i)$ essentially plays the role of our EMC $e_i(\tau_i)$. Weitzman (1979) showed that for a given $\boldsymbol{\tau}$, the optimal strategy consists of a scheduling rule and a stopping rule as described below:

Scheduling Rule If a task is to be performed, it should be the task with the highest Weitzman index $w_i(\tau_i)$ if it is positive.

Stopping Rule Terminate the project whenever the maximum sampled reward from the already attempted tasks exceeds the Weitzman index of all the remaining tasks.

We shall call the above the $\boldsymbol{\tau}$-Weitzman rule and denote it by $\gamma^*_{\boldsymbol{\tau}}$.

Now let us introduce the agency concerns into the above model of Weitzman. That is, the outcome of task i is a success only if agent i with privately known cost c_i exerts effort. It should be fairly evident that there is nothing conceptually new in analyzing this model via our approach. Indeed, proceeding analogously to the definition of $P(g \mid \mathbf{c})$ in Section 2, we can write $\hat{P}(\gamma^*_{\boldsymbol{\tau}}, \boldsymbol{\tau} \mid \mathbf{c})$ – the expected profit from following the $\gamma^*_{\boldsymbol{\tau}}$ rule when transfers are given by $\boldsymbol{\tau}$ but the true state is $\mathbf{c}$. Since the payoffs of both the agents and the Principal are quasi-linear in transfers, we may repeat every part of Proposition 1 almost ad verbatim to obtain its direct analogue:

Proposition 5 *The highest expected profit for the Principal across all Perfect Bayesian Equilibria of direct revelation games, is*

$$\Pi^* \quad = \quad E[\max\{P(\gamma^*_{\mathbf{v}}, \mathbf{v} \mid \mathbf{c}), 0\}]. \tag{22}$$

In particular at any cost realization $\mathbf{c}$*, the optimal scheduling-stopping strategy is the* $\mathbf{v}$*-Weitzman rule.*

References

Agastya, Murali K., Parimal Kanti Bag, and Nona Pepito, "Task Ordering in Incentives Under Externalities," 2016. Available at SSRN:.

Athey, Susan and Glenn Ellison, "Position Auctions with Consumer Search," *The Quarterly Journal of Economics*, 2011, *126* (3), 1213–1270.

Buvalio, Andrii, "Pharma R&D Outsourcing Is On The Rise," Online: https://www.biopharmatrend.com/post/30-pharma-rd-outsourcing-is-on-the-rise/ Accessed: Oct 25, 2019 2018.

Chakraborty, Indranil, Fahad Khalil, and Jacques Lawarree, "Competitive Procurement with Ex Post Moral Hazard," Technical Report, Forthcoming, RAND 2017.

Cooper, R.G., "Stage-Gate Systems: A new tool for managing new products," *Business Horizons*, 1990, *33* (3), 44–54.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz, "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, March 2007, *97* (1), 242–259.

Gassmann, Oliver, Maximilian von Zedtwitz, and Gerrit Reepmeyer, *Leading pharmaceutical innovation : trends and drivers for growth in the pharmaceutical industry*, Berlin ; New York: Springer, 2004.

Gibbons, Robert and Michael Waldman, "Chapter 36 Careers in organizations: Theory and evidence," in "Handbook of Labor Economics," Vol. 3 of *Handbook of Labor Economics*, Elsevier, 1999, pp. 2373–2437.

Gilbert, Richard J. and Michael H. Riordan, "Regulating Complementary Products: A Comparative Institutional Analysiss," *The RAND Journal of Economics*, 1995, *26* (2), 243–256.

Herroelen, Willy and Roel Leus, "Project scheduling under uncertainty: Survey and research potentials," *European Journal of Operational Research*, 2005, *165* (2), 289 – 306. Project Management and Scheduling.

Herroelen, Willy S., Patrick Van Dommelen, and Erik L. Demeulemeester, "Project network models with discounted cash flows a guided tour through recent developments," *European Journal of Operational Research*, 1997, *100* (1), 97 – 121.

Jeitschko, Thomas D., "Equilibrium price paths in sequential auctions with stochastic supply," *Economics Letters*, July 1999, *64* (1), 67–72.

Kremer, Michael, "The O-Ring Theory of Economic Development," *The Quarterly Journal of Economics*, 1993, *108* (3), 551–575.

Laffont, Jean-Jacques and David Martimort, *The Theory of Incentives: The Principal-Agent Model*, Princeton University Press, 2002.

_ **and Jean Tirole**, "Auctioning Incentive Contracts," *Journal of Political Economy*, October 1987, *95* (5), 921–937.

Li, Zukui and Marianthi Ierapetritou, "Process scheduling under uncertainty: Review and challenges," *Computers & Chemical Engineering*, 2008, *32* (4), 715 – 727. Festschrift devoted to Rex Reklaitis on his 65th Birthday.

Lynch, Maggie, "Outsourced Drug Discovery Market to Top $68bn in 2028," Online: https://www.outsourcing-pharma.com/Article/2019/07/12/Outsourced-drug-discovery-market-to-top-68bn-in-2028 Accessed: Oct 25, 2019 2019.

Lynch, Timothy J., *The Oxford Encyclopedia of American Military and Diplomatic History*, Oxford University Press, 2013. ISBN 0199759251.

McAfee, R. Preston and John McMillan, "Bidding for Contracts: A Principal-Agent Analysis," *The RAND Journal of Economics*, 1986, *17* (3), 326–338.

Mitra, Manipushpak, "Mechanism design in queueing problems," *Economic Theory*, 2001, *17* (2), 277–305.

Mueller, John, *Retreat from Doomsday: The Obsolescence of Major War*, New York: Basic Books, 1989.

Mumford, Andrew, *Proxy Warfare*, 1st ed., Polity, 2013.

Myerson, Roger B., "Optimal Auction Design," *Mathematics of Operations Research*, February 1981, *6* (1), 58–73.

_ , "Optimal coordination mechanisms in generalized principal-agent problems," *Journal of Mathematical Economics*, June 1982, *10* (1), 67–81.

Mylovanov, Tymofiy and Patrick Schmitz, "Task scheduling and moral hazard," *Economic Theory*, November 2008, *37* (2), 307–320.

Olszewski, Wojciech and Richard Weber, "A more general Pandora rule?," *Journal of Economic Theory*, 2015, *160* (C), 429–437.

Oprea, Tudor I., Julie E. Bauman, Cristian G. Bologa, Tione Buranda, Alexandre Chigaev, Bruce S. Edwards, Jonathan W. Jarvik, Hattie D. Gresham, Mark K. Haynes, Brian Hjelle, Robert Hromas, Laurie Hudson, Debra A. Mackenzie, Carolyn Y. Muller, John C. Reed, Peter C. Simons, Yelena Smagley, Juan Strouse, Zurab Surviladze, Todd Thompson, Oleg Ursu, Anna Waller, Angela Wandinger-Ness, Stuart S. Winter, Yang Wu, Susan M. Young, Richard S. Larson, Cheryl Willman, and Larry A. Sklar, "Drug repurposing from an academic perspective," *Drug Discovery Today: Therapeutic Strategies*, 2011, *8* (3), 61 – 69. Drug repurposing.

Prevost, Gary, "The "Contra" War in Nicaragua," *Conflict Quarterly*, June 1987, *7* (3).

Ready, Kevin, *Startup: An Insider's Guide to Launching and Running a Business*, Apress, 2011.

Reyck, Bert De and Roel Leus, "R&D project scheduling when activities may fail," *IIE Transactions*, 2008, *40* (4), 367–384.

Riedel, Bruce, *What We Won: America's Secret War in Afghanistan, 1979–1989*, Brookings Institution Press, 2014.

Rocha, J. M. Da and M. Angeles de Frutos, "A Note on the Optimal Structure of Production," *Journal of Economic Theory*, December 1999, *89* (2), 234–246.

Schmitz, Patrick W., "Allocating Control in Agency Problems with Limited Liability and Sequential Hidden Actions," *The RAND Journal of Economics*, 2005, *36* (2), 318–336.

Severinov, Sergei, "The value of information and optimal organization," *The RAND Journal of Economics*, 2008, *39* (1), 238–265.

Sobel, Joel, "How to Count to One Thousand," *The Economic Journal*, 1992, *102* (410), 1–8.

Suijs, Jeroen, "On incentive compatibility and budget balancedness in public decision making," *Economic design*, Dec 1996, *2* (1), 193–209.

Varian, Hal R., "Position auctions," *International Journal of Industrial Organization*, December 2007, *25* (6), 1163–1178.

Visiongain, "Drug Discovery Outsourcing Market Forecast 2016-2026," Technical Report, Visiongain September 2016. PRODUCT CODE: PHA0143.

Weitzman, Martin L., "Optimal Search for the Best Alternative," *Econometrica*, 1979, *47* (3), 641–654.

Winter, Eyal, "Optimal incentives for sequential production processes," *RAND Journal of Economics*, June 2006, *37* (2), 376–390.