

Incomplete Information VCG Contracts for Common Agency

Tal Alon* Ron Lavi† Elisheva S. Shamash‡ Inbal Talgam-Cohen§

June 1, 2021

Abstract

We study contract design for welfare maximization in the well-known “common agency” model of Bernheim and Whinston [1986]. This model combines the challenges of coordinating multiple principals with the fundamental challenge of contract design: that principals have *incomplete information* of the agent’s choice of action. Motivated by the significant social inefficiency of standard contracts for such settings (which we formally quantify using a price of anarchy/stability analysis), we investigate whether and how a recent toolbox developed for the first set of challenges under a complete-information assumption – VCG contracts [Lavi and Shamash, 2019] – can be extended to incomplete information.

We define and characterize the class of “incomplete information VCG contracts (IIVCG)”, and show it is the unique class guaranteeing truthfulness of the principals and welfare maximization by the agent. Our results reveal an inherent tradeoff between two important properties required to ensure participation in the contract: individual rationality (for the principals) and limited liability (for the agent). We design a polynomial-time algorithm for determining whether a setting has an IIVCG contract with both properties. As our main result we design a polynomial-time “algorithmic IIVCG” contract: given valuation reports from the principals it returns, if possible for the setting, a payment scheme for the agent that constitutes an IIVCG contract with all desired properties. We also give a sufficient graph-theoretic condition on the population of principals that ensures the existence of such an IIVCG contract.

Keywords: VCG; Principal; Common Agent; Contract; Equilibrium; Limited Liability; Individual Rationality; Polynomial Complexity

*Technion – Israel Institute of Technology. Email: alontal@campus.technion.ac.il.

†U. of Bath, U.K., and Technion – Israel Institute of Technology. Email: ronlavi@ie.technion.ac.il.

‡Technion – Israel Institute of Technology. Email: elisheva@campus.technion.ac.il.

§Technion – Israel Institute of Technology. Email: inbaltalgam@gmail.com.

1 Introduction

Common agency. The principal-agent model is a fundamental area of microeconomic theory, capturing natural contractual relations with important practical implications [5, 21, 29, 28]. At the heart of the model is a relationship between one entity (agent) who acts on behalf of another (principal); for example, corporate management and shareholders, a freelance worker and employers, the public regulator and citizens. In all of these examples, the agent (manager/freelancer/regulator) acts on behalf of *multiple* principals (shareholders/employers/citizens). The extension of the basic single-agent, single-principal setting to multiple principals mirrors the extension of the basic single-seller, single-buyer *auction* setting to multiple buyers. This extension was introduced in 1986 by Bernheim and Whinston [3], and is known as *common agency*. This fundamental theory has many additional applications to the study of markets, for example: a marketing agency that bids for ad space on behalf of multiple advertisers [8], a professional host on Booking.com who manages several properties for the owners, or platforms like Airbnb that have characteristics of an agent representing sellers (Airbnb works to promote listings of property owners).¹

The agency problem. The main challenge in the principal-agent model (with one or more principals) is the *agency problem* [32]: Actions that generate good outcomes for the principal(s) are costly for the agent, who rationally acts in his own self-interest. To align the interests of the principals with those of the agent, the principals compensate the agent depending on the outcome of his action. This compensation scheme, designed to incentivize the agent to choose a desirable action, is called a *contract*. The agency problem has two necessary ingredients (if one of these is missing, the problem is resolved): (i) *Incomplete information* of the agent’s actions; (ii) Principals are better *risk bearers* than the agent.

Incomplete information means that the principals cannot directly observe the agent’s choice of action, only its resulting (stochastic) outcome. The contractual payment scheme thus depends on the outcome, which serves as a noisy indicator of the action. This crucial feature of the model is known as *hidden action* or *moral hazard*. For example, employers of a freelancer are not aware of how much effort he invests, but do view the outcome of his efforts and compensate him accordingly. The second crucial feature is the asymmetry between principals and agents in their attitude towards risk. The principal typically has “deeper pockets” and can bear the risk that the agent’s actions will lead to losses. The way this is incorporated into our model follows classic works like [18, 12, 6]: we focus on contracts guaranteeing *limited liability (LL)* for the agent, i.e., payments in the contractual relationship flow only from the principals to the agent. Indeed, in our examples, a manager does not typically pay shareholders (even if he loses their investment), nor does a freelancer pay the employers to purchase projects.²

¹It was in fact in legal dispute in the EU whether Airbnb is legally an estate agent or not.

²Other works like [3] capture the asymmetry in risk attitudes among principal and agent by explicitly modeling the agent’s risk averseness. Without LL or risk averseness, the trivial but unrealistic solution of “buying the project”

1.1 Multiple Principals: Challenges and a Promising Approach

Having multiple principals raises intriguing new challenges:

“The principal-agent problem here is intensified: not only is there still asymmetric information between the principals and agent that can bring moral hazard, but there is also *asymmetric information between the principals* themselves [... S]ince *principals’ interests often diverge*, they face incentives to advance their individual interests instead of the joint interests by all principals[...] As a result, introducing governance to align the interests of the principals with those of the agent is much more difficult” [31].

Thus, misalignment among the principals combined with their private information (their types) requires novel contractual mechanisms. In this paper we focus on the design of *welfare-maximizing* contracts.

First price contracts. The classic literature on multiple principals has thoroughly explored one class of contracts, which can be referred to as *first price*. A first price contract is a simple mathematical object: it is a function t^ℓ for every principal ℓ that maps an outcome generated by the agent’s action to the principal’s payment. The payment is simply the principal’s (reported) *value* for this outcome, which can obviously be different than her true value. As we show in Theorem 1, first price contracts can be very socially inefficient (where social efficiency is considered in equilibrium, through the standard measures of price of anarchy/stability).

A solution designed for complete information. A possible solution arises from recent work of Lavi and Shamash [22]. Their work studies a model with *complete* information about the agent’s choice of action, and thus no agency problem. To demonstrate how different the complete and incomplete information models are, consider a single principal and an agent choosing among two actions (“shirking” and “working”), the second costing him more effort. There are two possible outcomes (“failure” and “success”), and the principal’s value for the second is higher. Shirking is more likely to produce failure than working. The precise costs, values and probabilities appear in Example 3.5, and are set up such that working maximizes the expected social welfare (the principal’s expected value for the action’s outcome, less the agent’s cost for the action). With full information the principal can simply tell the agent to work and pay him his cost, thus maximizing welfare and extracting it fully as her expected utility. But in the incomplete information world where the agent’s action is unknown, the socially-inefficient “null” contract, which pays nothing and incentivizes the agent to shirk, is preferable for the principal (see Example 3.5 for details).

Further evidence of the difference between the two models is the state of the literature in each. For complete information, the problem of welfare maximization with multiple principals and a single agent is largely settled by Bernheim and Whinston [4] (in a paper preceding their “Common Agency” paper [3]). Lavi and Shamash [22] address a generalization of this problem to

would apply [6].

multiple agents (the generalization is due to [26]). They instantiate an abstract concept known as *contractible contracts* [25], and introduce a novel class of *VCG contracts*. In contrast, in the incomplete information world, the problem of welfare maximization with multiple principals has remained largely open even for a single agent in the three decades since the common agency model was introduced by [3]. Several works have considered the design of contracts with incomplete information from a computational point of view [2, 16, 19, 20, 10, 9]; some of these works have multiple agents (e.g., [33, 1]), or agent private types [15], but none to our knowledge consider multiple principals with private types.

VCG contracts and incomplete information. In the VCG contracts of Lavi and Shamash [22], principal ℓ 's payment for an outcome is no longer just her own (reported) value from this outcome (as in first price contracts), but is a function of the other principals' reported values as well. In this sense, the payment functions in VCG contracts are reminiscent of payments in VCG *auctions* [30, 7, 14].³ Such payments are indeed used in practice – for example, contracts with “price match guarantee” clauses. Moreover, computerized platforms (such as freelancing platforms) make VCG contracts easier to coordinate. Such platforms typically have a central intermediary who can serve as “mechanism designer”. E.g., in online labor markets, the market platform serves as an intermediary between an agent (freelancer) and the principals who hire him. The platform has control over pricing structures and usually treats social welfare as a first-order concern [27, p. 55]. Other examples include marketing agencies (multiple advertisers and a single marketing agent), corporate management (multiple managers and shareholders, and a single lower-level manager), etc.

Our goal in this work is to import ideas from VCG contracts and auctions to the incomplete information world, towards making progress on the long-standing problem of welfare maximization in common agency settings. Our first challenge is that the class of VCG contracts is undefined for incomplete information settings. After devising an appropriate generalization of the definition (Section 3.4), we uncover a more serious challenge – the adaptation of VCG contracts to incomplete information introduces negative payments. In other words, under such contracts the agent is sometimes required to pay the principal, in violation of the LL property. Our results show this is no coincidence, and ultimately address the challenge through algorithmic methods (Section 4). Our main result is a polynomial-time algorithm that implements a contract with all the desired properties, including LL, in all settings for which such a contract exists. We believe our contribution can be a first step in a potentially rich research area of utilizing contractible contracts to address agency problems: Indeed, it is as hopeless to solve all agency problems with first price contracts as it is to solve all mechanism design problems with first price auctions. More general contract classes are called for, and as with auctions, the computational viewpoint comes in handy.

³In fact, Lavi and Shamash [22] note that VCG contracts generalize second price (VCG) auctions; principals in the former replace the buyers in the latter, and an agent replaces the social planner.

1.2 Our Results

Design objectives. Motivated by the inefficiency of first price contracts, we seek new contracts that satisfy the following desiderata. (i) Incentives: Since principals’ values are privately-known, we require that principals are incentivized to report their values truthfully in dominant strategies. We also require *individual rationality* (IR) so that principals’ participation is a dominant strategy, and as explained above, also the LL property so the agent faces no risk.⁴ (ii) Objective: The payment schemes should incentivize the agent to choose a welfare maximizing action (we leave revenue maximization⁵ for future work). (iii) Computational efficiency: Our contracts should be polynomial-time computable.

Beyond first price contracts. We begin with two examples showing that when the principals offer *first price* contracts to the agent, there can be a highly inefficient equilibrium (Example 1 shows the price of anarchy is $O(\frac{1}{n})$ for n principals), and it’s even possible that all equilibria are inefficient (Example 2 shows the price of stability is $\frac{1}{2} + O(\frac{1}{q-1})$ for q actions). As our first conceptual contribution, in Section 2 we formulate a game that allows us to consider more complex contract formats beyond first price: Every principal $\ell \in [n]$ reports a valuation b^ℓ mapping outcomes to values. (The reports can be viewed as offers made to the agent, or as “bids” to a centralized computerized platform.) The agent then chooses an action which generates an outcome o . A *contract* $t = (t^1, \dots, t^n)$ maps the reports $b = (b^1, \dots, b^n)$ and the realized outcome o to a payment $t^\ell(b, o)$ from every principal ℓ to the agent.

The IIVCG family and an impossibility. In Section 3 we define the *Incomplete Information VCG (IIVCG)* contract family, which is dominant strategy truthful and always induces welfare maximization. Dominant strategy truthfulness is a significant advantage of IIVCG over first price contracts: with the latter, a decades-old question is how the principals arrive at an equilibrium, since lacking coordination they face a bargaining-like problem of how to split the welfare. Building upon a classic result of Holmström [17], we show the IIVCG family is unique (note that all results stated in this section are informal):

Theorem (Uniqueness of IIVCG, see Theorem 2). *Every contract t such that (i) the agent is induced to maximize welfare, and (ii) the principals’ dominant strategy is truthful, must belong to the IIVCG family.*

Our definition of the IIVCG family (Definition 2) specifies only the *expected* payments to the agent for every action, where the expectation is over the action’s stochastic outcome. This is not enough – a contract must determine the payments for every possible outcome realization. We

⁴One could also consider relaxing the LL constraint by requiring not that *every* contractual payment is non-negative but rather that only the *aggregate* payment from the principals to the agent is non-negative for every outcome; this does not change our results (see Remark 7).

⁵The concept of “revenue” in a multi-principal settings may require new definitions.

address this by characterizing the outcome-specific payments that incentivize welfare maximization on behalf of the agent:

Lemma (Payment characterization, see Lemma 1). *A contract $t = (t^1, \dots, t^\ell)$ is IIVCG if and only if principal ℓ 's payment function t^ℓ can be expressed using $h^\ell(\cdot)$ and w^ℓ , where $h^\ell(\cdot)$ is a function of the other principals' reports $b^{-\ell}$, and $\sum_\ell w^\ell$ is a payment vector belonging to the polytope of payments incentivizing the socially efficient action.*

This lemma is useful since it exactly defines our “design space”, and in particular translates to linear constraints that are the cornerstone of our algorithmic constructions.

IIVCG contracts satisfy truthfulness and welfare maximization; what about the other desiderata of IR and LL? Whether or not these properties hold depends on the choices of $h^\ell(\cdot)$ and w^ℓ that determine the payments for every principal ℓ . A natural starting point is to choose these to make principal ℓ 's expected payment equal to the externality she imposes on the others, as in VCG auctions with the Clarke pivot rule [7], as well as VCG contracts [22]. We refer to this instantiation of the IIVCG family as *Auction-Inspired* contracts. In Section 3.4 we show that while Auction-Inspired contracts are always IR, the LL property is not necessarily satisfied. This is no coincidence – there is in fact an inherent clash between the two properties:

Theorem (LL-IR tradeoff, see Theorem 3). *There are common agency settings for which no IIVCG contract is both LL and IR.*

The above theorem can apply even to settings with a single principal. This impossibility result highlights how challenging it can be to design contracts under incomplete information, and how fundamentally different this model is from the complete information one.

An algorithmic approach and main results. In light of our impossibility result, we adopt an algorithmic approach that leverages our characterization of IIVCG contracts and uses it constructively. We give a polynomial-time algorithm (Algorithm 2) that determines for a given common agent setting whether or not there is an IIVCG contract for this setting that is both LL and IR. If such a contract exists, we design another polynomial-time algorithm (Algorithm 1) that implements it, i.e., upon receiving reports $b = (b^1, \dots, b^n)$ and an outcome o , it computes $h^\ell(b^{-\ell})$ and w^ℓ “on the fly” for every principal ℓ in order to determine the contractual payments. The algorithmic basis of this approach gives us flexibility that may not exist when trying to design a contract with a closed-form formula for payments.

Theorem (Main result, see Theorems 5-6). *Consider a common agency setting. Algorithm 1 implements a (truthful and welfare maximizing) IIVCG contract that satisfies LL and IR whenever such a contract exists for the setting; Algorithm 2 determines existence of such a contract. Both algorithms run in polynomial time.*

The next observation demonstrates that IIVCG contracts can achieve welfare maximization for settings in which first price contracts fail to achieve so even approximately.

Observation (The power of IIVCG beyond first price, see Appendix A.3). *Algorithm 2 returns “Exists” (i.e., there exists an LL, IR IIVCG contract) when run on Examples 1 and 2, which have bad price of anarchy (respectively, stability) for first price contracts.*

Finally, in Section 5 we give an instantiation of the IIVCG family called *Weighted IIVCG*. Weighted IIVCG contracts are always LL, and we characterize the class of settings for which they are also IR – namely, settings that satisfy a graph-theoretic condition on the graph whose nodes are the principals. This completes the picture, showing the full range within IIVCG: IR and not LL (Auction-Inspired IIVCG), LL and not IR (Weighted IIVCG), and both LL and IR whenever possible (Algorithmic IIVCG). One avenue we do not explore in this paper but leave as an open question is contracts that guarantee LL and IR always, but only approximately achieve the optimal social welfare.

2 Model and Preliminaries

Common agency settings. We adopt the classic model of Bernheim and Whinston [3], in which there is a single *agent* who is common to n *principals*. The agent chooses an unobservable *action* that incurs a *cost* on the agent himself, and stochastically leads to an observable *outcome* that benefits the principals. This is formalized by a *common agency setting* as we now describe.

Every setting is a tuple $(\mathcal{A}, \psi, \mathcal{O}, \mathcal{V}, F)$ as follows. \mathcal{A} is a set of q actions available to the agent. ψ maps an agent’s action to its cost (we assume for simplicity no two actions have the same cost). \mathcal{O} is a set of m outcomes. Each principal $\ell \in [n]$ has a multiparameter *valuation* function $v^\ell : \mathcal{O} \rightarrow \mathbb{R}_{\geq 0}$, which maps an outcome to a non-negative value, and can also be thought of as a vector in $\mathbb{R}_{\geq 0}^m$ (valuations in our model are thus succinct objects). Valuation v^ℓ belongs to a *domain* \mathcal{V}^ℓ , where for simplicity we assume \mathcal{V}^ℓ is a convex, polynomially-described polytope; all our results extend to general convex domains.⁶ Valuation *profile* v is the tuple (v^1, \dots, v^n) , and \mathcal{V} is the tuple of domains $(\mathcal{V}^1, \dots, \mathcal{V}^n)$. F is a set of distributions; an action a induces a probability distribution $F|_a$ over the outcomes (with probability $p_{o \sim F|_a}(o)$ for outcome o). Throughout, principals are indexed by ℓ , actions by j and outcomes by i for consistency. Towards the computational analysis, we assume that the representation size of $(\mathcal{A}, \psi, \mathcal{O}, \mathcal{V}, F)$ is polynomial in n, q and m .

Who knows what. The setting $(\mathcal{A}, \psi, \mathcal{O}, \mathcal{V}, F)$ is assumed to be public knowledge. The realized outcome i is publicly known, while the agent’s chosen action j and the principals valuations $\{v^\ell\}$ are privately known. This makes our model a *hidden action, hidden type* model (where the type is of the principals).

⁶In the general case, the convex domains are represented by polynomial-time membership oracles. Convex valuation domains appear, e.g., as the main application in [17].

Contracts. A contract profile $t = (t^1, \dots, t^n)$ (*contract* for short) includes a payment rule t^ℓ for each principal ℓ . We diverge from the classic model in that the payment rules depend not only on the outcome of the agent’s action, but also on the reported valuations of all principals: $t^\ell : \mathcal{V} \times \mathcal{O} \rightarrow \mathbb{R}$. Such contracts are known as *contractible*. We denote principal ℓ ’s report (a.k.a. *bid*) by $b^\ell \in \mathcal{V}^\ell$ (the report need not be truthful, i.e., b^ℓ is not necessarily v^ℓ).

Taken together, the n (publicly known) payment rules incentivize the agent when choosing an action. This introduces a tension among the n noncooperative principals to coordinate the agent’s choice of action, on top of the usual tension between each principal and the self-interested agent. What constitutes a *desirable* action depends on the principals’ valuations and on the design objective (welfare, revenue, etc.). In this work we focus on welfare, as detailed below.

The game. Fix a common agency setting. A contract t and valuation profile v define the following two-stage game. First, principals simultaneously report their bids $b = (b^1, \dots, b^n)$. After perceiving the bids, the agent chooses a hidden action a , which imposes a distribution $F_{|a}$ over the outcomes. An outcome o is realized according to this distribution, and each principal ℓ pays the agent $t^\ell(b, o)$. We assume quasi-linear utilities and risk neutrality for all parties. The utilities of the game are as follows: The agent’s utility is the sum of payments less cost of chosen action, i.e., $\sum_{\ell \in [n]} t^\ell(b, o) - \psi(a)$. Thus, the agent’s action $x^*(b)$ given bid profile b maximizes his expected utility:⁷

$$x^*(b) \in \arg \max_{a \in \mathcal{A}} \left\{ \sum_{\ell \in [n]} \mathbb{E}_{o \sim F_{|a}} [t^\ell(b, o)] - \psi(a) \right\}.$$

We assume the principals know the agent is rational and can anticipate his action $x^*(b)$ given bid profile b . As for the utility of the principals, principal ℓ ’s utility is her value less payment, i.e., $v^\ell(o) - t^\ell(b, o)$.

Equilibrium. A bid profile b is in *equilibrium* if no principal can strictly gain in expectation by changing her bid while holding other principals’ bids fixed, and given the agent’s anticipated choice of action $x^*(b)$. Put differently, given bid profile $b^{-\ell}$, principal ℓ ’s equilibrium bid b^ℓ maximizes her expected utility $\mathbb{E}_{o \sim F_{|x^*(b)}} [v^\ell(o)] - \mathbb{E}_{o \sim F_{|x^*(b)}} [t^\ell(b, o)]$ among all possible bids in \mathcal{V}^ℓ , where $b = (b^{-\ell}, b^\ell)$, and $x^*(b)$ is the agent’s anticipated action.

Consider now the incomplete information game defined by contract t when the valuation profile v is not fixed. A *strategy* of principal ℓ in this game is a mapping from every possible valuation v^ℓ to a bid b^ℓ . A strategy is called *dominant* if it is guaranteed to maximize the principal’s expected utility for every v^ℓ no matter what the other principals bid, i.e., for every $b^{-\ell} \in \mathcal{V}^{-\ell}$ (and still

⁷As is standard in the literature, we assume the agent tie-breaks in favor of the design objective (for an example in the context of revenue see [6]). In our case the objective is the (declared) social welfare. If among the utility-maximizing actions there are several with highest (declared) welfare, we assume the agent chooses to invest maximum effort, i.e., takes the one with the maximum cost. (While we feel this kind of “eager” tie-breaking is in line with the standard tie-breaking assumption, our results do not rely on it and would hold for any consistent tie-breaking rule.)

assuming the rationality of the agent). Principal ℓ is *truthful* as a dominant strategy if announcing her true valuation maximizes her expected utility, i.e., $v^\ell \in \arg \max_{b^\ell \in \mathcal{V}^\ell} \mathbb{E}_{o \sim F_{|x^*(b)^\ell}} [v^\ell(o) - t^\ell(b, o)]$ for every $v^\ell \in \mathcal{V}^\ell, b^{-\ell} \in \mathcal{V}^{-\ell}$. We say that contract t is *truthful* if every principal is truthful as a dominant strategy.

Design objectives. Our desiderata from a contract t concern (1) incentives, (2) social welfare, and (3) computational complexity. The first group of design objectives address incentives – principals’ and agent’s willingness to participate, and principals’ willingness to provide private information. For the latter, we focus on truthful contracts, in which every principal’s dominant strategy is to simply report her valuation. To ensure principals’ participation in such contracts, we require that every truthful principal has non-negative expected utility no matter how the others bid – this is known as the *individual rationality (IR)* property for principals. Formally, contract t satisfies IR if $\mathbb{E}_{o \sim F_{|x^*(b^{-\ell}, v^\ell)}} [v^\ell(o) - t^\ell((b^{-\ell}, v^\ell), o)] \geq 0$ for every $\ell \in [n], b^{-\ell} \in \mathcal{V}^{-\ell}, v^\ell \in \mathcal{V}^\ell$. We now turn to the agent’s incentives, for which we need the following central definition:

Definition 1. A contract t satisfies limited liability (LL) if every payment from a principal to the agent is non-negative: $t^\ell(b, o) \geq 0$ for every $\ell \in [n], b \in \mathcal{V}, o \in \mathcal{O}$.

The LL requirement guarantees that the agent never pays out-of-pocket. Under LL, to ensure the agent is guaranteed non-negative utility, we assume there exists an action with zero cost.⁸

The second design objective is social welfare. Given a valuation profile v , the (expected) *welfare* of action a , denoted by $\text{Wel}^a(v)$, is the total expected value less cost, and $a^*(v)$ is a welfare-maximizing action:

$$\text{Wel}^a(v) = \sum_{\ell \in [n]} \mathbb{E}_{o \sim F_{|a}} [v^\ell(o)] - \psi(a); \quad a^*(v) \in \arg \max_{a \in \mathcal{A}} \text{Wel}^a(v).$$

We seek contracts such that for every bid profile b , the agent’s action choice $x^*(b)$ is “declared welfare maximizing”; we denote this by $x^*(b) = a^*(b)$.⁹ Since we focus on truthful contracts, when playing according to their dominant strategies the principals report $b = v$, and the agent’s choice $a^*(v)$ is indeed socially efficient.

The last design objective is computational complexity: given a bid profile b and an outcome o , the payments $t^\ell(b, o)$ should be computable in polynomial time for every $\ell \in [n]$. To summarize our desiderata, we seek truthful, IR and LL contracts, which incentivize the agent to maximize welfare and whose payments are computable in polynomial time.

⁸This assumption is for simplicity; the alternative of requiring IR for the agent would not change our results.

⁹If there are multiple welfare-maximizing actions then by writing $x^*(b) = a^*(b)$ we mean the agent chooses one of them. Since $x^*(b)$ is uniquely defined by the tie-breaking rule (see Footnote 7), if $x^*(b) = a^*(b)$ then $a^*(b)$ is also uniquely defined.

3 The IIVCG Family of Contracts

In this section we introduce IIVCG contracts. After giving motivation in Section 3.1, we define the IIVCG family and show it is (uniquely) truthful and welfare-maximizing in Section 3.2. In Section 3.3 we give the payment format and in Section 3.4 we instantiate this format such that payments capture externalities, à la VCG *auctions*. We show this instantiation is IR but not LL, and in Section 3.5 establish an inherent tradeoff between these two properties.

3.1 Motivation: The Social Inefficiency of First Price Contracts

A *first price* contract is a contract where each principal simply pays her bid, that is, $t^\ell(b, o) = b^\ell(o)$. As discussed in the introduction, this is the most common contract format in the literature. It satisfies the IR and LL properties. However, as we next show, it suffers from inefficiency in terms of social welfare. Since first price contracts are not truthful, we address their inefficiency through the measures of price of anarchy and price of stability.¹⁰

Theorem 1. *The price of anarchy of first price contracts tends to zero when the number of principals tends to infinity. The price of stability of first price contracts tends to $\frac{1}{2}$ when the number of actions tends to infinity.*¹¹

The proof relies on the following examples.

Example 1 (Price of anarchy). *There are three actions a_1, a_2, a_3 and three outcomes o_1, o_2, o_3 . Action a_i deterministically yields outcome o_i . The costs are: $\psi(a_1) = 0, \psi(a_2) = \epsilon, \psi(a_3) = \gamma$, where $0 < \epsilon < \gamma < 1$. There are $n > 3$ principals with the following valuation domains: $\mathcal{V}^1 = \{(0, 0, v) \mid v \geq 0\}$, $\mathcal{V}^2 = \{(0, v, 0) \mid v \geq 0\}$, and $\mathcal{V}^\ell = \{(v, 0, 0) \mid v \geq 0\} \forall \ell > 2$. The valuations are: $v^1 = (0, 0, 1 + \gamma)$, $v^2 = (0, 1 + \epsilon, 0)$, and $v^\ell = (1, 0, 0) \forall \ell > 2$.*

In this example, the bid profile $b^1 = (0, 0, 1 + \gamma)$, $b^2 = (0, 1 + \epsilon, 0)$, and $b^\ell = (0, 0, 0) \forall \ell > 2$ induces the agent to take action $x^*(b) = a_3$, and forms an equilibrium of the game. Note that $\text{Wel}^{a_3}(v) = 1$, while $\text{Wel}^{a_1}(v) = n - 2$. Hence, the price of anarchy tends to zero when $n \rightarrow \infty$. The full analysis appears in Appendix A.1.

Example 2 (Price of stability). *There are $q \geq 2$ actions a_1, \dots, a_q and two outcomes o_1, o_2 . Action a_i yields outcome o_2 with probability $p_{o \sim F|a_i}(o_2) = \gamma^{q-i}$ where $0 < \gamma < \frac{1}{q}$. The cost for action a_i is $\psi(a_i) = \gamma^{1-i} - i + (i - 1)(\gamma + \epsilon)$, where $0 < \epsilon \leq \frac{1}{q} - \gamma$. There is one principal with valuation domain $\mathcal{V}^1 = \mathbb{R}_{\geq q}^2$ and valuation $v^1 = (q, q + \gamma^{1-q})$.*

¹⁰The *price of anarchy* is the ratio between the worst equilibrium welfare and the optimal welfare. The *price of stability* is the ratio between the best equilibrium welfare and the optimal welfare. For details see, e.g., [24, Chapter 17].

¹¹We do not know whether $\frac{1}{2}$ is tight.

In this example, $\text{Wel}^{a_i}(v) = q + i - (i - 1)(\gamma + \epsilon)$. Thus, the socially efficient action is a_q , while the action that minimizes welfare is a_1 . In Appendix A.2 we show that, when using a first price contract, the agent takes a_1 in all equilibria. The intuitive reason for this is that the principal bids to maximize her own utility rather than social welfare. Hence, the price of stability is $\text{Wel}^{a_1}(v)/\text{Wel}^{a_q}(v)$ which is at most $(q + 1)/(2q - 1)$. The last term tends to $\frac{1}{2}$ when $q \rightarrow \infty$. While the starting point of Example 2 is the “gap setting” construction of [10, p. 15], we must modify the construction in a nontrivial way and provide a new analysis (the gap in that construction is between the optimal welfare and revenue, whereas we seek a gap between the optimal welfare and equilibrium welfare).

3.2 Definition, Properties and Uniqueness of the IIVCG Family

As our main conceptual contribution, we propose to overcome the inefficiency of first price contracts by introducing the following family of contracts, which we show is (uniquely) truthful and welfare-maximizing.

Definition 2. *A contract $t = (t^1, \dots, t^n)$ is called Incomplete Information VCG (IIVCG) if the following two properties hold:*

1. **DECLARED WELFARE MAXIMIZATION.** *Given contract t , for every bid profile $b \in \mathcal{V}$, the agent chooses to maximize the declared social welfare, that is, $x^*(b) = a^*(b)$.*
2. **EXPECTED PAYMENT CHARACTERIZATION.** *There exist functions h^1, \dots, h^n where $h^\ell : \mathcal{V}^{-\ell} \rightarrow \mathbb{R}$, such that for every bid profile $b \in \mathcal{V}$, principal ℓ 's expected payment $\mathbb{E}_{o \sim F_{|x^*(b)}}[t^\ell(b, o)]$ for action $x^*(b) = a^*(b)$ is equal to $h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0})$.*

The first property of IIVCG contracts requires that the agent maximizes the *declared* social welfare. Therefore, when principals bid truthfully, the agent takes a welfare-maximizing action, and social efficiency is achieved. The second property characterizes the expected payment from each principal to the agent for his action $a^*(b)$. The characterization for principal ℓ uses h^ℓ , a function of the bids of the other principals besides ℓ , and $\text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0})$, the part of the welfare from action $a^*(b)$ that is due to the other principals.¹²

The characterization in Definition 2 determines only the payments in expectation (taken over the random outcome) for a certain action, not the payment rules themselves. We further develop the payment rules in Section 3.3. But Definition 2 suffices to establish truthfulness, and thus welfare maximization, for every IIVCG contract. We also show the other direction, i.e., a sense in which this family of contracts is uniquely truthful and welfare-maximizing.

Proposition 1. *Every IIVCG contract is truthful.*

¹²Technically, $\text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0})$ is the welfare from $a^*(b)$ had principal ℓ 's valuation been a zero vector, i.e., the welfare of all principals but ℓ .

Proof. Let t be an IIVCG contract. Consider principal ℓ and fix the other principals' bids $b^{-\ell}$. We show that reporting truthfully maximizes principal ℓ 's expected utility. The key observation is that by Property 2 of Definition 2 of IIVCG contracts, principal ℓ 's expected utility $\mathbb{E}_{o \sim F_{|a^*(b)}}[v^\ell(o) - t^\ell(b, o)]$ for reporting b^ℓ is equal to $\text{Wel}^{a^*(b)}(b^{-\ell}, v^\ell) - h^\ell(b^{-\ell})$. The first term is maximized when $b^\ell = v^\ell$, because it is $\text{Wel}^a(b^{-\ell}, v^\ell)$ when $a = a^*(b^{-\ell}, b^\ell)$, and due to the fact that $\max_a \text{Wel}^a(b^{-\ell}, v^\ell) = \text{Wel}^{a^*(b^{-\ell}, v^\ell)}(b^{-\ell}, v^\ell)$. As for the second term, $h^\ell(b^{-\ell})$ does not depend on principal ℓ 's report b^ℓ . We conclude that the total expected utility is maximized when $b^\ell = v^\ell$. Therefore, reporting truthfully is a dominant strategy for principal ℓ . \square

The next corollary follows immediately from Proposition 1 and from Property 1 of IIVCG contracts.

Corollary 1. *In every IIVCG contract the agent maximizes social welfare.*

We now show a uniqueness result for IIVCG: Whenever the agent maximizes declared welfare, IIVCG contracts are the only contracts that incentivize principals to bid truthfully as a dominant strategy. This result is reminiscent of the well-known result by Holmström [17], establishing uniqueness of VCG *auctions* for convex domains.¹³ In fact, our proof makes use of a lemma from [17], in addition to the convexity of the valuation domains. We give here a sketch of the proof and defer the full proof to Appendix B.1.

Theorem 2 (Uniqueness). *A truthful contract t that satisfies $x^*(b) = a^*(b) \forall b \in \mathcal{V}$ is an IIVCG contract.*

Proof sketch. To prove t is IIVCG we must show existence of functions $\{h^\ell\}$ that satisfy Property 2 in Definition 2 of IIVCG contracts. We first define $h^\ell : \mathcal{V} \rightarrow \mathbb{R}$, i.e., as a function of the *full* bid profile. The definition is as follows: $h^\ell(b) = \mathbb{E}_{o \sim F_{|a^*(b)}}[t^\ell(b, o)] + \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0}) \forall b \in \mathcal{V}$. In order to establish that t is an IIVCG contract, it suffices to show that $h^\ell(b^{-\ell}, \cdot)$ is constant on $b^\ell \forall b^{-\ell} \in \mathcal{V}^{-\ell}$. To do so, we fix $b^{-\ell}$ and apply a key lemma from [17]. Roughly, this lemma states that if two functions $f : \mathcal{V}^\ell \times \mathcal{V}^\ell \rightarrow \mathbb{R}$, $g : \mathcal{V}^\ell \rightarrow \mathbb{R}$ satisfy: (i) $v^\ell \in \arg \max_{b^\ell} f(b^\ell, v^\ell)$, and (ii) $v^\ell \in \arg \max_{b^\ell} f(b^\ell, v^\ell) + g(b^\ell) \forall b^\ell, v^\ell \in \mathcal{V}^\ell$, then $g(\cdot)$ is constant on b^ℓ . Take $f(b^\ell, v^\ell) = \text{Wel}^{a^*(b)}(b^{-\ell}, v^\ell) \forall b^\ell, v^\ell \in \mathcal{V}^\ell$. It is immediate that (i) holds. Take $g(b^\ell) = -h^\ell(b) \forall b^\ell \in \mathcal{V}^\ell$. We show that (ii) also holds. Since principal ℓ is truthful as a dominant strategy, $v^\ell \in \arg \max_{b^\ell \in \mathcal{V}^\ell} \mathbb{E}_{o \sim F_{|a^*(b)}}[v^\ell(o) - t^\ell(b, o)]$. That is, $v^\ell \in \arg \max_{b^\ell \in \mathcal{V}^\ell} \text{Wel}^{a^*(b)}(b^{-\ell}, v^\ell) - h^\ell(b)$. Thus, according to the key lemma, $h^\ell(b^{-\ell}, \cdot)$ is constant on b^ℓ . \square

3.3 A Characterization of the Outcome-Specific Payment Scheme

The definition of IIVCG contracts in the previous section includes a characterization of expected payments for action $a^*(b)$. In this section we complete the picture by characterizing the outcome-

¹³Holmström's result is a generalization of the one by Green and Laffont [13], and in fact applies more generally to smoothly connected domains.

specific payments for all actions (Lemma 1). The characterization relies on Property 1 of IIVCG contracts – that the agent always maximizes the declared social welfare. Since the agent acts in his own self-interest, the payments must *incentivize* the agent to choose the welfare-maximizing action. We use this fact to deduce the format of the payments.

The characterization makes use of the following definition. For every action a define

$$\mathcal{L}_a = \left\{ w \in \mathbb{R}_{\geq 0}^m \mid a \in \arg \max_{a'} \{ \mathbb{E}_{o \sim F|_{a'}} [w(o)] - \psi(a') \} \right\},$$

where $w(o)$ for $o \in [m]$ is the o th entry of vector w . \mathcal{L}_a can be interpreted as follows: It is the set of all non-negative payment vectors w (with one payment per outcome) such that action a maximizes the agent’s expected utility given w . Indeed, if the principals were cooperating rather than acting each in their own self-interest, to incentivize the agent to take action a they could have jointly paid the agent $w(o)$ for every outcome o .

The following lemma uses \mathcal{L}_a to characterize the payments:

Lemma 1 (Outcome-specific payment characterization). *A contract t is an IIVCG contract if and only if there exist functions h^1, \dots, h^n where $h^\ell : \mathcal{V}^{-\ell} \rightarrow \mathbb{R}$, and for every $b \in \mathcal{V}$ there exist vectors $w^1, \dots, w^n \in \mathbb{R}_{\geq 0}^m$ where $\sum_{\ell \in [n]} w^\ell \in \mathcal{L}_{a^*(b)}$, such that*

$$t^\ell(b, o) = h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^{-\ell}, w^\ell) + w^\ell(o) \quad \forall \ell \in [n], o \in \mathcal{O}.$$

Note that the payment format is parameterized by functions $\{h^\ell\}$ and vectors $\{w^\ell\}$, which must still be specified to get a concrete instance of an IIVCG contract, as we do in Section 3.4. The functions $\{h^\ell\}$ in Lemma 1 are the same functions $\{h^\ell\}$ in Property 2 of IIVCG contracts (Definition 2). The vectors $\{w^\ell\}$ are specific to a bid profile b . Note that for every action $a^*(b)$ that maximizes the declared welfare, $\mathcal{L}_{a^*(b)}$ is nonempty.

The proof of Lemma 1 appears in Appendix B.2. We outline here one of the directions, i.e., that the conditions in Lemma 1 are sufficient for contract t to be IIVCG: Consider for simplicity a bid profile b for which there is a single action $a^*(b)$ that maximizes the declared welfare. The agent’s chosen action $x^*(b)$ maximizes his expected utility, i.e., belongs to $\arg \max_{a' \in \mathcal{A}} \sum_{\ell \in [n]} \mathbb{E}_{o \sim F|_{a'}} [t^\ell(b, o)] - \psi(a')$. Since the only term in $t^\ell(b, o)$ that depends on the agent’s choice of action is $w^\ell(o)$, and since summing up this term over the principals gives vector $\sum_{\ell \in [n]} w^\ell \in \mathcal{L}_{a^*(b)}$, it follows that $x^*(b) = a^*(b)$.¹⁴ Thus Property 1 of IIVCG holds. Property 2 can also be verified, by observing that the expected payment is $\mathbb{E}_{o \sim F|_{a^*(b)}} [t^\ell(b, o)] = h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0}) \quad \forall b \in \mathcal{V}$. See Appendix B.2 for details.

¹⁴Recall that the agent tie-breaks in favor of the declared welfare-maximizing action $a^*(b)$.

3.4 IIVCG Instantiation: Auction-Inspired IIVCG

The payments developed in the previous section (Lemma 1) are parameterized by functions $\{h^\ell\}$ and bid-dependent vectors $\{w^\ell\}$. In this section we give a first instantiation of the IIVCG family of contracts by specifying these functions and vectors. The instantiation is closely related to the VCG contracts of [22] as well as to VCG *auctions*, so we refer to the resulting contracts as “Auction-Inspired IIVCG”.

Definition 3. *Contract t is an Auction-Inspired IIVCG contract if*

$$t^\ell(b, o) = \text{Wel}^{a^*(b^{-\ell}, o)}(b^{-\ell}, \mathbf{0}) - \text{Wel}^{a^*(b)}(b) + b^\ell(o) \quad \forall \ell \in [n], b \in \mathcal{V}, o \in \mathcal{O}.$$

Observation 1. *Every Auction-Inspired IIVCG contract belongs to the IIVCG family of contracts.*

Proof. Follows from Lemma 1 by setting $h^\ell(b^{-\ell}) = \text{Wel}^{a^*(b^{-\ell}, \mathbf{0})}(b^{-\ell}, \mathbf{0})$, and setting $w^\ell = b^\ell$ for bid profile b . \square

Since the IIVCG family is guaranteed to be truthful (Proposition 1) and welfare-maximizing (Corollary 1), our main concern in what follows is with the properties of IR and LL.

It is not hard to see that in Auction-Inspired IIVCG contracts, the expected payment of principal ℓ given bid profile b is her *externality* on the other principals: the difference between their (declared) expected welfare had she not participated, i.e., $\text{Wel}^{a^*(b^{-\ell}, \mathbf{0})}(b^{-\ell}, \mathbf{0})$, and their (declared) expected welfare with her participation, i.e., $\text{Wel}^{a^*(b)}(\mathbf{0}, b^\ell)$. This is exactly as in VCG auctions with the Clarke pivot rule. Interestingly, we now show that this natural instantiation of IIVCG does *not* satisfy all our desiderata – while it guarantees IR, it does not necessarily guarantee LL. In Section 3.5 we show this is no coincidence – there is an inherent tradeoff among IR and LL.

Observation 2. *Every Auction-Inspired IIVCG contract is IR.*

Proof Sketch. It can be verified that principal ℓ 's expected utility when she truthfully bids $b^\ell = v^\ell$ and when others bid $b^{-\ell}$ is $\max_{a \in \mathcal{A}} \text{Wel}^a(b^{-\ell}, v^\ell) - \max_{a \in \mathcal{A}} \text{Wel}^a(b^{-\ell}, \mathbf{0})$, which is clearly non-negative. \square

Proposition 2. *There exist common agency settings for which Auction-Inspired IIVCG contracts do not satisfy LL.*

Proof Sketch. Consider a common agency setting with a valuation profile such that no principal has a positive externality on the others. Assume the principals are truthful (LL should hold for any bid profile, including truthful ones). As mentioned above, in any Auction-Inspired IIVCG, principal ℓ 's expected payment is her externality. The principals' expected payments are therefore all zero. The LL requirement that the agent is never paid a negative amount then means that he cannot be paid a positive amount either. When all payments are zero, the agent will not take any action with nonzero cost, in contradiction to the requirement that the agent maximizes welfare. \square

A specific setting that illustrates Proposition 2 appears in the next section (proof of Theorem 3).

3.5 The LL-IR Tradeoff

As we have seen, a natural IIVCG instantiation (namely Auction-Inspired IIVCG) is IR, but does not guarantee LL. In Section 5 we show a different instantiation that guarantees LL but not IR. We now show that the clash between LL and IR is unavoidable (see Appendix B.3 for the full proof):

Theorem 3. *There exist common agency settings for which no IIVCG contract satisfies both LL and IR, even with a single principal.*

Proof Sketch. Consider the following common agency setting: There are two actions a_1, a_2 and two possible outcomes o_1, o_2 . The distributions imposed by the actions are $p_{o \sim F|a_1}(o_1) = \frac{1}{2}$, $p_{o \sim F|a_1}(o_2) = \frac{1}{2}$, $p_{o \sim F|a_2}(o_1) = 0$, $p_{o \sim F|a_2}(o_2) = 1$. The costs for the actions are $\psi(a_1) = 0$ and $\psi(a_2) = \epsilon$. There is one principal with the following valuation domain $\mathcal{V}^1 = \mathbb{R}_{\geq 0}^2$.

Suppose towards a contradiction that there exists an IIVCG contract t for this setting which satisfies IR and LL. For every valuation profile v , contract t is truthful and incentivizes the agent to take the socially efficient action (Proposition 1 and Corollary 1). Consider the valuation profile $v^1 = (0, \epsilon')$, where $\frac{\epsilon'}{2} > \epsilon$. Note that $\text{Wel}^{a_1}(v) = \frac{\epsilon'}{2}$ and $\text{Wel}^{a_2}(v) = \epsilon' - \epsilon$. Thus, the socially efficient action is a_2 . Since the agent maximizes welfare, his expected payoff from a_2 is at least his expected payoff from action a_1 . That is,

$$\begin{aligned} t^1(v, o_2) - \psi(a_2) &\geq \frac{1}{2} \cdot t^1(v, o_1) + \frac{1}{2} \cdot t^1(v, o_2) - \psi(a_1) \iff \\ \frac{1}{2} \cdot t^1(v, o_2) &\geq \frac{1}{2} \cdot t^1(v, o_1) + \epsilon. \end{aligned}$$

This implies that

$$t^1(v, o_2) > t^1(v, o_1) \geq 0. \tag{1}$$

Since t also satisfies IR, the principal's expected value $\mathbb{E}_{o \sim F|a_2}[v^1(o)] = \epsilon'$ from her truthful bid is at least her expected payment $\mathbb{E}_{o \sim F|a_2}[t^1(v, o)]$, which by (1) is strictly positive. By Property 2 of IIVCG contracts, $\mathbb{E}_{o \sim F|a_2}[t^1(v, o)] = h^1(v^{-1}) - \text{Wel}^{a^*(v)}(v^{-1}, \mathbf{0})$. Putting everything together we get:

$$\epsilon' \geq h^1(v^{-1}) - \text{Wel}^{a_2}(v^{-1}, \mathbf{0}) > 0. \tag{2}$$

Recall that $\text{Wel}^{a_2}(v^{-1}, \mathbf{0}) = -\epsilon$, thus $\epsilon' \geq h^1(v^{-1}) + \epsilon > 0$. However, $h^1(\cdot)$ is independent of v^1 and so cannot depend on ϵ' . This is a contradiction. \square

Remark 4. *With two or more principals, Theorem 3 is related to the classic result of Myerson and Satterthwaite, by which there do not exist socially-efficient and budget-balanced mechanisms [23]. Indeed, the agent in our setting can be viewed as the mechanism designer in the classic result, and the principals in our setting can be viewed as the players. However, the impossibility result*

of Myerson and Satterthwaite crucially relies on the existence of two players, whereas Theorem 3 applies even to settings with a single principal.

4 Algorithmic IIVCG

Motivated by the LL-IR tradeoff (Theorem 3), in this section we formulate a necessary and sufficient condition on a common agency setting for the existence of an IIVCG contract that is both LL and IR (Lemma 2). Our characterization is constructive, yielding an IIVCG contract that satisfies both properties whenever such a contract exists. There is a conceptual difference between this contract design and Auction-Inspired IIVCG (presented in Section 3.4) or Weighted IIVCG (presented in Section 5). In the latter classes, the contract mapping $t : \mathcal{V} \times \mathcal{O} \rightarrow \mathbb{R}^n$ is expressed using a short mathematical formula. We introduce *Algorithmic IIVCG* contracts, in which t is algorithm-based:

Definition 4 (Contract ALG1). *For a given common agency setting, $\text{ALG1} : \mathcal{V} \times \mathcal{O} \rightarrow \mathbb{R}^n$ is the contract that for every bid profile $b \in \mathcal{V}$ and outcome $o \in \mathcal{O}$, maps (b, o) to Algorithm 1's output on b and o for this setting.*

The following two theorems constitute our main results:

Theorem 5. *For every common agency setting for which an IIVCG contract that satisfies LL and IR exists, contract ALG1 is such a contract. Furthermore, for every bid profile $b \in \mathcal{V}$ and outcome $o \in \mathcal{O}$, Algorithm 1 computes the payment profile $\text{ALG1}(b, o)$ in polynomial time.*

To use contract ALG1, one must determine for a given setting whether there exists an IIVCG contract with both properties. Algorithm 2 leverages our necessary and sufficient condition and provides a procedure to do so. Formally:

Theorem 6. *Given a common agency setting, Algorithm 2 determines in polynomial time whether there exists an IIVCG contract that is both LL and IR for this setting.*

Proof overview. The proofs of both theorems are given in Sections 4.1 and 4.2, respectively. At a high-level they proceed as follows. Given a common agency setting, we present two bounds on principals' payments:

1. $m^\ell(b) \forall b \in \mathcal{V}$, is an upper bound on principal ℓ 's expected payment in any IIVCG contract t that satisfies IR, i.e., $\mathbb{E}_{o \sim F_{|a^*(b)}}[t^\ell(b, o)] \leq m^\ell(b) \forall \ell \in [n], b \in \mathcal{V}$ (Claim 1).
2. $k_{a^*(b)} \forall b \in \mathcal{V}$, is a lower bound on the sum of principals' expected payments in any IIVCG contract t that satisfies LL, i.e., $k_{a^*(b)} \leq \sum_{\ell \in [n]} \mathbb{E}_{o \sim F_{|a^*(b)}}[t^\ell(b, o)] \forall b \in \mathcal{V}$ (Claim 2).

It follows that if there exists an IIVCG contract that is both LL and IR for the given setting, then

$$k_{a^*(b)} \leq \sum_{\ell \in [n]} m^\ell(b) \quad \forall b \in \mathcal{V}. \quad (3)$$

Surprisingly, the reverse direction is also true, i.e., if the predicate in Eq. (3) holds, one can construct an IIVCG contract that is both LL and IR (Claim 3 and Algorithm 1). Section 4.2 shows that this predicate can be computed in polynomial time.

4.1 Proof of Theorem 5

In this section we explain Algorithm 1 and prove Theorem 5 step by step. As a first step, we provide the following upper bound on principals' payments.

Claim 1. *Consider a common agency setting. Define:*

$$m^\ell(b) = \min_{\tilde{v}^\ell \in \mathcal{V}^\ell} \text{Wel}^{a^*(b^{-\ell}, \tilde{v}^\ell)}(b^{-\ell}, \tilde{v}^\ell) - \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0}) \quad \forall \ell \in [n], b \in \mathcal{V}.$$

An IIVCG-contract t satisfies IR if and only if $\mathbb{E}_{o \sim F_{|a^(b)}}[t^\ell(b, o)] \leq m^\ell(b) \quad \forall \ell \in [n], b \in \mathcal{V}$.*

Proof. By definition, t satisfies IR if and only if each principal's expected payment when she bids truthfully is at most her expected value. That is, $\mathbb{E}_{o \sim F_{|a^*(b^{-\ell}, v^\ell)}}[t^\ell((b^{-\ell}, v^\ell), o)] \leq \mathbb{E}_{o \sim F_{|a^*(b^{-\ell}, v^\ell)}}[v^\ell(o)] \quad \forall \ell \in [n], b \in \mathcal{V}, v^\ell \in \mathcal{V}^\ell$. By Property 2 in the definition of IIVCG (Definition 2), the above holds if and only if $h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b^{-\ell}, v^\ell)}(b^{-\ell}, \mathbf{0}) \leq \mathbb{E}_{o \sim F_{|a^*(b^{-\ell}, v^\ell)}}[v^\ell(o)] \quad \forall \ell \in [n], b \in \mathcal{V}, v^\ell \in \mathcal{V}^\ell$. By adding $\text{Wel}^{a^*(b^{-\ell}, v^\ell)}(b^{-\ell}, \mathbf{0}) - \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0})$ to both sides of the inequality, we have that t satisfies IR if and only if $h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0}) \leq \text{Wel}^{a^*(b^{-\ell}, v^\ell)}(b^{-\ell}, v^\ell) - \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0}) \quad \forall \ell \in [n], b \in \mathcal{V}, v^\ell \in \mathcal{V}^\ell$. Since the left-hand side is independent of v^ℓ , the above is equivalent to the following condition (obtained by minimizing the right-hand side over all $v^\ell \in \mathcal{V}^\ell$ and dropping the $\forall v^\ell \in \mathcal{V}^\ell$ quantifier): $h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0}) \leq \min_{\tilde{v}^\ell \in \mathcal{V}^\ell} \text{Wel}^{a^*(b^{-\ell}, \tilde{v}^\ell)}(b^{-\ell}, \tilde{v}^\ell) - \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0}) \quad \forall \ell \in [n], b \in \mathcal{V}$. That is, by Property 2 in the definition of IIVCG, and by the definition of $m^\ell(b)$, if and only if $\mathbb{E}_{o \sim F_{|a^*(b)}}[t^\ell(b, o)] \leq m^\ell(b) \quad \forall \ell \in [n], b \in \mathcal{V}$. \square

Proposition 3. *$m^\ell(b)$ can be computed in polynomial time by solving the following linear program and subtracting $F_{|a^*(b)} \cdot b^{-\ell} - \psi(a^*(b))$ from its solution.*

$$\begin{aligned} & \underset{\tilde{v}^\ell \in \mathcal{V}^\ell, h^\ell \in \mathbb{R}}{\text{minimize}} \quad h^\ell & & \text{(LP1)} \\ & \text{subject to} \quad F_{|a_j} \cdot (b^{-\ell} + \tilde{v}^\ell) - \psi(a_j) \leq h^\ell & & \forall j \in [q], \end{aligned}$$

where $b^{-\ell}(o) = \sum_{\ell' \neq \ell \in [n]} b^{\ell'}(o) \quad \forall o \in \mathcal{O}$.

Proof. The linear program above is solvable in polynomial time since it consists of $O(m)$ variables: $0 \leq h^\ell, \tilde{v}^\ell \in \mathcal{V}^\ell \subseteq \mathbb{R}^m$; and polynomially-many constraints: $O(q)$ constraints in addition to the polynomially-many constraints ensuring that $v \in \mathcal{V}$ (recall that \mathcal{V}^ℓ is a convex, polynomially-described polytope for every ℓ). Note that $\text{Wel}^{a_j}(b^{-\ell}, \tilde{v}^\ell) = F_{|a_j} \cdot (b^{-\ell} + \tilde{v}^\ell) - \psi(a_j)$. Thus, the j th constraint ensures that the welfare from action a_j given $(b^{-\ell}, \tilde{v}^\ell)$ is at most h^ℓ . Altogether

we have that the optimal welfare $\text{Wel}^{a^*(b^{-\ell}, \tilde{v}^\ell)}(b^{-\ell}, \tilde{v}^\ell) \leq h^\ell$ at the feasible region. Minimizing h^ℓ over all $\tilde{v}^\ell \in \mathcal{V}^\ell$ gives the first term in $m^\ell(b)$. To complete $m^\ell(b)$'s computation, subtract $\text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0}) = F_{|_{a^*(b)}} \cdot b^{-\ell} - \psi(a^*(b))$ from h^ℓ . \square

Claim 2. Consider a common agency setting. Define: $k_a = \min_{w \in \mathcal{L}_a} \mathbb{E}_{o \sim F_{|_a}}[w(o)] \forall a$. If an IIVCG-contract t satisfies LL, then $k_{a^*(b)} \leq \sum_{\ell \in [n]} \mathbb{E}_{o \sim F_{|_{a^*(b)}}}[t^\ell(b, o)] \forall b \in \mathcal{V}$.

Proof. By linearity of expectation it suffices to show that $(\sum_{\ell \in [n]} t^\ell(b, o))_{o \in \mathcal{O}} \in \mathcal{L}_{a^*(b)} \forall b \in \mathcal{V}$. By definition of $\mathcal{L}_{a^*(b)}$, we need to show that: (i) $\sum_{\ell \in [n]} t^\ell(b, o) \geq 0 \forall o \in \mathcal{O}$. This follows from LL. (ii) $a^*(b) \in \arg \max_{a'} \{\mathbb{E}_{o \sim F_{|_{a'}}}[\sum_{\ell \in [n]} t^\ell(b, o)] - \psi(a')\}$. This follows from the definition of IIVCG, where the agent takes $a^*(b)$. \square

Proposition 4. $k_{a^*(b)}$ can be computed in polynomial time by solving the following linear program.

$$\begin{aligned} & \underset{w \in \mathbb{R}^m}{\text{minimize}} && F_{a^*(b)} \cdot w && \text{(LP2)} \\ & \text{subject to} && F_{|_{a_j}} \cdot w - \psi(a_j) \leq F_{|_{a^*(b)}} \cdot w - \psi(a^*(b)) && \forall j \in [q], \\ & && w_k \geq 0 && \forall k \in [m]. \end{aligned}$$

Proof. The linear program above is solvable in polynomial time since it consists of $O(m)$ variables: $w \in \mathbb{R}^m$, and $O(q + m)$ constraints. The first set of constraints ensures that $a^*(b) \in \arg \max_{a'} \{\mathbb{E}_{o \sim F_{|_{a'}}}[w(o)] - \psi(a')\}$. The second set of constraints ensures that $w \in \mathbb{R}_{\geq 0}^m$. Altogether, $w \in \mathcal{L}_{a^*(b)}$ at the feasible region. Minimizing $F_{a^*(b)} \cdot w = \mathbb{E}_{o \sim F_{|_{a^*(b)}}}[w(o)]$ computes $k_{a^*(b)}$. \square

Lemma 2 (Necessary and sufficient condition). Consider a common agency setting. There exists an IIVCG contract that satisfies LL and IR if and only if $k_{a^*(b)} \leq \sum_{\ell \in [n]} m^\ell(b) \forall b \in \mathcal{V}$.

Proof. For the forward direction, suppose that there exists an IIVCG-contract t that satisfies both LL and IR. By claims 2, 1, we conclude that $k_{a^*(b)} \leq \sum_{\ell \in [n]} \mathbb{E}_{o \sim F_{|_{a^*(b)}}}[t^\ell(b, o)] \leq \sum_{\ell \in [n]} m^\ell(b) \forall b \in \mathcal{V}$. The backwards direction is constructive and relies on the following claim. \square

Claim 3. If a setting satisfies $k_{a^*(b)} \leq \sum_{\ell \in [n]} m^\ell(b) \forall b \in \mathcal{V}$, an IIVCG contract that satisfies LL and IR may be constructed by using Lemma 1, and

1. $h^\ell(b^{-\ell}) = \min_{\tilde{v}^\ell \in \mathcal{V}^\ell} \text{Wel}^{a^*(b^{-\ell}, \tilde{v}^\ell)}(b^{-\ell}, \tilde{v}^\ell) \forall \ell \in [n], b^{-\ell} \in \mathcal{V}^{-\ell}$.
2. For every $b \in \mathcal{V}$, take $w^b \in \arg \min_{w \in \mathcal{L}_{a^*(b)}} \mathbb{E}_{o \sim F_{|_{a^*(b)}}}[w(o)]$ and choose w^1, \dots, w^n as follows: $w^\ell = x_\ell^b \cdot w^b$ where $0 \leq x_\ell^b, x_\ell^b \cdot k_{a^*(b)} \leq m^\ell(b) \forall \ell \in [n]$, and $\sum_{\ell \in [n]} x_\ell^b = 1$.

Proof of Claim 3. We first show that this contract is in IIVCG, by proving that Lemma 1 can be applied. Note that $\sum_{\ell \in [n]} w^\ell = w^b \in \mathcal{L}_{a^*(b)} \forall b \in \mathcal{V}$, and that h^ℓ is indeed independent of b^ℓ . Thus, the constructed contract is an IIVCG contract. To show LL, we must show that

$t^\ell(b, o) = h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^{-\ell}, w^\ell) + w^\ell(o) \geq 0 \forall \ell \in [n], b \in \mathcal{V}, o \in \mathcal{O}$. Since $w^\ell(o)$ is non-negative, the last inequality holds if $h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^{-\ell}, w^\ell) \geq 0 \forall \ell \in [n], b \in \mathcal{V}$. By adding $\mathbb{E}_{o \sim F_{|a^*(b)}}[w^\ell(o)]$ to both sides of the inequality, we have that the above holds if $h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0}) \geq \mathbb{E}_{o \sim F_{|a^*(b)}}[w^\ell(o)] \forall \ell \in [n], b \in \mathcal{V}$. By our choice of $h^\ell(b^{-\ell}) = \min_{\tilde{v}^\ell \in \mathcal{V}^\ell} \text{Wel}^{a^*(b^{-\ell}, \tilde{v}^\ell)}(b^{-\ell}, \tilde{v}^\ell)$ and of $w^\ell = x_\ell^b \cdot \arg \min_{w \in \mathcal{L}_{a^*(b)}} \mathbb{E}_{o \sim F_{|a^*(b)}}[w(o)]$, this is equivalent to $m^\ell(b) \geq x_\ell^b \cdot k_{a^*(b)} \forall \ell \in [n], b \in \mathcal{V}$, which holds by construction. IR follows from Claim 1, the choice of $h^\ell(b^{-\ell})$, and IIVCG expected payment formula: $\mathbb{E}_{o \sim F_{|a^*(b)}}[t^\ell(b, o)] = h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0})$. \square

Observation 3. *The following linear program computes x from Claim 3 (2) in polynomial time.*

$$\begin{aligned} & \underset{x^b \in \mathbb{R}^n}{\text{maximize}} && \sum_{\ell \in [n]} x_\ell^b && \text{(LP3)} \\ & \text{subject to} && \sum_{\ell \in [n]} x_\ell^b \leq 1, \quad x_\ell^b \cdot k_{a^*(b)} \leq m^\ell, \quad 0 \leq x_\ell^b && \forall \ell \in [n]. \end{aligned}$$

Lemma 3. *If a setting satisfies $k_{a^*(b)} \leq \sum_{\ell \in [n]} m^\ell(b) \forall b \in \mathcal{V}$, Algorithm 1 computes the construction in Claim 3 in polynomial time.*

Proof. The first loop computes $h^\ell(b^{-\ell}) \forall \ell \in [n]$. Line 4 computes w^b , and line 6 computes x . The last loop computes the exact payments according to Lemma 1. All linear programs are solvable in polynomial time as argued above (Propositions 3, 4 and Observation 3). \square

Proof of Theorem 5. According to Lemma 2, Claim 3 and Lemma 3, ALG1 is an IIVCG contract that satisfies LL and IR for every setting where such a contract exists. Truthfulness and social efficiency follow from the definition of IIVCG. By Lemma 3, Algorithm 1 computes the payments of ALG1 in polynomial time. \square

Remark 7. *One may consider relaxing the LL constraint, by requiring that only the aggregate payment from all principals to the agent is non-negative. The proof of Claim 2 shows that requiring the payments from each individual principal to the agent to be non-negative is without loss of generality.*

4.2 Proof of Theorem 6

In this section we prove Theorem 6, providing a method for determining when there exists an IIVCG contract that is both LL and IR. According to the characterization in Lemma 2, there exists such a contract if and only if $k_{a^*(b)} \leq \sum_{\ell \in [n]} m^\ell(b) \forall b \in \mathcal{V}$. A key observation is that in order to test this predicate, it suffices to test only $O(q)$ bid profiles as follows (where q is the number of actions).

Observation 4. *Define $\mathcal{V}_a = \{v \in \mathcal{V} \mid a = a^*(v)\} \forall a$. Then $k_{a^*(b)} \leq \sum_{\ell \in [n]} m^\ell(b) \forall b \in \mathcal{V}$ if and only if $k_a \leq \min_{b \in \mathcal{V}_a} \{\sum_{\ell \in [n]} m^\ell(b)\} \forall a$.*

ALGORITHM 1: Returns the payments of an IIVCG contract that satisfies IR and LL for the setting, or Impossible if no such contract exists.

Input: Bid profile $b \in \mathcal{V}$, outcome $o \in \mathcal{O}$ and an access to the setting $(\mathcal{A}, \mathcal{O}, \mathcal{V}, F, \psi)$.

```

1 for  $\ell \in [n]$  do
2   | Solve LP1 to find  $h^\ell$ ;
3   |  $m^\ell \leftarrow h^\ell - \{F_{|a^*(b)} \cdot b^{-\ell} - \psi(a^*(b))\}$ ;
4   | Solve LP2 to find  $k_{a^*(b)}$  and  $w^b$ ;
5   | if  $k_{a^*(b)} > \sum_{\ell \in [n]} m^\ell$  then return Impossible;
6   | Solve LP3 to find  $x$ .
7   | for  $\ell \in [n]$  do
8   |   |  $t^\ell(b, o) \leftarrow h^\ell - \{F_{|a^*(b)} \cdot (b^{-\ell} + x_\ell^b \cdot w^b) - \psi(a^*(b))\} + x_\ell^b \cdot w^b(o)$ ;
9   | return  $\{t^\ell(b, o)\}_{\ell \in [n]}$ ;

```

Observation 4 follows directly from the definition of \mathcal{V}_a , and will be used in the proof of Theorem 6.

Claim 4. Consider an action a . The following linear program computes $\min_{b \in \mathcal{V}_a} \sum_{\ell \in [n]} m^\ell(b)$ in polynomial time.¹⁵

$$\begin{aligned}
& \underset{b \in \mathcal{V}_a, \tilde{v} \in \mathcal{V}, h \in \mathbb{R}^n}{\text{minimize}} && \sum_{\ell \in [n]} \{h^\ell - (F_{|a} \cdot b^{-\ell} - \psi(a))\} && \text{(LP4)} \\
& \text{subject to} && F_{|a_j} \cdot (b^{-\ell} + \tilde{v}^\ell) - \psi(a_j) \leq h^\ell && \forall j \in [q], \ell \in [n].
\end{aligned}$$

Proof. This linear program is similar to LP1, which computes the first term in $m^\ell(b)$, and is referred to as h^ℓ . We use the same set of constraints for each principal to ensure that $\text{Wel}^{a^*(b^{-\ell}, \tilde{v}^\ell)}(b^{-\ell}, \tilde{v}^\ell) \leq h^\ell \forall \ell \in [n]$ at the feasible region. However, the search space and the objective are different. Instead of computing the first term in $m^\ell(b)$, we subtract $F_{|a} \cdot b^{-\ell} - \psi(a) = F_{|a^*(b)} \cdot b^{-\ell} - \psi(a^*(b))$ and sum over all principals to compute $\sum_{\ell \in [n]} m^\ell(b)$. We can do so since \tilde{v}^ℓ s are independent. Further, the search space is all $b \in \mathcal{V}_a$ on top of all $\tilde{v} \in \mathcal{V}$ to find $\min \sum_{\ell \in [n]} m^\ell(b)$. \square

Proof of Theorem 6. Recall Algorithm 1. Note that upon receiving a bid profile b and an outcome o , in addition to computing the payments of ALG1, it computes $k_{a^*(b)} \leq \sum_{\ell \in [n]} m^\ell(b)$, and returns *Impossible* if this predicate does not hold (see line 5 in Algorithm 1). According to Observation 4, given a common agency setting, it suffices to test if $k_a \leq \min_{b \in \mathcal{V}_a} \sum_{\ell \in [n]} m^\ell(b) \forall a$ to determine whether there exists an IIVCG contract that satisfies both LL and IR. This is done in Algorithm 2. For each action a , line 2 uses LP4 to find $\arg \min_{b \in \mathcal{V}_a} \sum_{\ell \in [n]} m^\ell(b)$. Then, using Algorithm 1, we test if $k_a \leq \sum_{\ell \in [n]} m^\ell(b)$. Thus, Algorithm 2 computes $k_a \leq \min_{b \in \mathcal{V}_a} \sum_{\ell \in [n]} m^\ell(b) \forall a$ and determines whether a “good” IIVCG contract exists. \square

¹⁵To set $b \in \mathcal{V}_a$ as a linear constraint we use the constraints of \mathcal{L}_a (as in Algorithm 1) with weak and strict inequalities according to the tie-breaking in $a^*(\cdot)$. A standard technique for strict inequalities in a linear program is adding a small constant that depends on the bit precision, see e.g., [11].

ALGORITHM 2: Determines if there exists an IIVCG contract that satisfies LL and IR for a given setting.

Input: Setting $(\mathcal{A}, \mathcal{O}, \mathcal{V}, F, \psi)$

```

1 for  $a \in \mathcal{A}$  do
2   | Solve LP4 to find  $\arg \min_{b \in \mathcal{V}_a} \sum_{\ell \in [n]} m^\ell(b)$ ;
3   | if Algorithm 1 on  $b$  and arbitrary  $o$  returns Impossible then
4   |   | return Impossible;
5 return Possible;

```

Remark 8. When the principals have no private information (i.e., their types are fixed and publicly known so V_i is a singleton), there are always contracts in IIVCG that satisfy LL and IR. I.e., our Algorithm 2 will return “Possible” and our Algorithm 1 will find the best such contract for the principals. As a concrete example, one LL and IR IIVCG contract is as follows. Take $h^\ell = \text{Wel}^{a^*(v)}(v)$, and $w^\ell = v^\ell$ in Lemma 1 of the IIVCG characterization. One can verify that a principal’s payments in this contract are exactly her values. LL holds since the principal’s values are non-negative, and IR holds since the principal’s utility is zero. We mention that this contract favors the agent’s utility over that of the principals’. At the other extreme, one can use Algorithmic IIVCG. Our LP approach to this class assures that the minimum payment required to incentivize the agent is charged from the principals, thus favoring their utility over the agent’s.

5 IIVCG Instantiation: Weighted IIVCG Contracts

In this section we introduce a class within IIVCG which guarantees LL but not necessarily IR, and give a condition on common agency settings (“ G -correlation”) that suffices for the existence of IR as well. The high-level idea is as follows. By Proposition 2, in order to satisfy LL, principals’ expected payment should be more than their externality on others. While increasing principals’ payments may damage IR, a type of contract which we term *Weighted IIVCG* succeeds in providing IR for a wide family of common agency settings. The central idea is to approximate each principal’s valuation according to other principals using the concept of a *correlation graph*, and use this to set the payments. Correlation graphs and G -correlated settings are defined in Section 5.1, and in Section 5.2 we use these to define Weighted IIVCG.

5.1 Correlation Graphs and G -Correlated Settings

A correlation graph (Definition 5) describes how principals’ valuations are correlated. (As one interpretation, principals’ valuations can be intuitively thought of as reflecting their socioeconomic status, and the correlation graph as a natural way to describe it.)

Definition 5. A correlation graph G is a weighted directed graph with vertex set $[n]$ and weights on edges $d : [n]^2 \rightarrow [0, 1]$ such that $\sum_{k \in [n]} d(k, \ell) = 1$, and $d(\ell, \ell) = 0 \forall \ell \in [n]$.

We introduce the following “running example” for this section to demonstrate our definitions. Figure 1 depicts a correlation graph of the setting in the following example.

Example 3. *There are 3 principals, 2 actions and 2 outcomes. The valuation domain is $\mathcal{V}^\ell = [10, 15]^2 \forall \ell \in [3]$. The costs are: $\psi(a_1) = 0$, and $\psi(a_2) = 1$. The probabilities over outcomes are: $p_{o \sim F|a_1}(o_1) = 1$, and $p_{o \sim F|a_2}(o_1) = 0.25$, and principals’ valuations are: $v^1 = (11, 13)$, $v^2 = (12, 14)$, and $v^3 = (10, 11)$.*

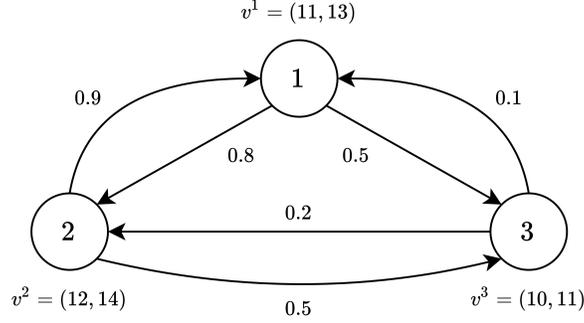


Figure 1: A possible correlation graph of the common agency setting in Example 3.

We mention that a correlation graph may consist of several connected components. (In our intuitive interpretation, think of a state with multiple cities, within each city individuals are socioeconomically correlated.)

Given a correlation graph G , we define principals’ approximate valuations, referred to as their G -weighted valuations, as follows.

Definition 6. *Consider a correlation graph G , and a valuation profile v . The G -weighted valuation of principal ℓ is: $\hat{v}_G^\ell = \sum_{k \in [n]} d(\ell, k) \cdot \hat{v}^k$, where $\hat{v}^k = v^k(o) - \min_{o'} v^k(o')$ $\forall o \in \mathcal{O}, k \in [n]$.*

In words, the G -weighted valuation of principal ℓ is a weighted average on other principals valuations, when their valuations are shifted downwards. We will demonstrate this for Example 3 and the correlation graph G in Figure 1: The G -weighted valuation of principal 1 is $\hat{v}_G^1 = d(1, 2) \cdot \hat{v}^2 + d(1, 3) \cdot \hat{v}^3$. It can be verified that $\hat{v}^2 = (0, 2)$, $\hat{v}^3 = (0, 1)$, and since $d(1, 2) = 0.8$, and $d(1, 3) = 0.5$, the G -weighted valuation of 1 is $\hat{v}_G^1 = (0, 1.6) + (0, 0.5) = (0, 2.1)$.

Definition 7. *Given a correlation graph G , a setting is G -correlated if $\text{Wel}^{a^*(v^{-\ell}, \hat{v}_G^\ell)}(v^{-\ell}, \hat{v}_G^\ell) \leq \text{Wel}^{a^*(v)}(v) \forall \ell \in [n], v \in \mathcal{V}$.*

That is, a setting is G -correlated if the welfare when principal ℓ ’s valuation is replaced with her G -weighted valuation, is bounded by the actual welfare (for every valuation profile). We demonstrate this condition for the setting and the valuation profile in Example 3, and the correlation

graph in Figure 1. Recall that $\hat{v}_G^1 = (0, 2.1)$. It can thus be verified that $\text{Wel}^{a_1}(v) = 33$, $\text{Wel}^{a_2}(v) = 35.75$ and that $\text{Wel}^{a_1}(v^{-\ell}, \hat{v}_G^\ell) = 22$, $\text{Wel}^{a_2}(v^{-\ell}, \hat{v}_G^\ell) = 24.825$. Thus, the welfare when principal ℓ 's valuation is replaced with her G -weighted valuation is bounded by the actual welfare (for v).

Examples of G -correlated settings. The following claims describe natural special cases for G -correlated settings.

Claim 5. *Consider a common agency setting. If principals share the same expected value for each action, there exists a correlation graph G for which the setting is G -correlated.*¹⁶

Proof. Take G as any graph on $[n]$ vertices that has a weighted in-degree and out-degree of 1 for all vertices.¹⁷ Consider a valuation profile v and a principal ℓ . To satisfy the condition for G -correlation it suffices to have $\mathbb{E}_{o \sim F|_a}[\hat{v}_G^\ell(o)] \leq \mathbb{E}_{o \sim F|_a}[v^\ell(o)] \forall a$. Since it implies that $\text{Wel}^{a^*(v^{-\ell}, \hat{v}_G^\ell)}(v^{-\ell}, \hat{v}_G^\ell) \leq \text{Wel}^{a^*(v^{-\ell}, \hat{v}_G^\ell)}(v^{-\ell}, v^\ell) \leq \text{Wel}^{a^*(v)}(v)$. By the definition of \hat{v}^ℓ and of \hat{v}^k , $\mathbb{E}_{o \sim F|_a}[\hat{v}_G^\ell(o)] = \mathbb{E}_{o \sim F|_a}[\sum_{k \in [n]} d(\ell, k) \cdot \hat{v}^k(o)] \leq \mathbb{E}_{o \sim F|_a}[\sum_{k \in [n]} d(\ell, k) \cdot v^k(o)]$. Since principal's share the same expected valuation, $\mathbb{E}_{o \sim F|_a}[\sum_{k \in [n]} d(\ell, k) \cdot v^k(o)] = \sum_{k \in [n]} d(\ell, k) \cdot \mathbb{E}_{o \sim F|_a}[v^\ell(o)]$, and since $\sum_{k \in [n]} d(\ell, k) = 1$, the last term is $\mathbb{E}_{o \sim F|_a}[v^\ell(o)]$. We conclude that $\mathbb{E}_{o \sim F|_a}[\hat{v}_G^\ell(o)] \leq \mathbb{E}_{o \sim F|_a}[v^\ell(o)] \forall a$ as desired. This completes the proof. \square

Claim 6. *Consider a common agency setting. If principals' valuation domain is $[a, b]^m$ where $b - a \leq a$, there exists a correlation graph G for which the setting is G -correlated.*

We mention that in this case, first price contracts may fail.

Proof. The proof proceeds as in Claim 5, but requires a different argument for proving that $\mathbb{E}_{o \sim F|_a}[\hat{v}_G^\ell(o)] \leq \mathbb{E}_{o \sim F|_a}[v^\ell(o)] \forall a$. By definition, $\mathbb{E}_{o \sim F|_a}[\hat{v}_G^\ell(o)] = \mathbb{E}_{o \sim F|_a}[\sum_{k \in [n]} d(\ell, k) \cdot \hat{v}^k(o)]$. Note that since $v^k(o) \leq b$, and $a \leq \min_{o' \in \mathcal{O}} v^k(o')$, we have that $\hat{v}^k(o) \leq b - a \leq a \forall k \in [n], o \in \mathcal{O}$. This implies that, $\mathbb{E}_{o \sim F|_a}[\sum_{k \in [n]} d(\ell, k) \cdot \hat{v}^k(o)] \leq \sum_{k \in [n]} d(\ell, k) \cdot a$, and since $\sum_{k \in [n]} d(\ell, k) = 1$ the last term is a . Recall that principal's valuations are at least a . Thus, $a \leq \mathbb{E}_{o \sim F|_a}[v^\ell(o)]$. We conclude that $\mathbb{E}_{o \sim F|_a}[\hat{v}_G^\ell(o)] \leq \mathbb{E}_{o \sim F|_a}[v^\ell(o)] \forall a$. This completes the proof. \square

5.2 G -Weighted IIVCG Contracts

We can now use the concepts introduced in the previous section to define Weighted IIVCG contracts and analyze their properties.

¹⁶While this is natural, one can show that first price contracts do well in this case: Since principals have the same expected value for each action, and principals' domains are independent, the welfare of each action is fixed. Thus, the welfare maximizing action is unified for all valuation profiles (and bid profiles). In first price contracts, the agent maximizes declared welfare. Thus the socially efficient action is always chosen by the agent.

¹⁷As two simple examples for such graphs, consider the graph defined by $d(\ell, \ell + 1) = 1 \forall \ell \in [n]$, where $\ell + 1$ is calculated modulo n , and the graph is defined by $d(\ell, k) = \frac{1}{n-1} \forall \ell, k \in [n]$.

Definition 8. Consider a correlation graph G . A G -Weighted IIVCG Contract is given by

$$t^\ell(b, o) = \text{Wel}^{a^*(b^{-\ell}, \hat{b}_G^\ell)}(b^{-\ell}, \hat{b}_G^\ell) - \text{Wel}^{a^*(b)}(b^{-\ell}, \hat{b}_G^\ell) + \hat{b}_G^\ell(o) \quad \forall \ell \in [n], b \in \mathcal{V}, o \in \mathcal{O}.$$

Proposition 5. Consider a correlation graph G . Every G -Weighted IIVCG contract is in IIVCG.

Proof. The proof relies on Lemma 1 when $h^\ell(b^{-\ell}) = \text{Wel}^{a^*(b^{-\ell}, \hat{b}_G^\ell)}(b^{-\ell}, \hat{b}_G^\ell)$ and $w^\ell = \hat{b}_G^\ell(o)$. We prove that the conditions of Lemma 1 are met. First, h^ℓ is independent of b^ℓ , since \hat{b}_G^ℓ is independent of b^ℓ ($d(\ell, \ell) = 0 \forall \ell \in [n]$). Second, we show that $\sum_{\ell \in [n]} w^\ell \in \mathcal{L}_{a^*(b)}$. To do so, we establish that $\sum_{\ell \in [n]} w^\ell = \sum_{\ell \in [n]} \hat{b}^\ell$ using the following chain of equalities:

$$\sum_{\ell \in [n]} w^\ell \stackrel{(1)}{=} \sum_{\ell \in [n]} \sum_{k \in [n]} d(\ell, k) \hat{b}^k \stackrel{(2)}{=} \sum_{k \in [n]} \sum_{\ell \in [n]} d(\ell, k) \hat{b}^k \stackrel{(3)}{=} \sum_{k \in [n]} \hat{b}^k.$$

Equality (1) follows from the choice of w^ℓ s above, and the definition of \hat{b}_G^ℓ (Definition 6). Equality (2) changes the order of summands, and Equality (3) follows from the fact that $\sum_{k \in [n]} d(k, \ell) = 1 \forall \ell \in [n]$ (Definition 7). Thus, $\sum_{\ell \in [n]} w^\ell = \sum_{\ell \in [n]} \hat{b}^\ell$, and it suffices to show that $\sum_{\ell \in [n]} \hat{b}^\ell \in \mathcal{L}_{a^*(b)}$. To do so, we show that $\sum_{\ell \in [n]} \hat{b}^\ell \in \mathbb{R}_{\geq 0}^m$ and that $a^*(b) \in \arg \max_{a'} \{\mathbb{E}_{o \sim F_{|a'}}[\sum_{\ell \in [n]} \hat{b}^\ell(o)] - \psi(a')\}$. First, by definition $\hat{b}^\ell(o) \geq 0 \forall o \in \mathcal{O}, \ell \in [n]$. Thus, $\sum_{\ell \in [n]} \hat{b}^\ell \in \mathbb{R}_{\geq 0}^m$. Second, by definition of $a^*(b)$, it holds that $a^*(b) \in \arg \max_{a'} \{\mathbb{E}_{o \sim F_{|a'}}[\sum_{\ell \in [n]} b^\ell(o)] - \psi(a')\}$. Thus $a^*(b) \in \arg \max_{a'} \{\mathbb{E}_{o \sim F_{|a'}}[\sum_{\ell \in [n]} b^\ell(o) - y] - \psi(a')\} \forall y \in \mathbb{R}$, and specifically for $y = \sum_{\ell \in [n]} \min_{o \in \mathcal{O}} b^\ell(o)$. We conclude that $a^*(b) \in \arg \max_{a'} \{\mathbb{E}_{o \sim F_{|a'}}[\sum_{\ell \in [n]} \hat{b}^\ell(o)] - \psi(a')\}$. This completes the proof. \square

Theorem 9. Consider a correlation graph G . Every G -Weighted IIVCG contract satisfies LL. Furthermore, IR is satisfied if and only if the setting is G -correlated.

Proof. Truthfulness and social efficiency are immediate since G -Weighted IIVCG is in IIVCG. LL follows from the fact that the first term in the payment formula, $\text{Wel}^{a^*(b^{-\ell}, \hat{b}_G^\ell)}(b^{-\ell}, \hat{b}_G^\ell)$, is (weakly) higher than the second term, $\text{Wel}^{a^*(b)}(b^{-\ell}, \hat{b}_G^\ell)$, and since $\hat{b}_G^\ell(o) \geq 0 \forall o$. We prove IR if and only if the setting is G -correlated. By definition, IR holds if and only if each principal ℓ 's expected payment when she bids truthfully is at most her expected value. That is, by the definition of G -Weighted IIVCG IR holds, if and only if $\text{Wel}^{a^*(b^{-\ell}, \hat{b}_G^\ell)}(b^{-\ell}, \hat{b}_G^\ell) - \text{Wel}^{a^*(b^{-\ell}, v^\ell)}(b^{-\ell}, \mathbf{0}) \leq \mathbb{E}_{o \sim F_{|a^*(b^{-\ell}, v^\ell)}}[v^\ell(o)] \forall b^{-\ell} \in \mathcal{V}^{-\ell}, v^\ell \in \mathcal{V}^\ell, \ell \in [n]$. By adding $\text{Wel}^{a^*(b^{-\ell}, v^\ell)}(b^{-\ell}, \mathbf{0})$ to both sides of the inequality, we get $\text{Wel}^{a^*(b^{-\ell}, \hat{b}_G^\ell)}(b^{-\ell}, \hat{b}_G^\ell) \leq \text{Wel}^{a^*(b^{-\ell}, v^\ell)}(b^{-\ell}, v^\ell) \forall b^{-\ell} \in \mathcal{V}^{-\ell}, v^\ell \in \mathcal{V}^\ell, \ell \in [n]$. Denoting $b = (b^{-\ell}, v^\ell)$, this condition is equivalent to Definition 7 of the G -correlated setting. This completes the proof. \square

References

- [1] Tal Alon, Magdalen Dobson, Ariel D. Procaccia, Inbal Talgam-Cohen, and Jamie Tucker-Foltz. 2020. Multiagent Evaluation Mechanisms. In *The Thirty-Fourth AAAI Conference on*

- Artificial Intelligence, AAAI 2020*. 1774–1781.
- [2] Moshe Babaioff, Michal Feldman, and Noam Nisan. 2006. Combinatorial agency. In *Proceedings of the 7th ACM conference on Electronic commerce*. 18–28.
 - [3] B. Douglas Bernheim and Michael Whinston. 1986. Common Agency. *Econometrica* 54, 4 (1986), 923–42.
 - [4] B. Douglas Bernheim and Michael Whinston. 1986. Menu auctions, resource allocation, and economic influence. *The Quarterly Journal of Economics* 101, 1 (1986), 1–31.
 - [5] Patrick Bolton and Mathias Dewatripont. 2005. *Contract theory*. MIT press.
 - [6] Gabriel Carroll. 2015. Robustness and linear contracts. *American Economic Review* 105, 2 (2015), 536–63.
 - [7] Edward H. Clarke. 1971. Multipart pricing of public goods. *Public Choice* 11, 1 (1971), 17–33.
 - [8] Francesco Decarolis, Maris Goldmanis, and Antonio Penta. 2020. Marketing agencies and collusive bidding in online ad auctions. *Management Science* 66, 10 (2020), 4433–4454.
 - [9] Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. 2019. Simple versus Optimal Contracts. In *Proceedings of the 2019 ACM Conference on Economics and Computation EC*, Anna Karlin, Nicole Immorlica, and Ramesh Johari (Eds.). ACM, 369–387.
 - [10] Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. 2020. The Complexity of Contracts. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020*, Shuchi Chawla (Ed.). SIAM, 2688–2707.
 - [11] Miguel A. Goberna, Valentín Jornet, and Margarita M.L. Rodríguez. 2003. On linear systems containing strict inequalities. *Linear Algebra Appl.* 360 (2003), 151 – 171.
 - [12] Christian Gollier, Pierre-François Koehl, and Jean-Charles Rochet. 1997. Risk-taking behavior with limited liability and risk aversion. *Journal of Risk and Insurance* 64, 2 (1997), 347–370.
 - [13] J. Green and J. J. Laffont. 1977. Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods. *Econometrica* 45 (1977), 727–738.
 - [14] Theodore Groves. 1973. Incentives in Teams. *Econometrica* 41, 4 (1973), 617–631.
 - [15] Guru Guruganesh, Jon Schneider, and Joshua Wang. 2020. Contracts under Moral Hazard and Adverse Selection. (2020). <https://arxiv.org/abs/2010.06742>
 - [16] Chien-Ju Ho, Aleksandrs Slivkins, and Jennifer Wortman Vaughan. 2016. Adaptive Contract Design for Crowdsourcing Markets: Bandit Algorithms for Repeated Principal-Agent Problems. *Journal of Artificial Intelligence Research* 55 (2016), 317–359.

- [17] Bengt Holmström. 1979. Groves' Scheme on Restricted Domains. *Econometrica* 47, 5 (1979), 1137–1144.
- [18] Robert D Innes. 1990. Limited liability and incentive contracting with ex-ante action choices. *Journal of economic theory* 52, 1 (1990), 45–67.
- [19] Jon Kleinberg and Robert Kleinberg. 2018. Delegated Search Approximates Efficient Search. In *Proceedings of the 2018 ACM Conference on Economics and Computation EC*. ACM, 287–302.
- [20] Jon Kleinberg and Manish Raghavan. 2019. How Do Classifiers Induce Agents to Invest Effort Strategically?. In *Proceedings of the 2019 ACM Conference on Economics and Computation EC*. ACM, 825–844.
- [21] Jean-Jacques Laffont and David Martimort. 2009. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press.
- [22] Ron Lavi and Elisheva S. Shamash. 2019. Principal-Agent VCG Contracts. In *Proceedings of the 2019 ACM Conference on Economics and Computation EC*. ACM, 783.
- [23] Roger B. Myerson and Mark A. Satterthwaite. 1983. Efficient Mechanisms for Bilateral Trading. *Journal of Economic Theory* 29, 2 (1983), 265–281.
- [24] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V.Editors Vazirani (Eds.). 2007. *Algorithmic Game Theory*. Cambridge University Press.
- [25] Michael Peters and Balázs Szentes. 2012. Definable and contractible contracts. *Econometrica* 80, 1 (2012), 363–411.
- [26] Andrea Prat and Aldo Rustichini. 2003. Games played through agents. *Econometrica* 71, 4 (2003), 989–1026.
- [27] Tim Roughgarden. 2016. *Twenty lectures on algorithmic game theory*. Cambridge University Press.
- [28] Royal Swedish Academy of Sciences. 2016. Scientific Background on the 2016 Nobel Prize in Economic Sciences.
- [29] Bernard Salanie. 2017. *The Economics of Contracts: A Primer*. MIT press.
- [30] William Vickrey. 1961. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance* 16, 1 (1961), 8–37.

- [31] Wikipedia contributors. 2021. Multiple principal problem — Wikipedia, The Free Encyclopedia. https://en.wikipedia.org/wiki/Multiple_principal_problem [Online; accessed 5-February-2021].
- [32] Wikipedia contributors. 2021. Principal-agent problem — Wikipedia, The Free Encyclopedia. https://en.wikipedia.org/wiki/Principal-agent_problem [Online; accessed 5-February-2021].
- [33] Shenke Xiao, Zihe Wang, Mengjing Chen, Pingzhong Tang, and Xiwang Yang. 2020. Optimal Common Contract with Heterogeneous Agents. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*. AAAI Press, 7309–7316.

A Social Inefficiency of First Price Contracts

A.1 Price of Anarchy: Analysis of Example 1

In this section we prove that in Example 1, action a_3 and the bid profile $b^1 = (0, 0, 1+\gamma)$, $b^2 = (0, 1+\epsilon, 0)$ and $b^\ell = (0, 0, 0) \forall \ell > 2$ constitute an equilibrium. To do so, we show that the agent maximizes his utility, and that each principal cannot benefit from deviating and switching her bid. Note that $\text{Wel}^{a_3}(v) = 1$, while the action which maximizes social welfare is a_1 with $\text{Wel}^{a_1}(v) = n - 2 > 1$. Therefore, the price of anarchy is $\text{Wel}^{a_3}(v)/\text{Wel}^{a_1}(v) = 1/(n - 2)$, which tends to zero when $n \rightarrow \infty$.

For the agent, the payoff from taking action a_3 or action a_2 is 1, while the payoff from action a_1 is 0. Recall that the tie-breaking rule is in favor of the declared welfare and then the cost (see footnote 7). Since a_2 and a_3 have the same declared welfare the agent takes action a_3 which incurs more cost. Therefore, $x^*(b) = a_3$. Principal 1 cannot raise her utility by changing the bid for o_3 . If she raises her bid, the agent still takes a_3 and her utility will strictly decrease since she pays more to the agent. If principal 1 lowers her bid, then the agent takes a_2 , and principal 1’s utility will not vary. Principal 2 cannot raise her utility by changing the bid for o_2 . If she raises her bid, the agent takes a_2 and her utility will strictly decrease since she pays more than her value to the agent. If principal 2 lowers her bid, then the agent still takes a_3 , and principal 2’s utility does not vary. As for principal $\ell > 2$, she cannot raise her utility by announcing a different bid for o_1 . Assume she announces a positive bid $\tilde{b}^\ell(o_1) > 0$ and consider two cases: (i) If $\tilde{b}^\ell(o_1) \leq 1$, the agent takes a_3 and the utility of principal ℓ does not vary. (ii) If $\tilde{b}^\ell(o_3) > 1$, the agent takes a_1 , yet the utility of principal ℓ becomes negative. We have shown that neither the agent nor the principals may benefit from deviation. Hence, action a_3 and the bid profile b constitute an equilibrium.

A.2 Price of Stability: Analysis of Example 2

In this section we prove that in Example 2, the agent takes a_1 in all equilibria. The proof relies on the following Claim.

Claim 7. Consider Example 2. For every bid profile that incentivizes the agent to take a_i for $i > 1$, the principal's expected utility is at most $1 - \epsilon/(1 - \gamma)$.

Observe that the smallest possible bid profile $b^1 = (q, q)$ incentivizes the agent to take action a_1 . Also, the principal's expected utility given b^1 is $1 > 1 - \epsilon/(1 - \gamma)$. Thus, there is no equilibrium where the agent takes a_i for $i > 1$.

Proof of Claim 7. Let b^1 be a bid profile that incentivizes a_i for $i > 1$. It must hold that $\mathbb{E}_{o \sim F_{|a_i}}[b^1(o)] - \psi(a_i) \geq \mathbb{E}_{o \sim F_{|a_{i-1}}}[b^1(o)] - \psi(a_{i-1})$. That is, $(\gamma^{q-i} - \gamma^{q-(i-1)})(b^1(o_2) - b^1(o_1)) \geq \gamma^{1-i}(1 - \gamma) - 1 + (\gamma + \epsilon)$. Since $b^1(o_1)$ must be at least q , it follows from the last inequality that $b^1(o_2) \geq q + \gamma^{1-q} - \gamma^{i-q} + \epsilon/(\gamma^{q-i} - \gamma^{q-(i-1)})$. Using this, the following chain of inequalities shows that the principal's expected utility is at most $1 - \epsilon/(1 - \gamma)$:

$$\begin{aligned}
& \mathbb{E}_{o \sim F_{a_i}}[v^1(o) - t^1(b, o)] \\
&= \gamma^{q-i}(v^1(o_2) - b^1(o_2)) + (1 - \gamma^{q-i})(v^1(o_1) - b^1(o_1)) \\
&\leq \gamma^{q-i}(v^1(o_2) - q - \gamma^{1-q} + \gamma^{i-q} - \epsilon/(\gamma^{q-i} - \gamma^{q-(i-1)})) + (1 - \gamma^{q-i})(v^1(o_1) - q) \\
&= \gamma^{q-i}(\gamma^{i-q} - \epsilon/(\gamma^{q-i} - \gamma^{q-(i-1)})) \\
&= 1 - \epsilon/(1 - \gamma).
\end{aligned}$$

□

A.3 Existence of IIVCG for Examples 1 and 2

Proposition 6. There exists an IIVCG contract for the setting in Example 1.

Proof. We must show that $k_{a^*(b)} \leq \sum_{\ell \in [n]} m^\ell(b) \forall b \in \mathcal{V}$. We use the following case analysis:

1. $a^*(b) = a_1$. Note that $k_{a_1} = 0$, and that $m^\ell(b) \geq 0 \forall \ell \in [n], b \in \mathcal{V}$. Thus, $k_{a^*(b)} \leq \sum_{\ell \in [n]} m^\ell(b)$.
2. $a^*(b) = a_2$. Note that $k_{a_2} = \epsilon$ by using $w = (0, \epsilon, 0)$. Further, since $a^*(b) = a_2$, and by the definition of the domain, $\text{Wel}^{a^*(b)}(b) = b^2(o_2) - \epsilon$. Thus, $\text{Wel}^{a^*(b)}(b^{-1}, \mathbf{0}) = \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0}) = b^2(o_2) - \epsilon \forall \ell > 2$, and $\min_{\tilde{v}^1 \in \mathcal{V}^1} \text{Wel}^{a^*(b^{-1}, \tilde{v}^1)}(b^{-1}, \tilde{v}^1), \min_{\tilde{v}^\ell \in \mathcal{V}^\ell} \text{Wel}^{a^*(b^{-\ell}, \tilde{v}^\ell)}(b^{-\ell}, \tilde{v}^\ell) \geq b^2(o_2) - \epsilon \forall \ell \in [n]$. This implies that $m^1(b) = m^\ell(b) = 0 \forall \ell > 2$. Further, $\text{Wel}^{a^*(b)}(b^{-2}, \mathbf{0}) = -\epsilon$. Thus, $m^2(b) \geq \epsilon$. It follows that $k_{a^*(b)} \leq \sum_{\ell \in [n]} m^\ell(b)$.
3. $a^*(b) = a_3$. Note that $k_{a_3} = \gamma$ by using $w = (0, 0, \gamma)$. Further, since $a^*(b) = a_3$, and by the definition of the domain, $\text{Wel}^{a^*(b)}(b) = b^1(o_3) - \gamma$. Thus, $\text{Wel}^{a^*(b)}(b^{-2}, \mathbf{0}) = \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0}) = b^1(o_3) - \gamma \forall \ell > 2$, and $\min_{\tilde{v}^2 \in \mathcal{V}^2} \text{Wel}^{a^*(b^{-2}, \tilde{v}^2)}(b^{-2}, \tilde{v}^2), \min_{\tilde{v}^\ell \in \mathcal{V}^\ell} \text{Wel}^{a^*(b^{-\ell}, \tilde{v}^\ell)}(b^{-\ell}, \tilde{v}^\ell) \geq b^1(o_3) - \gamma \forall \ell \in [n]$. This implies that $m^2(b) = m^\ell(b) = 0 \forall \ell > 2$. Further, $\text{Wel}^{a^*(b)}(b^{-1}, \mathbf{0}) = -\gamma$. Thus, $m^1(b) \geq \gamma$. It follows that $k_{a^*(b)} \leq \sum_{\ell \in [n]} m^\ell(b)$.

This completes the proof. □

Proposition 7. *There exists an IIVCG contract for the setting in Example 2.*

To prove this, we first state the following claims.

Claim 8. $k_{a_j} = \gamma^{1-j} - 1 + \epsilon/(1 - \gamma) \forall j \in [q]$.

Claim 9. $m^\ell(b) = q + \psi(a^*(b)) \forall b \in \mathcal{V}$.

Proof of Proposition 7. As shown in Section 4.1, the setting is applicable if $k_{a^*(b)} \leq m^1(b) \forall b \in \mathcal{V}$. Equivalently, according to Claims 8 and 9, the setting is applicable if $\gamma^{1-j} - 1 + \epsilon/(1 - \gamma) \leq q + \gamma^{1-j} - j + (j-1)(\gamma + \epsilon)$, where $a_j = a^*(b)$, for all $b \in \mathcal{V}$. Note that $(j-1)(1-\gamma-\epsilon) + \epsilon/(1-\gamma) \leq q$ for all $j \in [q]$, since the left-hand side is at most $(q-1)(1-\gamma-\epsilon) + 1$ which, in turn, is at most q . This completes the proof. \square

Proof of Claim 8. Recall that $k_{a_j} = \min_{w \in \mathcal{L}_{a_j}} \mathbb{E}_{o \sim F_{|a_j}}[w(o)]$. We show that one w which minimizes $\mathbb{E}_{o \sim F_{|a_j}}[w(o)]$ is $w = (0, \gamma^{1-q} - \gamma^{j-q} + \gamma^{j-q} \cdot \epsilon/(1 - \gamma))$. Thus, $k_{a_j} = \gamma^{1-j} - 1 + \epsilon/(1 - \gamma)$.

First, we prove that $w \in \mathcal{L}_{a_j}$. To do so, we must show that $a_j \in \arg \max_{a' \in \mathcal{A}} \mathbb{E}_{o \sim F_{|a'}}[w(o)] - \psi(a')$. Equivalently, that $\mathbb{E}_{o \sim F_{|a_j}}[w(o)] - \psi(a_j) \geq \mathbb{E}_{o \sim F_{|a_{j'}}}[w(o)] - \psi(a_{j'}) \forall j' \in [q]$. That is, $(\gamma^{q-j} - \gamma^{q-j'})(w_2 - w_1) \geq \psi(a_j) - \psi(a_{j'}) \forall j' \in [q]$. Replacing w and ψ , the last inequality holds if

$$(\gamma^{q-j} - \gamma^{q-j'})(\gamma^{1-q} - \gamma^{j-q} + \gamma^{j-q} \cdot \epsilon/(1 - \gamma)) \geq \gamma^{1-j} - \gamma^{1-j'} - (j - j')(1 - \gamma - \epsilon) \forall j' \in [q].$$

Reorganizing both sides, this is equal to

$$\gamma^{1-j} - \gamma^{1-j'} + (\gamma^{q-j} - \gamma^{q-j'})(-\gamma^{j-q} + \gamma^{j-q} \cdot \epsilon/(1 - \gamma)) \geq \gamma^{1-j} - \gamma^{1-j'} - (j - j')(1 - \gamma - \epsilon) \forall j' \in [q].$$

Which, by subtracting $\gamma^{1-j} - \gamma^{1-j'}$ from both sides, holds if

$$(\gamma^{q-j} - \gamma^{q-j'})(-\gamma^{j-q} + \gamma^{j-q} \cdot \epsilon/(1 - \gamma)) \geq -(j - j')(1 - \gamma - \epsilon) \forall j' \in [q]. \quad (4)$$

Observe the left-hand side in (4), by dividing the first term by γ^{q-j} and multiplying the second term by γ^{q-j} , it is equal to $(1 - \gamma^{j-j'})(-1 + \epsilon/(1 - \gamma)) = -(1 - \gamma^{j-j'})(1 - \gamma - \epsilon)/(1 - \gamma)$. Then, by dividing both sides in (4) by $-(1 - \gamma - \epsilon)$, we conclude that $w \in \mathcal{L}_{a_j}$ if

$$(1 - \gamma^{j-j'})/(1 - \gamma) \leq (j - j') \forall j' \in [q]. \quad (5)$$

The following case analysis shows the above inequality holds:

1. For $j - j' = 0, j - j' = 1$, this is immediate.
2. For $j - j' \geq 2$, (5) holds since $1 - \gamma^{j-j'} \leq 1 \leq 2 \cdot (1 - \frac{1}{q}) \leq 2 \cdot (1 - \gamma) \leq (j - j')(1 - \gamma)$.
3. For $0 > j - j'$, (5) holds since $1 - \gamma^{j-j'} = -(\frac{1}{\gamma}^{j'-j} - 1) \leq -(2^{j'-j} - 1) \leq -(j' - j) \leq (j - j')(1 - \gamma)$, where the first inequality follows from the fact that $\gamma \leq \frac{1}{q} \leq \frac{1}{2}$.

We have shown that $w \in \mathcal{L}_{a_j}$. It is now left to show that w minimizes $\mathbb{E}_{o \sim F_{|a_j}}[w(o)]$. Observe that for every $w' \in \mathcal{L}_a$, it must hold that $\mathbb{E}_{o \sim F_{|a_j}}[w'(o)] \geq \mathbb{E}_{o \sim F_{|a_{j-1}}} [w'(o)]$. That is, $(\gamma^{q-j} - \gamma^{q-j+1})(w'_2 - w'_1) \geq \gamma^{1-j} - \gamma^{2-j} - (1 - \gamma - \epsilon)$. By reorganizing both sides, the last inequality equals to $\gamma^{q-j}(1 - \gamma)(w'_2 - w'_1) \geq (\gamma^{1-j} - 1)(1 - \gamma) + \epsilon$. Dividing both sides by $\gamma^{q-j}(1 - \gamma)$, and since $w'_1 \geq 0$, it holds that $w'_2 \geq \gamma^{1-q} - \gamma^{j-q} + \epsilon/(\gamma^{q-j} - \gamma^{q-(j-1)})$. Since w meets these lower bounds, it minimizes $\mathbb{E}_{o \sim F_{|a_j}}[w(o)]$. This completes the proof. \square

Proof of Claim 9. Recall that $m^1(b) = \min_{\tilde{v}^1 \in \mathcal{V}^1} \text{Wel}^{a^*(b^{-1}, \tilde{v}^1)}(b^{-1}, \tilde{v}^1) - \text{Wel}^{a^*(b)}(b^{-1}, \mathbf{0})$. Note that $\text{Wel}^{a^*(b)}(b^{-1}, \mathbf{0}) = -\psi(a^*(b))$. Further, $\min_{\tilde{v}^1 \in \mathcal{V}^1} \text{Wel}^{a^*(b^{-1}, \tilde{v}^1)}(b^{-1}, \tilde{v}^1) = q$, when $\tilde{v}^1 = (q, q)$: Since for every $v^1 \neq \tilde{v}^1 \in \mathcal{V}^1$ it holds that $\tilde{v}^1 \leq v^1$ (element wise). Thus, $\text{Wel}^{a^*(b^{-1}, \tilde{v}^1)}(b^{-1}, \tilde{v}^1) \leq \text{Wel}^{a^*(b^{-1}, v^1)}(b^{-1}, v^1) \leq \text{Wel}^{a^*(b^{-1}, v^1)}(b^{-1}, v^1)$. This completes the proof. \square

To complete the picture, one IIVCG contract for Example 2 may be contracted using $h^1(b^{-1}) = q$ and $w = (0, \gamma^{1-q} - \gamma^{j-q} + \gamma^{j-q} \cdot \epsilon/(1 - \gamma))$ where $a_j = a^*(b) \forall b \in \mathcal{V}$. That is,

$$\begin{aligned} t^1(b, o_1) &= q - \{(j-1)(1 - \gamma - \epsilon) + \epsilon/(1 - \gamma)\}. \\ t^1(b, o_2) &= q - \{(j-1)(1 - \gamma - \epsilon) + \epsilon/(1 - \gamma)\} + \gamma^{1-q} - \gamma^{j-q} + \gamma^{j-q} \cdot \epsilon/(1 - \gamma). \end{aligned}$$

This contract satisfies limited liability because $q \geq (q-1)(1 - \gamma - \epsilon) + 1 \geq (j-1)(1 - \gamma - \epsilon) + \epsilon/(1 - \gamma)$, and $\gamma^{j-q} - \gamma^{j-q} \cdot \epsilon/(1 - \gamma) = \gamma^{1-q} \cdot \gamma^{j-1} \cdot (1 - \epsilon/(1 - \gamma)) \leq \gamma^{1-q}$, since $\gamma^{j-1}, (1 - \epsilon/(1 - \gamma)) \leq 1$. This contract also satisfies individual rationality since for $a_j = a^*(v)$, it must hold that $\mathbb{E}_{o \sim F_{|a_j}}[v^1(o)] \geq q + \psi(a_j) = \mathbb{E}_{o \sim F_{|a_j}}[t^1(v, o)]$. The inequality is immediate for $j = 1$, since $\psi(a_j) = 0$ and $v \in \mathcal{V} = \mathbb{R}_{\geq q}^2$. Suppose towards a contradiction that it does not hold for $j > 1$, then $\text{Wel}^{a^1}(v) \geq q \geq \mathbb{E}_{o \sim F_{|a_j}}[v^1(o)] - \psi(a_j) = \text{Wel}^{a^j}(v)$, in contradiction to the fact that $a^*(v) = a_j$.

B IIVCG Contracts: Missing Proofs from Section 3

B.1 Uniqueness of IIVCG Contracts: Proof of Theorem 2

To prove Theorem 2 we rely on a result of Holmstrom [17]:

Lemma 4 (Lemma 1 in [17]). *Let $f : [0, 1]^2 \rightarrow \mathbb{R}$ and $g : [0, 1] \rightarrow \mathbb{R}$ satisfy:*

1. $y \in \arg \max_{m \in [0, 1]} f(m, y)$,
2. $y \in \arg \max_{m \in [0, 1]} f(m, y) + g(m)$,
3. $|\frac{\partial f(m, y)}{\partial y}| \leq K < \infty$ for all $m, y \in [0, 1]$.

Then $g(m) = \text{constant}$ on $[0, 1]$.

Proof of Theorem 2. Define $h^\ell : \mathcal{V} \rightarrow \mathbb{R}$ as follows: $h^\ell(b) = \mathbb{E}_{o \sim F_{|a^*(b)}}[t^\ell(b, o)] + \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0})$ $\forall b \in \mathcal{V}$ and fix an arbitrary $b^{-\ell} \in \mathcal{V}^{-\ell}$. To show that t is an IIVCG contract, it suffices to show that $h^\ell(b^{-\ell}, \cdot)$ is constant on b^ℓ . We make the following observations: First, bidding truthfully is socially efficient, i.e.,

$$v^\ell \in \arg \max_{b^\ell \in \mathcal{V}^\ell} \text{Wel}^{a^*(b)}(b^{-\ell}, v^\ell) \quad \forall v^\ell \in \mathcal{V}^\ell. \quad (6)$$

Second, since principal ℓ is truthful, $v^\ell \in \arg \max_{b^\ell \in \mathcal{V}^\ell} \mathbb{E}_{o \sim F_{|a^*(b)}}[v^\ell(o) - t^\ell(b, o)] \quad \forall v^\ell \in \mathcal{V}^\ell$. Replacing $\mathbb{E}_{o \sim F_{|a^*(b)}}[t^\ell(b, o)]$ with $h^\ell(b) - \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0})$, it holds that

$$v^\ell \in \arg \max_{b^\ell \in \mathcal{V}^\ell} \text{Wel}^{a^*(b)}(b^{-\ell}, v^\ell) - h^\ell(b) \quad \forall v^\ell \in \mathcal{V}^\ell. \quad (7)$$

After establishing (6) and (7), we show that $h^\ell(b^{-\ell}, \cdot)$ is constant on b^ℓ . Equivalently, we show that for every pair of bid profiles $b^\ell, \tilde{b}^\ell \in \mathcal{V}^\ell$ $h^\ell(b^{-\ell}, b^\ell) = h^\ell(b^{-\ell}, \tilde{b}^\ell)$. To do so, define $\mathcal{V}^\ell(b^\ell, \tilde{b}^\ell) = \{v^\ell(y) \mid y \in [0, 1]\}$, where $v^\ell(y) : \mathcal{O} \rightarrow \mathbb{R}$ as follows $v^\ell(y; o) = yb^\ell(o) + (1 - y)\tilde{b}^\ell(o)$ $\forall o \in \mathcal{O}, y \in [0, 1]$. Note that $v^\ell(1) = b^\ell$ and that $v^\ell(0) = \tilde{b}^\ell$. Thus, it suffices to show that $h^\ell(b^{-\ell}, v^\ell(1)) = h^\ell(b^{-\ell}, v^\ell(0))$. By the assumption that \mathcal{V} is convex, $\mathcal{V}^\ell(b^\ell, \tilde{b}^\ell)$ is a sub domain of \mathcal{V}^ℓ . Hence, it follows from (6) and (7) that, using the notation of $a(m) = a^*(b^{-\ell}, v^\ell(m))$,

$$y \in \arg \max_{m \in [0, 1]} \text{Wel}^{a(m)}(b^{-\ell}, v^\ell(y)) \quad \forall y \in [0, 1]. \quad (8)$$

$$y \in \arg \max_{m \in [0, 1]} \text{Wel}^{a(m)}(b^{-\ell}, v^\ell(y)) - h^\ell(b^{-\ell}, v^\ell(m)) \quad \forall y \in [0, 1]. \quad (9)$$

Define $f(m, y) = \text{Wel}^{a(m)}(b^{-\ell}, v^\ell(y))$ and $g(m) = -h^\ell(b^{-\ell}, v^\ell(m)) \quad \forall m, y \in [0, 1]$. With these notations, note that (8), (9) are exactly 1,2 in Lemma 4 (respectively). To apply Lemma 4 we show that 3 holds. Define $K = \arg \max_{a \in \mathcal{A}} |\mathbb{E}_{o \sim F_{|a}}[b^\ell(o) - \tilde{b}^\ell(o)]|$. Since \mathcal{A} is a finite set, $K < \infty$. We show that $|\frac{\partial f(m, y)}{\partial y}| \leq K$ for all $m, y \in [0, 1]$ as follows. By definition of f and of the social welfare $\text{Wel}(\cdot)$,

$$\left| \frac{\partial f(m, y)}{\partial y} \right| = \left| \frac{\partial \text{Wel}^{a(m)}(b^{-\ell}, v^\ell(y))}{\partial y} \right| = \left| \frac{\partial \mathbb{E}_{o \sim F_{|a(m)}}[v^\ell(y, o)]}{\partial y} \right| \quad (10)$$

By definition of $v^\ell(y, o)$,

$$\mathbb{E}_{o \sim F_{|a(m)}}[v^\ell(y, o)] = y \cdot \mathbb{E}_{o \sim F_{|a(m)}}[b^\ell(o)] + (1 - y) \cdot \mathbb{E}_{o \sim F_{|a(m)}}[\tilde{b}^\ell(o)]. \quad (11)$$

According to (10) and (11) and definition of K ,

$$\left| \frac{\partial \mathbb{E}_{o \sim F_{|a(m)}}[v^\ell(y, o)]}{\partial y} \right| = |\mathbb{E}_{o \sim F_{|a^*(m)}}[b^\ell(o) - \tilde{b}^\ell(o)]| \leq K.$$

By Lemma 4, $g(m)$ is constant on m . That is, $h^\ell(b^{-\ell}, v^\ell(m))$ is constant on m . We conclude that $h^\ell(b^{-\ell}, v^\ell(1)) = h^\ell(b^{-\ell}, v^\ell(0))$ and the proof is complete. \square

B.2 Characterization of IIVCG: Proof of Lemma 1

Proof of Lemma 1. We start with the forward direction. Let t be an IIVCG contract. Note that $t^\ell : \mathcal{V} \times \mathcal{O} \rightarrow \mathbb{R}$ can be re-composed as the sum of two functions; $c^\ell : \mathcal{V} \rightarrow \mathbb{R}$, that depends only on b , and $g^\ell : \mathcal{V} \times \mathcal{O} \rightarrow \mathbb{R}$ that depends on b as well as on the realised outcome, i.e., $t^\ell(b, o) = c^\ell(b) + g^\ell(b, o)$. With these notations, for every principal ℓ and every bid profile b ,

$$\mathbb{E}_{o \sim F_{|a^*(b)}}[t^\ell(b, o)] = c^\ell(b) + \mathbb{E}_{o \sim F_{|a^*(b)}}[g^\ell(b, o)]. \quad (12)$$

By IIVCG Definition, there exist $\{h^\ell\}$ such that

$$\mathbb{E}_{o \sim F_{|a^*(b)}}[t^\ell(b, o)] = h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^{-\ell}, \mathbf{0}), \quad (13)$$

where h^ℓ does not depend on b^ℓ . Equalizing (12), (13) we get $c^\ell(b) = h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^{-\ell}, g^\ell(b, o))$. Since $t^\ell(b, o) = c^\ell(b) + g^\ell(b, o)$, $t^\ell(b, o) = h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^{-\ell}, g^\ell(b, o)) + g^\ell(b, o)$. By IIVCG contract definition, the agent maximizes the (declared) social welfare, i.e., $x^*(b) = a^*(b)$. Thus, $a^*(b) \in \arg \max_{a \in \mathcal{A}} \mathbb{E}_{o \sim F_{|a}}[\sum_{\ell \in [n]} t^\ell(b, o)] - \psi(a)$. The only outcome-dependent term in $t^\ell(b, o)$ is $g^\ell(b, o)$.¹⁸ Hence,

$$a^*(b) \in \arg \max_{a \in \mathcal{A}} \mathbb{E}_{o \sim F_{|a}}[\sum_{\ell \in [n]} g^\ell(b, o)] - \psi(a). \quad (14)$$

Define $w^\ell : \mathcal{O} \rightarrow \mathbb{R}_{\geq 0}^m$ as follows: $w^\ell(o) = g^\ell(b, o) - \min_{o \in \mathcal{O}} g^\ell(b, o) \forall o \in \mathcal{O}$. Hence, $t^\ell(b, o) = h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^\ell, w^\ell) + w^\ell(o)$, where $\sum_{\ell \in [n]} w^\ell \in \mathcal{L}_{a^*(b)}$ by (14) and the definition of \mathcal{L}_a .

For the backward direction, let t be a contract as specified in Lemma 1. Note that the only outcome dependent term in $t^\ell(b, o)$ is $w^\ell(o)$. Thus,

$$\arg \max_{a \in \mathcal{A}} \mathbb{E}_{o \sim F_{|a}}[\sum_{\ell \in [n]} t^\ell(b, o)] - \psi(a) = \arg \max_{a \in \mathcal{A}} \mathbb{E}_{o \sim F_{|a}}[\sum_{\ell \in [n]} w^\ell(o)] - \psi(a). \quad (15)$$

Since, $\sum_{\ell \in [n]} w^\ell \in \mathcal{L}_{a^*(b)}$, $a^*(b) \in \arg \max_{a \in \mathcal{A}} \mathbb{E}_{o \sim F_{|a}}[\sum_{\ell \in [n]} w^\ell(o)] - \psi(a)$. Thus, and according to (15) we conclude that $a^*(b)$ maximizes the agent's utility. Furthermore, Property 2 in the definition of IIVCG (Definition 2) is also met. since principal ℓ 's expected payment is $\mathbb{E}_{o \sim F_{|a^*(b)}}[t^\ell(b, o)] = h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^\ell, w^\ell) + \mathbb{E}_{o \sim F_{|a^*(b)}}[w^\ell(o)]$. After rearranging the above $\mathbb{E}_{o \sim F_{|a^*(b)}}[t^\ell(b, o)] = h^\ell(b^{-\ell}) - \text{Wel}^{a^*(b)}(b^\ell, \mathbf{0})$ as desired. \square

¹⁸If $h^\ell(b^{-\ell})$ depends on the realised outcome, then it would also depend on principal ℓ 's bid (since ℓ 's bid influences the realised outcome indirectly by influencing the chosen action).

B.3 A Trade-off between LL and IR: Proof of Theorem 3

Here, we formally prove Theorem 3.

Example 4. *There are two actions a_1, a_2 and two possible outcomes o_1, o_2 . The distributions imposed by actions are: $p_{o \sim F|a_1}(o_1) = \frac{1}{2}$, $p_{o \sim F|a_2}(o_1) = 1$. There is one principal with valuation domain $\mathcal{V}^1 = \mathbb{R}_{\geq 0}^2$. The cost for each action is: $\psi(a_1) = 0$, $\psi(a_2) = \epsilon$.*

Proof of Theorem 3. Consider the common agency setting in Example 4 and the valuation profile $v^1 = (3\epsilon, 0)$. Note that $\text{Wel}^{a_1}(v) = 1.5\epsilon$, and $\text{Wel}^{a_2}(v) = 2\epsilon$. Thus, the socially efficient action is a_2 . Suppose toward a contradiction that there exists an IIVCG contract t that satisfies LL and IR. Since t is an IIVCG contract, in a dominant strategy equilibrium the principal is truthful, and the agent takes the socially efficient action, i.e., a_2 . Since the agent maximizes the sum of payments minus cost, it must hold that $\sum_{\ell \in [n]} t^\ell(b, o_1) - \epsilon \geq \sum_{o \in \mathcal{O}} \frac{1}{2} \cdot \sum_{\ell \in [n]} t^\ell(b, o)$. Thus, $t^1(b, o_1) - t^1(b, o_2) \geq 2\epsilon$. By the assumption that t satisfies LL, $t^1(b, o_2) \geq 0$, so, $t^1(b, o_1) \geq 2\epsilon$. That is, $t^1(b, o_1) = \mathbb{E}_{o \sim F|a_2}[t^1(b, o)] = h^1(b^{-1}) - \text{Wel}^{a_2}(b^{-1}, \mathbf{0}) \geq 2\epsilon$. Since $\text{Wel}^{a_2}(b^{-1}, \mathbf{0}) = -\epsilon$, we conclude that $h^1(b^{-1}) \geq \epsilon$. Now, consider the following valuation profile: $\tilde{v}^1 = (0, 0)$. Again, in a dominant strategy equilibrium the principal is truthful $\tilde{b} = \tilde{v}$ and the agent takes the socially efficient action, which is now a_1 . By definition of IIVCG, h^1 is independent of b^1 . Thus, $h^1(b^{-1}) = h^1(\tilde{b}^{-1}) \geq \epsilon$. Recall that $\mathbb{E}_{o \sim F|a_1}[t^1(\tilde{b}, o)] = h^1(\tilde{b}^{-1}) - \text{Wel}^{a_1}(\tilde{b}^{-1}, \mathbf{0})$, where $\text{Wel}^{a_1}(\tilde{b}^{-1}, \mathbf{0}) = 0$. Thus, $\mathbb{E}_{o \sim F|a_1}[t^1(\tilde{b}, o)] = h^1(\tilde{b}^{-1}) \geq \epsilon$, while $\mathbb{E}_{o \sim F|a_2}[\tilde{v}^1(o)] = 0$. This is a contradiction to the fact that t satisfies IR. \square