

Buying Opinions

Supplementary Appendix (For Online Publication)

Mark Whitmeyer* Kun Zhang†

September 28, 2022

Contents

A	When the Agent “Skips Town” (An Example)	1
B	Salvage Value	3
C	Efficient Implementation With a Restricted Set of Distributions	5
D	An Analog of Proposition 5.1 for a Prior with a Density	6
E	Interim IR in the Canonical Problem (Section 5.1.1.)	8
F	Non-genericity of the STP Contract	9

A When the Agent “Skips Town” (An Example)

The purpose of this subsection is to illustrate how the principal may wish to have the agent take his outside option with positive probability if such an exit does not hurt the principal (too much). First, a minor remark that justifies one of our assumptions in the text:

*Arizona State University, Email: mark.whitmeyer@gmail.com.

†Arizona State University, Email: kunzhang@asu.edu.

Remark A.1. If the principal’s payoff from the agent taking his outside option is $< -v_0$, there is no optimal contract in which the agent takes his outside option with positive probability.

Proof. The result is nearly immediate: suppose for the sake of contradiction there is an optimal contract in which the agent takes his outside option with strictly positive probability. Replace that contract with an identical one with one exception: now the principal offers a constant payout of v_0 to the agent for reporting the previous “null message belief.” The agent’s incentives are unaffected and the principal’s payoff is strictly higher, contradicting the original contract’s purported optimality. ■

Next we will illustrate that if the disutility incurred by the principal when the agent takes his outside option is not too low, the principal may prefer that the agent not have an interim participation constraint,¹ which allows the principal to write contracts in which the agent exits the relationship with positive probability.

Consider the following example, which is illustrated in Figure 1. The state is binary, $\Theta = \{0, 1\}$, and $P(1) = (3 + 2e) / (5 + 5e) \approx .45$. The principal’s decision problem is a simple “match the state” task: she has two actions $a \in \{0, 1\}$ and obtains a payoff of 1 if $a = \theta$ and 0 otherwise. The agent’s cost of acquiring information is the entropy cost, the cost parameter is $\kappa = 1$ for simplicity, and she has an outside option of .05. The principal suffers no disutility if the agent takes her outside option.

If the principal controlled information herself, her distribution over posteriors would have support $\{1/(1 + e), e/(1 + e)\} \approx \{.27, .73\}$. As we note in Proposition 5.1, the principal can still implement this distribution efficiently even with the interim IR constraint. This distribution is optimal, therefore, when the agent may not exit the relationship after learning, which yields the principal an approximate payoff of $.62 - .05 = .57$ (the payoff in the decision problem net of the information cost minus the agent’s outside option). The value function for the agent induced by the STP contract, which is optimal when there is no interim participation constraint, and the agent’s resulting optimal learning are depicted in Figure 1a.

¹A moment’s reflection reveals that this statement is either incorrect or misleading: introducing an additional constraint in an optimization problem can never strictly increase the optimal payoff. What engenders the (potential) improvement here is that the addition of the interim IR constraint is also the addition of new contractual possibilities (contracts in which the agent exits the relationship after learning).

Now let us allow the agent to exit the relationship after learning and construct a contract that strictly improves the principal’s payoff. The contract consists of a single message, $.73$, and a transfer from sending this message of $2/3$ in state 1 and $-3/4$ in state 0. The agent’s optimal learning now has (approximate) support $\{.39, .73\}$, after which she exits the relationship, taking her outside option, or sends the single offered message, respectively. This yields the principal a payoff of approximately $.58$, a strict improvement. Moreover, the agent obtains strictly positive surplus as well, approximately $.06$. The value function for the agent induced by this contract as well as the agent’s optimal learning are depicted in Figure 1b.


This contract is not optimal for the principal, but it serves its purpose: the optimal contract in this example must be one in which the agent exits the relationship with positive probability. Nor does the agent exit the relationship with probability one (for this would correspond to him acquiring no information), and so in the optimal contract the agent acquires strictly positive surplus. Thus, allowing an interim exit may engender a strict Pareto improvement.

That enabling an agent to exit the relationship may strictly improve welfare is a consequence of the special kind of output the agent is asked to produce in our setup. Namely, he is asked to provide the principal with information: the informational content of an agent’s action (on path) is the same regardless of whether it is conveyed via a message or the agent’s resignation. One natural interpretation for the cost an agent’s exit imposes on the principal is that it is the cost of finding a new advisor. Consequently, if it is relatively cheap to do so, the principal prefers this to compensating the advisor.²


B Salvage Value

Our framework can also easily be adapted to allow the agent to have a “salvage” value for a posterior (upper semicontinuous function) $p(\mathbf{x})$, from exiting the relationship after obtaining posterior \mathbf{x} . This corresponds, for instance, to a scenario in which the agent is able to (credibly)

²One interpretation of this observation is as rationalizing the phenomenon of “shooting the messenger” or blaming the bearer of bad news—“...Yet the first bringer of unwelcome news, Hath but a losing office...” (Henry IV, pt. II) Rather than inform the principal himself (and suffer her wrath), the agent prefers to skip town, which nevertheless informs the principal about the situation.



(a) An optimal contract (the STP contract) when the agent may not exit *ex interim*



(b) A contract yielding a strict Pareto improvement to any “no exit” contract

Figure 1: In both 1a and 1b, $\mu \approx .45$. In the former, $x_L \approx .27$, $x_H \approx .73$, and the concavifying line $f_{\mathcal{H}}(x) = v_0 = .05$. In the latter, $x_L \approx .39$, $x_H \approx .73$ and $f_{\mathcal{H}}(x)$ is approximately $.44x - .14$.

sell information to a third party. In this case, incentive compatibility is left unchanged, but the participation constraint becomes

$$f_{\mathcal{H}}(\mathbf{x}) \geq p(\mathbf{x}) - \kappa c(\mathbf{x}) \quad \text{for all } \mathbf{x} \in \Delta(\Theta). \quad (IR - p)$$

An analog of Lemma 3.1 is immediate:

Lemma B.1. *A contract (M, t) implements distribution F if and only if*

- (i) $\text{supp}(F) = P_{(M,t)}$; and
- (ii) Constraint $IR - p$ holds; and
- (iii) if there is limited liability, $t(m, \theta) \geq 0$ for all $\theta \in \Theta$ and $m \in M$.

Given this, it is easy to see that the other results of the paper go through with the outside option function $v_0 - \kappa c$ replaced with the concavification of $p - \kappa c$.

C Efficient Implementation With a Restricted Set of Distributions

In principle, the agent could be further constrained in how she learns: perhaps F must be chosen not from the set of Bayes-plausible distributions \mathcal{F}_μ but from some (weak-*) compact subset \mathcal{P} of the Bayes' plausible distributions with finite support. One example of this is if there are just two states, $\Theta = \{0, 1\}$, the prior is $\frac{1}{2}$, and the agent has access to collection of binary signals $\pi_\alpha(1|1) = \pi_\alpha(0|0) = \alpha \in [\frac{1}{2}, 1]$. In this case \mathcal{P} is the collection of binary distributions with support $\{1 - \alpha, \alpha\}_{\alpha \in [\frac{1}{2}, 1]}$ and mean $\frac{1}{2}$.

We further assume that the cost functional restricted to subset \mathcal{P} is posterior separable. In the parametrized binary experiment example, any strictly convex, twice continuously differentiable function $c : [\frac{1}{2}, 1] \rightarrow \mathbb{R}_+$ with $c(\frac{1}{2}) = 0$ would do; e.g., $c(\alpha) = \kappa \left(\frac{\alpha^2}{2} - \frac{1}{8} \right)$ (with $\kappa > 0$).

Even though the set of distributions available to the agent is limited, the principal is nevertheless unconstrained by the agency problem. Namely, an analog of Proposition 5.1 holds:

Proposition C.1. *If the agent is risk neutral and not protected by limited liability, every (feasible) distribution $F \in \mathcal{P}$ with $\text{supp}(F) \subseteq \text{int } \Delta(\Theta)$ can be implemented efficiently.*

Proof. Fix any $F \in \mathcal{P}$ with $\text{supp}(F) \subseteq \text{int } \Delta(\Theta)$. As we note in Corollary 4.2 (and which is a consequence of Choquet’s theorem), this F can be obtained by randomizing over Bayes-plausible distributions each with support on at most n points. This collection of distributions is $(F_i)_{i \in I}$. By Proposition 5.1, each F_i can be implemented efficiently. Moreover, efficiency is synonymous with the tangency of the corresponding hyperplanes $f_{\mathcal{H},i}$ with the outside option curve $v_0 - \kappa c$ at the prior. However, this hyperplane is unique so each $f_{\mathcal{H},i}$ must equal some common $f_{\mathcal{H}}$. Furthermore, for each contract $(M_i, t_i)_{i \in I}$, the other part of Lemma 3.1 also must hold: $\text{supp } F_i = P(M_i, t_i)$ for all i . Accordingly, the contract (M, t) —defined as $M = \cup_{i \in I} M_i$ and $t(m, \theta) = t_i(m, \theta)$ for all $m \in M_i$ for all $i \in I$ —implements $F \in \mathcal{P}$ efficiently. ■

D An Analog of Proposition 5.1 for a Prior with a Density

Now let us derive an analog of Proposition 5.1 when there are uncountably many states and the prior admits a density.³ Suppose the state θ is distributed on the unit interval according to absolutely continuous cdf F . Given any distribution, G , we stipulate that the principal’s utility function is a convex function of the first z (non-centered) moments

$$m_1 = \int_0^1 x dG(x), m_2 = \int_0^1 x^2 dG(x), \dots, m_z = \int_0^1 x^z dG(x) .$$

We write $V(m_1, m_2, \dots, m_z)$. In addition, the agent’s cost of acquiring information is also a convex function of the first z non-centered moments $\kappa c(m_1, m_2, \dots, m_z)$ (where $\kappa > 0$ is a scaling parameter).

Associated with the prior F is a joint distribution over (x, x^2, \dots, x^z) , \hat{F} , and so the principal’s problem—should she control information acquisition herself—is

$$\max_{H \in \mathcal{F}(\hat{F})} \int (V - \kappa c) dH ,$$

where $\mathcal{F}(\hat{F})$ denotes the set of fusions of \hat{F} .⁴ For simplicity we also assume that the principal also has just $t < \infty$ actions, which ensures that the optimal fusion has support on at most t points.

Then,

³The idea for this sort of moment-based problem originates from a joint project of the first author with Andreas Kleiner, Benny Moldovanu, and Philipp Strack (Kleiner et al. (2022)), whom we thank.

⁴Kleiner et al. (2022) explore this problem in detail.

Proposition D.1. *If the agent is risk neutral and not protected by limited liability, every distribution $H \in \mathcal{F}(\hat{F})$ with support on t or fewer points can be implemented efficiently.*

Proof. By definition, distribution $H \in \mathcal{F}(\hat{F})$ and has support on t or fewer points in the z -dimensional unit hypercube. Denote by $\mu := (\mu_1, \dots, \mu_z)$ the barycenter of measure \hat{F} , and let f_μ denote the hyperplane tangent to $u_0 - \kappa c(m)$ at μ . The principal contracts the contract as follows: for each posterior m' in support of H , the contract is such that the agent's payoff gross of the information acquisition cost is of the form

$$\tau_{m'}(m) = \alpha_{m',1}m_1 + \alpha_{m',2}m_2 + \dots + \alpha_{m',z}m_z + \beta_{m'} ,$$

where each $\alpha_{m',i}$ and $\beta_{m'}$ are scalars;⁵ and such that $\tau_{m'}(m) - \kappa c(m)$ is tangent to f_μ at m' .

Evidently, given this contract it is optimal for the agent to acquire distribution H in the relaxed problem in which she simply chooses a distribution that is a dilation of the prior μ , and therefore it must be optimal for the agent to acquire distribution H in her fusion problem (and by construction $H \in \mathcal{F}(\hat{F})$). Moreover, the agent obtains zero rents. \blacksquare

Note that this results holds regardless of whether there is an interim participation constraint— if no such constraint exists, the principal can offer the STP contract, and if not (as above), the STP does not satisfy the interim IR constraint generically.

We can also illustrate this result with the following example. The prior is the uniform distribution on the unit interval. The principal has just two actions, and the payoff in the principal's decision problem can be written as a function of the posterior's first moment. Specifically, $V(m) = \max\{1 - m, m\}$. The agent's cost of acquiring information can also be written as a function of the posterior's first moment: $c(m) = (m - 1/2)^2$. The agent has an outside option of

⁵To elaborate, for each $m' \in \text{supp}(H)$, the transfer is given by

$$t(m', \theta) = \alpha_{m',1}\theta + \alpha_{m',2}\theta^2 + \dots + \alpha_{m',z}\theta^z + \beta_{m'} ,$$

and hence if the agent receives some signal ψ and reports m' , his expected payoff gross of the information acquisition cost is

$$\int_0^1 t(\theta) dG_\psi(\theta) = \alpha_{m',1} \int_0^1 \theta dG_\psi(\theta) + \alpha_{m',2} \int_0^1 \theta^2 dG_\psi(\theta) + \dots + \alpha_{m',z} \int_0^1 \theta^z dG_\psi(\theta) + \beta_{m'} ,$$

where G_ψ is the distribution over states conditional on the agent's signal. The right-hand side of the equation above is just $\tau_{m'}(m)$.




Figure 2: A “moment acquisition” example.

0. The principal can implement any (feasible) pair of posterior first moments m_L and m_H (with $m_L < 1/2 < m_H$) by offering the contract

$$\left(m_L; \left(\kappa (2m_L - 1), \frac{\kappa}{4} (1 - 4m_L^2) \right) \right), \left(m_H; \left(\kappa (2m_H - 1), \frac{\kappa}{4} (1 - 4m_H^2) \right) \right),$$

where each element of the contract is of the form $(m'; (\alpha_{m'}, \beta_{m'}))$.

Figure 2 illustrates the induced decision problem for the agent, $W(m)$ that implements the pair $m_L = .3$ and $m_H = .8$ when $\kappa = 1$.

E Interim IR in the Canonical Problem (Section 5.1.1.)

Consider the basic moral hazard setting with no limited liability and risk neutral principal and agent. For simplicity the agent chooses effort $a \in [0, 1]$ and her output, x , takes values in $[0, 1]$. Now, however, the agent’s output is private. After it realizes she may choose whether to turn in the output and receive the promised remuneration $w(x)$ or exit the relationship to take her outside option v_0 . As is standard in the literature, the agent’s cost of effort $c(\cdot)$ is strictly increasing and strictly convex, with $c(0) = 0$. Output takes values, x , in the unit interval. The family of conditional densities of output realizations $f(x|a)$ has full support for each $a \in [0, 1]$ and satisfies

the MLRP. Let $\mathcal{X} \subseteq [0, 1]$ denote the subset of agents that the principal wants to report, and a^* is the desired effort level. Both the principal and the agent are risk neutral.

The result:

Proposition E.1. *Unless the principal is observing the lowest possible effort ($a = 0$) and/or the agent is almost never turning in her output (\mathcal{X} has Lebesgue measure 0), the agent gets strictly positive rents.*

Proof. The principal solves

$$\min_{w(\cdot)} \int_{\mathcal{X}} w(x) f(x|a^*) dx ,$$

subject to

$$\int_{\mathcal{X}} w(x) f(x|a^*) dx + v_0 \int_{[0,1] \setminus \mathcal{X}} f(x|a^*) dx - c(a^*) \geq v_0 , \quad (IR)$$

$$a^* \in \arg \max_a \left\{ \int_{\mathcal{X}} w(x) f(x|a) dx + v_0 \int_{[0,1] \setminus \mathcal{X}} f(x|a) dx - c(a) \right\} , \quad (IC)$$

and

$$w(x) \geq v_0 \text{ for all } x \in \mathcal{X} . \quad (IIR)$$

Observe that unless the principal is implementing $a = 0$ and or \mathcal{X} has (Lebesgue) measure 0, $w(x) > v_0$ for a positive measure subset of \mathcal{X} . Thus, if the principal implements a nontrivial outcome $a > 0$ and \mathcal{X} has strictly positive measure, by *IC* and *IIR*,

$$\int_{\mathcal{X}} w(x) f(x|a^*) dx + v_0 \int_{[0,1] \setminus \mathcal{X}} f(x|a^*) dx - c(a^*) > v_0 \int_{\mathcal{X}} f(x|0) dx + v_0 \int_{[0,1] \setminus \mathcal{X}} f(x|0) dx = v_0 ,$$

and hence *IR* holds strictly. Consequently, the agent gets strictly positive rents. ■

F Non-genericity of the STP Contract

Here we establish that the “selling the project to the agent” contract is generically violated by the agent’s interim participation constraint. Formally,

Remark F.1. Identify a (t -action) decision problem as a point in Euclidean space $\mathbb{R}^{n \times t}$. For any decision problem with bounded payoffs $y \in \mathbb{R}^{n \times t}$ in which the principal optimally induces a non-degenerate distribution in the first-best benchmark and any neighborhood U of y , there exists a decision problem $y' \in U$ for which the principal cannot attain efficiency via the STP contract.

Proof. Suppose WLOG that in decision problem y , the principal can attain efficiency via a “sell the product to the agent” contract, and let $f_{\mathcal{H}}$ be the tangent hyperplane corresponding to optimal learning in the first-best benchmark. Let a be an action that is taken with strictly positive probability by the principal in the first-best benchmark. Construct decision problem y' by increasing the payoff from taking action a in some state l whose probability is nonzero at the belief at which the principal takes action a in her optimal learning by ε . Evidently, the tangent hyperplane corresponding to optimal learning in the first-best benchmark for decision problem y' , $f'_{\mathcal{H}}$ is not equal to $f_{\mathcal{H}}$. Otherwise, the principal’s payoff would be the same, which is a contradiction since her payoff at some on-path posterior has strictly increased. ■

References

Andreas Kleiner, Benny Moldovanu, Philipp Strack, and Mark Whitmeyer. Applications of power diagrams to mechanism and information design. *Mimeo*, 2022.