Existence of Optimal Mechanisms in Principal-Agent Problems*

Ohad Kadan,† Philip J. Reny,‡ and Jeroen M. Swinkels§

First Draft, October 2009
This Draft, February 2014

Abstract

We provide general conditions under which principal-agent problems admit mechanisms that are optimal for the principal. Our result covers as special cases pure moral hazard and pure adverse selection. We allow multi-dimensional types, actions, and signals, as well as both financial and non-financial rewards. Our results extend to situations in which there are ex-ante or interim restrictions on the mechanism, and allows the principal to have decisions in addition to choosing the agent's contract. Beyond measurability, we require no *a priori* restrictions on the space of mechanisms. It is not unusual for randomization to be necessary for optimality and so it (should be and) is permitted. Randomization also plays an essential role in our proof. We also provide conditions under which some forms of randomization are unnecessary.

1 Introduction

A principal wishes to incentivize an agent to behave optimally from the principal's point of view. The agent has private information summarized by his "type" and can take an action that the principal cannot directly observe. On the other hand, the principal can observe a signal whose distribution depends on the agent's type and action, and can choose (possibly randomly) a reward for the agent. This principal-agent setting is quite general, incorporating as special cases pure moral-hazard (one possible type), pure adverse selection (one possible action) and settings with both, as for example, a health insurance provider that worries not only about what the agent may know about her health but also about any actions the agent may take that affect her health.

The purpose of the present paper is to provide general conditions under which an optimal mechanism for the principal exists. By the revelation principle (Myerson 1982), it is without loss of generality to restrict attention to mechanisms of the following form. First, the agent reports his type. Given the reported type, the mechanism recommends (possibly randomly) an action for the agent to take and, as a function of the reported type and recommended action, specifies a contract -

*We thank John Conlon, Wojciech Olszewski, Max Stinchcombe, and Timothy Van Zandt for very helpful comments. Reny gratefully acknowledges financial support from the National Science Foundation (SES-0922535, SES-1227506), as well as the hospitality of MEDS at Northwestern University during his 2009-2010 visit.

†Olin Business School, Washington University in St. Louis

‡Department of Economics, University of Chicago

§Kellogg School of Management, Northwestern University

a mapping from signals into (distributions over) rewards. The agent, knowing the contract, chooses an action, and a signal is generated (with distribution determined by the true type and action of the agent). Given the signal, rewards are then generated according to the contract.

Existence of an optimal mechanism in such settings is a significant question. Principal-agent problems are of course entirely central to economics. And, it is possible in many principal-agent problems to derive useful predictions about the environment using a partial characterization of an optimal solution (e.g. through first-order conditions).[1] But, none of this is relevant unless one knows that an optimal solution exists.[2] So it is troubling that Mirrlees (1999) provides an example of a surprisingly simple economic setting (pure moral hazard, logarithmic utility, normally distributed signals) in which an optimal mechanism does not in fact exist, even in the case of pure moral hazard.[3]

To cover a wide array of economic settings, we permit types, actions, signals, and rewards to be multi-dimensional and we impose no particular order structure. We do not rely on (but permit) the usual structure of separability of utility in income and effort and $MLRP$ that permeates this literature. The signal space can describe multiple dimensions, such as which product the salesperson sold and what price was negotiated, and the reward space can similarly include whether the agent is promoted, how much she is paid, and the desirability of her office. The utility of both the agent and the principal can depend in a general manner (with appropriate continuity) on the type, action, signal, and reward. The principal and agent can have common or opposing interests. We permit the utility of the agent to be unbounded above, and similarly permit the utility of the principal to be unbounded below. Signal supports can vary with the agent's action and type and can contain atoms.

We make four substantive assumptions, each of which has a clear economic interpretation. First, we require that there is a limit to how severely the agent can be punished, and to how much profit the principal can earn.[4] Second, we require that if the utility of the agent can be made unboundedly high, it becomes very expensive for the principal to do so at the margin. Third, we introduce a new form of continuity of information that ensures that it does not become discontinuously more difficult to reward a compliant agent without also rewarding a deviating agent as the action of the compliant agent is varied. This condition is very mild. In particular (see Remark 1 below), it is satisfied whenever densities are continuous, which is common in applications. Finally, we require that if the principal observes a signal that is inconsistent with the reported type and the

[1] See Laffont and Martimort (2002) for a wealth of examples of the use of the principal agent model in moral hazard, adverse selection, and mixed settings. See Kadan and Swinkels (2013) for an example in which a partial characterization of equilibrium allows the derivation of several interesting comparative statics.

[2] In particular, first-order conditions are not applicable and, therefore, comparative statics results become virtually impossible to obtain.

[3] See also Moroni and Swinkels (2013) for a distinct class of counter-examples to existence that do not depend on an unbounded likelihood ratio, but rather on the behavior of risk aversion as utility diverges to negative infinity.

[4] With some separability, these restrictions can be substantially weakened.

recommended action, then there is a simple way to punish the agent, one that the principal can be sure minimizes the utility of the agent independent of his true type and action.

Previous work in this area (e.g., Holmström 1979; Page 1987,1991; Balder 1996, Section 3.2) differs from ours in two key respects. First, previous work typically presumes from the outset that, given the reported type, the recommended action, and the signal, the mechanism does not randomize over the agent's reward. In pure moral hazard settings Holmström (1979) shows that this presumption is innocuous when the principal and agent are risk averse and the agent's utility is additively separable in effort and rewards.[5] In general however, randomization over rewards can be necessary for optimality in both moral-hazard and adverse-selection settings (see e.g., Gjesdal 1982 and Arnott and Stiglitz 1988).[6] Second, the literature does not always employ an appropriate form of continuity of information. Consequently, to obtain a useful compactness property, the conventional approach has been to impose rather strong and ad hoc restrictions on the set of allowable contracts. We discuss this in detail in Section 4.

In contrast, we introduce a new continuity of information condition and permit randomization over the agent's reward. We allow randomization over the agent's reward first and foremost because, as noted above, it can be strictly optimal. But there are important technical advantages as well. In particular, such randomization helps eliminate the need to impose ad hoc and restrictive conditions on the space of contracts. The advantage so obtained is that, aside from measurability, our contract space is entirely unrestricted, resulting in a truly optimal solution to the principal's problem. None of the previous literature achieves this.

A second form of randomization, namely randomization over the contract assigned to the agent after he reports his type, is equally important for optimality.[7] Both Page (1991) and Balder (1996) allow randomization over the contract, as do we. However, in our setup, randomization over the contract arises entirely through the contract's dependence on the randomly selected recommended action, which, as discussed in Section 6.6, entails no loss of generality. As first recognized by Page (1991), randomization over the contract assigned to the agent allows one to take full advantage of a powerful sequential compactness result that originates with Komlos (1967) and was generalized by Balder (1990).

Our informational continuity assumption, combined with permitting randomization both over the recommended action/contract and over the reward assigned to the agent, plays a central role in allowing us to make use of Balder's (1990) sequential compactness result. Moreover, our informational assumptions are sufficiently general to permit both finite and continuous signal spaces, as well as signaling technologies in which distinct actions by the agent induce signal distributions whose supports only partially overlap, something the previous literature rules out.

[5] We provide a generalization of Holmström's result in Section 11.

[6] Thus the mechanisms that are shown to exist in Page (1987, 1991) and Balder (1996) can potentially be substantially improved upon by permitting randomization over rewards.

[7] An example showing that such randomizaiton can be necessary for optimality is provided in Section 3.

In some settings of course, both the principal and the agent are risk averse, and the payoffs to the agent are sufficiently separable in reward, action, and type that randomization over rewards is undesirable. Our existence result remains relevant even in these settings because the optimal mechanism whose existence it guarantees can, under the conditions described, be costlessly adjusted so that it employs only deterministic contracts. (see Section 11).

Finally, we show how our analysis allows one to study the question of optimality under additional constraints on the mechanism. For example, a regulator may insist that a health insurance provider offer plans that are acceptable to a certain fraction of a risk pool, or one that earns no more than a certain margin on a specified subset of risk types. These restrictions may well be ex-ante or interim in nature, in that they are restrictions on the whole mechanism, and not, for example, simply a restriction on the payments in any given situation. Our model and results (see Section 10) readily adapt to such restrictions. Using this tool, we show how to model an ex-ante participation constraint, a setting in which the principal can choose to exclude the agent in some circumstances, and a setting where the principal has decisions to take beyond the choice of the agent's contract.

The remainder of the paper proceeds as follows. Section 2 illustrates the role of informational continuity in our setting. Section 3 provides examples of the strict optimality of randomization. Section 4 discusses related literature. Section 5 presents our assumptions on spaces, payoffs, and information, defines a mechanism, and states our main result. Section 6 discusses the various assumptions and formalizations. Section 7 provides an example illustrating the main ideas of our proof. Section 8 proves our main result. Section 9 introduces a metric on the space of mechanisms and uses the proof of our main result to show that, under this metric, the principal's objective function is lower semicontinuous on a compact subset of the space of incentive compatible mechanisms that is sufficient for optimality. Section 10 takes advantage of the metric introduced in Section 9 to show how to adapt our model to ex-ante and interim restrictions on the set of mechanisms. In particular, this permits us to handle additional decisions for the principal. Finally, Section 11 discusses when optimal contracts can be chosen to be deterministic.

2 Informational Continuity - An Example

In this section, we present a pair of simple examples in which existence fails. Between them, they highlight the need for an appropriate notion of informational continuity, and serve as a launching point for our discussion of the literature.

Example 1 *Let the compact set of available actions be* $A = \{1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, ..., 0\}$, *let the set of signals be* $S = [0, 1]$ *and let the set of feasible rewards be* $R = [0, 3]$. *The agent's utility is equal to his reward* r *if he takes any action* $a < 1$ *and* $r + 1$ *if he takes action* $a = 1$.[8] *The principal's losses*

[8] Since $a = 1$ is an isolated action, there is no discontinuity here.

are a if the agent takes action $a \in A$ (rewarding the agent is costless to the principal).[9] *The signal s is uniform on $[0,1]$ if $a = 0$ or 1. If $a = \frac{1}{k} \in \{\frac{1}{2}, \frac{1}{3}, \frac{1}{4}, ...\}$, then the signal s has density 2 when it is in $S_k = [\frac{1}{k}, \frac{1}{2} + \frac{1}{k}]$ and has density 0 otherwise.*

For each k, the action $a = 1/k$ can be implemented by paying the agent $r = 3$ when a signal in S_k is observed and paying $r = 0$ otherwise, at loss $1/k$ for the principal. But, $a = 0$ cannot be implemented because $a = 1$ is strictly preferred by the agent to $a = 0$ regardless of the contract offered. Hence, there is no optimal mechanism in this setting.

The payoff functions in this example are continuous and all spaces are compact. The culprit driving non-existence is that information changes discontinuously at $a = 0$. In particular, the distribution of signals conditional on the agent choosing action a fails to be continuous in the weak* topology at $a = 0$, since, for example, the probability of the open set $(1/2, 1)$ jumps from near 0 to $\frac{1}{2}$.

In view of this, it is no surprise that the extant literature assumes, at a minimum, continuity of information in the weak* topology, thereby ruling out the above example. But, as the next example illustrates, this is not enough.

Example 2 *Modify Example 1 only in that S_k is the subset of $[0,1]$ where the k-th digit in the binary expansion of s is 0. Now, as $a \to 0$, the signal distribution not only converges in the weak* topology to the uniform distribution, the convergence is strong enough so that the informational assumptions of Page (1987,1991) are satisfied. But, exactly as before, $a = \frac{1}{k}$ can be implemented by paying $r = 3$ on S_k and paying $r = 0$ otherwise, yet $a = 0$ cannot be implemented.*

For each k, these two examples are fundamentally the same, just with the signals re-shuffled in an irrelevant way. A satisfactory informational assumption should thus rule both of them out. To see what such an assumption might look like, let $f(s|a)$ denote the density of the signal s given the action a, and note that on S_k, $\frac{f(s|1)}{f(s|\frac{1}{k})} = \frac{1}{2}$. This ratio is critical, because it determines how difficult it is to reward the agent for choosing $a = \frac{1}{k}$ without also making $a = 1$ attractive. But, $\frac{f(s|1)}{f(s|0)} = 1$ on $[0,1]$, and hence this ratio is discontinuous at $a = 1$.

The informational continuity assumption we introduce in Section 5 (see Assumption 10) rules out such upward jumps in how difficult it is to reward one action without making another action more attractive. Its motivation is explored in Section 6 and 7. For many applications, the full generality of our informational continuity assumption will not be necessary and the following remark will suffice.

Remark 1 *Assume the signal, action, and type spaces are Euclidean, and that the distribution of the signal s as a function of the agent's action a and type t is given by a jointly measurable density*

[9] We will think about the principal as minimizing an expected loss rather than (as is completely equivalent) maximizing an expected gain.

$f(s|a,t)$ that, for each t, is jointly continuous in (s,a) at every (s_0,a_0) such that $f(s_0|a_0,t)>0$. Our informational assumptions are then satisfied.[10]

The conditions of Remark 1 are satisfied in many commonly used models.

3 Randomization

The principal has two potentially useful opportunities for randomization. First, after the agent reports his type, the principal can randomize over the recommended action which, because the contract might depend on this recommendation, can have the effect of randomizing over the agent's contract.[11] Second, after the agent chooses his action and a signal is generated, the contract might randomize over the agent's reward. The purpose of this section is to remind the reader of how each form of randomization can strictly benefit the principal.

In our first example,[12] randomizing over rewards is advantageous to the principal because both the principal's losses and the agent's aversion to risky rewards increase with the agent's action.

Example 3 ***Randomization over rewards.*** *There is pure moral hazard. The signal is null. The action space is $A=[0,1]$. The set of feasible rewards is $R=[0,4]$. When the agent takes action a and the reward is r, the agent's utility is $a(1+\sqrt{r})+(1-a)r$ and the principal's loss is $r+3a$. The unique optimal contract is to pay each of 0 and 4 with probability $1/2$. This implements $a=0$ at expected loss 2.*

To see this, begin by restricting the agent's actions to only $a=0$ and $a=1$. The optimal (i.e., expected loss minimizing) way to implement $a=1$ is to pay $r=0$ with probability 1. The question then comes down to how to optimally implement $a=0$. Given that the agent is risk neutral when $a=0$, but risk averse when $a=1$, this must involve rewards only of $r=0$ and $r=4$ (else by applying a mean-preserving spread to the reward distribution the principal can costlessly increase the risk to $a=1$ and relax the incentive constraint). Letting p be the probability of a reward of 4, and solving $(1-p)\,0+p4\geq(1-p)\,1+p\left(1+\sqrt{4}\right)$ gives $p=\frac{1}{2}$ as the optimal solution.[13]

Note that in this example, the reward and action spaces are convex, and the agent's utility function is jointly concave (and the principal's loss function is jointly convex) in rewards and actions. So, the need to randomize here is not about convexifying choice spaces or preferences.[14]

[10] Note that there is no requirement here of a constant support or that the density goes continuously to zero at the boundaries of its support for any given (a,t).

[11] In fact, randomization over the recommended action that does not induce randomization over the contract is never necessary for optimality. See Section 6.6 for details.

[12] A slight simplification of an example due to Gjesdal (1982) p.382.

[13] To see that this mechanism is optimal with the full action space, note first that because the signal is null, any randomization over R results in payoffs to the agent and to the principal that are linear in a. Since for each r the principal strictly prefers $a=0$ to $a=1$, the principal will thus never implement anything other than $a=0$ or $a=1$.

[14] Such examples are easy to generate.

In the next example, randomization over the recommended action induces randomization over the agent's contract, and this strictly improves the agent's incentives to truthfully report his type. We describe the example here in terms of discrete type and action spaces. In Appendix I, this is extended to convex spaces.

Example 4 ***Randomization over recommended actions.*** *There are two equiprobable types t_1 and t_2, and two actions a_1 and a_2. Actions are observable. Rewards are in $[0, \infty)$. Payoffs are described by the following matrix, where in each cell, the top left number is the agent's utility, and the bottom right number is the loss to the principal if the given type takes the given action and the reward is r.*

	t_1	t_2
a_1	r r	$1+r$ $2+r$
a_2	$1+r$ $2+r$	r r

In this example, the principal would like to, but cannot, ensure that t_1 chooses a_1 and t_2 chooses a_2. Indeed, if r_1 is the agent's reward for choosing a_1 after reporting t_1, and r_2 is the reward for choosing a_2 after reporting t_2, then for t_1 not to want to imitate t_2 we need $r_1 \geq 1 + r_2$, while for t_2 not to want to imitate t_1 we need $r_2 \geq 1 + r_1$.[15]

The following incentive compatible mechanism, which involves randomization over recommended actions, is uniquely optimal for the principal. If t_i is reported, then a lottery occurs. Half of the time the principal recommends a_i and pays the agent 1 if a_i is observed and pays 0 otherwise. The other half of the time the principal recommends $a_j \neq a_i$ and pays the agent 0 regardless of the action chosen.[16]

These examples illustrate that a general solution to the principal-agent problem requires the consideration of mechanisms that allow randomization over recommended actions and rewards. In contrast, restricting attention to deterministic mechanisms is with significant loss of both generality and optimality.[17]

But, randomization is not just the right approach from an economic point of view. Permitting randomization leads naturally to a space of mechanisms whose averages are themselves

[15] Randomizing over rewards for any given action does not help, as the participants are risk neutral and the set of feasible rewards is convex. Thus, we will restrict attention to mechanisms with deterministic rewards here.

[16] Adding the two truthtelling constraints (assuming compliance) implies that any IC mechanism must be at least as likely to recommend a_i after a report of t_j than after a report of $t_i \neq t_j$; and because t_i can obtain utility 1 or more by choosing $a_j \neq a_i$, the reward for compliantly choosing a_i must be at least 1. The principal's expected loss must therefore be at least 3/2, which the given mechanism acheives exactly. Uniqueness can then be shown rather easily.

[17] While economically meaningful conditions exist under which the principal can fully optimize without randomizing over rewards (see Section 11), Example 3 illlustrates that they are not without loss of generality. We know of no such conditions that ensure that randomization over recommended actions is not needed for optimality.

mechanisms.[18] This is important because it permits the use of results concerning the existence of pointwise-convergent subsequences of the averages of sequences of mechanisms, though not necessarily the sequences themselves. The significance of pointwise convergence is that it renders payoffs continuous at the limit. Establishing the sequentially compact nature of our space of mechanisms is a central technical contribution of this paper.

4 Literature

This paper offers a general existence result that does not rely on *a priori* restricting the set of contracts that are available to the principal. Several earlier papers provide existence results by imposing such restrictions, whose purpose is to compactify the space of contracts.

A typical first such restriction is to insist that contracts are deterministic. That is, a contract must be a function mapping the signal into a reward, and not a lottery over rewards.[19] This restriction, by itself, would entail no loss of generality were the mechanism permitted to randomize over the agent's contract after the agent chooses his action. However, the previous literature assumes that the agent is informed of the contract before taking his action. This combination of assumptions sometimes entails a substantial loss of generality and optimality (see e.g. Example 3).

Much of the previous literature assumes that the agent's reward is a real-valued quantity such as salary. Consequently, a contract is a real-valued function of the signal. Additional restrictions are then placed on this real-valued function space to ensure compactness. These restriction too can lead to suboptimality.

For example, in addition to assuming a deterministic contract space, Holmström (1979) compactifies the space of contracts by imposing upper and lower bounds on salary and by restricting the set of feasible contracts to those whose variation is bounded by some prespecified bound. Page (1987,1991), first for the pure moral hazard problem, and then for a problem that also includes adverse selection, proves existence in a model with substantively more general actions, signals, and preferences than in Holmström.[20] Similar to Holmström, Page imposes the *a priori* restriction that contracts are deterministic and come from a compact space (e.g., functions that are uniformly bounded and either equicontinuous, monotone, or of bounded variation). Balder (1996) extends and unifies a variety of these results, but follows Page by assuming *a priori* that the feasible set of contracts happens to be deterministic and compact.[21] The difficulty with these various restrictions

[18] By "naturally" we mean that compactness arises without the need to impose ad hoc restrictions on the set of contracts that are available for the principal as previously done in the literature.

[19] The agent is typically not assumed to be risk neutral with respect to rewards.

[20] See also Chen and Page (2013) for a similar approach.

[21] Although Balder follows Page and uses deterministic contracts, he needn't have done so. Indeed, for special cases of our model (most critically, with the space of rewards compact, bounded payoffs, and with the distribution over signals for each type and action absolutely continuous with respect to some fixed distribution over signals), one can obtain existence of a fully optimal mechanism by combining Balder's existence results (Propositions 3.1 and 3.2) with our idea that a contract with randomization over rewards is tightly linked to a joint distribution over the space of

is that they are ad hoc – there is no guarantee that otherwise reasonable contracts that do not meet the variation bound or the bounds on rewards, etc., would not be substantially better for the principal (see. e.g., Example 5). Thus, the restrictions that have routinely been imposed on contracts to obtain compactness call into question the optimality of the mechanisms whose existence is established.

It is important to note that the previous literature must forcibly restrict contracts to a compact set in part because a suitable form of continuity of information has been lacking. Indeed, the reason that Example 2, which satisfies the information assumptions of Page (1987) but does not satisfy our continuity of information condition, does not constitute a counterexample to his main result is that no set of contracts that includes, for every k, those that are 3 on S_k and 0 elsewhere is compact in the topology of almost everywhere pointwise convergence. Hence, for any set of contracts satisfying Page's compactness condition, an optimal contract exists (within that compact set) only because, for that set of contracts, there is a $\hat{k}$ such that for $k > \hat{k}$, $a = \frac{1}{k}$ is simply not implementable, and one is left choosing over the finite set of actions that remain.

Based on this example, one might argue that bounding the variation of contracts, or restricting to a finite dimensional closed and bounded subset, is reasonable since doing so has little effect in terms of the principal's payoff. But this is not in general true. In the following example, the unique optimal contract has unbounded variation, and any other contract (e.g., any with bounded variation) yields losses to the principal that are bounded away from the optimum.[22]

Example 5 *There are two equiprobable types, t_1 and t_2, and two actions, a_1 and a_2. The set of signals is $S = [0,1]$ and the set of rewards is $R = [0,2]$. The agent's utility is $u(t_1, a_2, r) = 2$, $u(t_2, a_2, r) = 1$, and $u(t_1, a_1, r) = u(t_2, a_1, r) = r$, and the principal's losses are $l(t_1, a_1) = l(t_2, a_2) = 0$, and $l(t_1, a_2) = l(t_2, a_1) = 1$. The signaling technology is as follows. If t_1 chooses a_1, the density of the signal is 2 on the set $S_1 = \left[0, \frac{1}{4}\right] \cup \left[\frac{1}{2}, \frac{5}{8}\right] \cup \left[\frac{3}{4}, \frac{9}{16}\right] \cdots = \cup_{j=0}^{\infty} \left[\frac{2^j-1}{2^j}, \frac{2^j-1}{2^j} + \frac{1}{2^{j+2}}\right]$, but if t_2 chooses a_1 or if either type chooses a_2 the signal is uniform on $[0,1]$. The unique optimal contract pays $r = 2$ if the signal is in S_1 and pays 0 otherwise. This incentivizes t_i to choose a_i and achieves an expected loss of 0. Any other contract either fails to get t_1 to choose a_1, or fails to get t_2 to choose a_2 and so generates an expected loss of at least $1/2$.*

But even if with some extra conditions one could show that a particular compact set of contracts did achieve ε-optimality, it would not be clear what the point of such an exercise would be, as the existence of ε-optimal solutions is trivial to establish under more general conditions (e.g. a lower bound on the principal's expected losses). Moreover, while the optimal contract within a pre-specified compact set of contracts may come close in payoffs to the optimal contract, it might

signals and rewards.

[22] The example is robust in the sense that, for any given bound on the variation of contracts, a small enough perturbation of the parameters of the example does not affect the conclusion.

involve behavior by the principal that is very far from the (in Example 2 trivial) behavior that the principal should adopt.

Grossman and Hart (1983) establish existence of an optimum in a pure moral hazard problem without ad hoc restrictions on contracts. They do this by restricting attention to a finite set of signals each of which occurs with probability bounded away from zero regardless of the agent's action. With a continuum of signals, Carlier and Dana (2005) and Jewitt, Kadan and Swinkels (2008) solve the existence problem in a pure moral hazard setting while avoiding ad hoc restrictions on contracts by assuming that effort is one-dimensional, that likelihood ratios are monotone and bounded, and that the first-order approach (Mirrlees (1976), Rogerson (1984), Jewitt (1988), Sinclair-Desgagne (1994), and Conlon (2009)) is valid. All three papers require signals and rewards to be one-dimensional and the principal's losses to be additively separable in them, and none permits the agent to possess private information. Finally, Kahn (1993) shows existence in a pure adverse selection problem, avoiding restrictions on contracts, but relying on restrictions on the set of types, and on the distributions and utilities considered.

5 The Model and Main Result

Let us now formalize the model and state our main result. We will provide examples, motivation, and discussion in the next section, especially for the informational conditions.

5.1 Preliminaries

A Polish space is any complete, separable, metrizable topological space. For any Polish space, X, denote by $\Delta(X)$ the space of Borel probability measures on X endowed with the topology of weak convergence (the weak* topology). When we say that a subset of X is measurable, we will always mean Borel measurable. The symbol δ_x denotes the Dirac measure placing probability one on $x \in X$. All integrals in the sequel are well-defined, though their values may be infinite, because under our assumptions the functions being integrated will always be nonnegative and measurable.

Suppose Y is also Polish, and that $\nu(\cdot|x) \in \Delta(Y)$ for each $x \in X$. If for every measurable subset C of Y, $\nu(C|x)$ is a measurable real-valued function of x on X, then we say that ν is a *regular conditional probability (r.c.p.) from X to Y*. For any $B \subseteq X \times Y$ and any $y \in Y$, let $B_y = \{x \in X : (x, y) \in B\}$ denote the "slice through y."

5.2 Assumptions

We maintain the following assumptions throughout.[23]

Assumption 1 *The set of types, T, is a nonempty Polish space. The prior is $H \in \Delta(T)$.*

[23] We remind the reader that Remark 1 provides simple sufficient conditions for Assumptions 8 - 10.

Assumption 2 *The set of actions, A, is a nonempty compact metric space.*

Assumption 3 *The set of signals, S, is a nonempty Polish space.*

Assumption 4 *The set of all possible rewards, R, is a nonempty Polish space. The set of feasible rewards can depend on the signal and is captured by Φ, a closed subset of $R \times S$, where Φ_s, the slice through s, is the set of feasible rewards given signal s.*

Assumption 5 *The agent's von Neumann-Morgenstern utility function is $u : R \times S \times A \times T \to \mathbb{R}$, and the principal's von Neumann-Morgenstern loss (disutility) function is $l : R \times S \times A \times T \to \mathbb{R}$. Both functions are bounded below, without loss of generality by 0; $u(r, s, a, t)$ is continuous in (r, s, a) for each t; $l(r, s, a, t)$ is lower semicontinuous in (r, s, a) for each t; and both functions are measurable in t for each (r, s, a).*

Assumption 6 *For every $t \in T$, every $c \in \mathbb{R}$ and every compact subset Y of S, $\{(r, s, a) \in \Phi \times A : s \in Y \text{ and } l(r, s, a, t) \leq c\}$ is compact.*

Assumption 7 *If $u(r, s, a, t) \to \infty$ for some sequence (r, s, a, t) in $\Phi \times A \times T$, then $l(r, s, a, t)/u(r, s, a, t) \to \infty$.*

The signal technology is given by P, an r.c.p. from $A \times T$ to S, where $P(\cdot|a, t) \in \Delta(S)$ is the probability measure over the signal space S when the agent takes action $a \in A$ and his type is $t \in T$. We write $P_{a,t}$ instead of $P(\cdot|a, t)$ whenever convenient.

Assumption 8 *If $a_n \to a$, then $P_{a_n,t} \to P_{a,t}$ for every $t \in T$ where convergence of measures is in the weak* topology.*

Assumption 9 *There is a collection $\{S_x\}_{x \in A \times T}$ and a nonnegative measurable function ξ on $S \times A \times T \times A \times T$ such that*

(i) for each $x \in A \times T$, S_x is a measurable subset of S with $P_x(S_x) = 1$,

(ii) $\{(s, x) \in S \times (A \times T) : s \in S_x\}$ is measurable, and

(iii) for every $x, x' \in A \times T$, and every measurable subset B of S,

$$P(B \cap S_{x'}|x) = \int_B \xi(s, x, x') dP(s|x').$$

Assumption 9 (iii) implies that the restriction of P_x to $S_{x'}$ is absolutely continuous with respect to $P_{x'}$ and, in particular, that $\xi(\cdot, x, x')$ is (a version of) its Radon-Nikodym derivative with respect to $P_{x'}$. The measurability of ξ requires that versions of these Radon-Nikodym derivatives can be chosen so that they vary in a mildly regular fashion with x and x'. Because $\xi(s, x, x')$ can be interpreted as the likelihood of $s \in S_{x'}$ conditional on x relative to its likelihood conditional on x', we take the mnemonically convenient step of writing $\xi_{x/x'}(s)$ for $\xi(s, x, x')$.

Assumption 10 *For all $(a,t),(a',t')$ and every $a'_n \to a'$, there is a sequence $a_n \to a$ such that for $P_{a',t'}$ a.e. $s \in S$, the inequality*

$$\underline{\lim}_n \xi_{a_n,t/a'_n,t'}(s_n) \geq \xi_{a,t/a',t'}(s) \tag{1}$$

holds for every sequence $s_n \in S_{a'_n,t'}$ converging to s.

Assumption 11 *Let $S^0 = \cap_{a,t} S_{a,t}$ be the set of signals that are consistent with every type and every action.*[24] *There is a nonnegative $u_* : S \times A \times T \to [0,\infty)$ with $u_*(s,a,t)$ continuous in (s,a) for each t, and a measurable $r_* : S \to R$ such that for all $(r,s,a,t) \in \Phi \times A \times T$, $u(r,s,a,t) \geq u_*(s,a,t)$ and if $s \in S_{a,t} \backslash S^0$ then $r_*(s) \in \Phi_s$ and $u_*(s,a,t) = u(r_*(s),s,a,t)$.*

5.3 Mechanisms and the Principal's Problem

Fix functions u_* and r_* satisfying Assumption 11. Our space of mechanisms is as follows.[25]

Definition 1 A *mechanism* is any (α,κ) such that α is an r.c.p. from T to A and κ is an r.c.p. from $S \times A \times T$ to R, and where for all $(s,a,t) \in S \times A \times T$, $\kappa(\Phi_s|s,a,t) = 1$ if $s \in S_{a,t}$ and $\kappa(\cdot|s,a,t) = \delta_{r_*(s)}$ if $s \notin S_{a,t}$. Denote the set of all mechanisms by M.

The interpretation is that for any type t' reported by the agent when his type is t, the mechanism (α,κ) generates an action a' according to $\alpha(\cdot|t')$ to recommend to the agent. After receiving the recommended action a', the agent chooses an action a from A. Finally, given the signal s generated according to $P(\cdot|a,t)$ the mechanism generates the agent's reward $r \in \Phi_s$ according to $\kappa(\cdot|s,a',t')$. If s is inconsistent with (a',t'), the agent receives $r_*(s)$.

Definition 2 A mechanism (α,κ) is *incentive compatible* if for H-almost every $t \in T$, the following inequality holds for every $t' \in T$.[26]

$$\int_{R\times S\times A} u(r,s,a,t)\, d\kappa(r|s,a,t)\, dP(s|a,t)\, d\alpha(a|t) \geq \int_A \left(\sup_{a\in A} \int_{R\times S} u(r,s,a,t)\, d\kappa(r|s,a',t')\, dP(s|a,t) \right) d\alpha(a'|t')$$

The left-hand side of the inequality is the utility to the agent of type t from reporting his true type to the mechanism and taking the recommended action, while the right-hand side is the largest utility the agent of type t can achieve by reporting that his type is t', and then, as a function of the action a' recommended by the mechanism, choosing any action a.

[24] It is possible that S^0 is empty.

[25] Section 6.6 explains why this is the appropriate space of mechanisms.

[26] The right-hand side integral is well-defined because the nonnegative function in parentheses is measurable in a'. See Footnote 42.

Letting M^* denote the set of incentive compatible mechanisms, the principal's problem is

$$\min_{(\alpha,\kappa)\in M^*} \int_{R\times S\times A\times T} l\left(r,s,a,t\right) d\kappa\left(r|s,a,t\right) dP(s|a,t) d\alpha\left(a|t\right) dH\left(t\right). \tag{2}$$

Note that an outside option that is always available to the agent can, as usual, be modeled by including it as an action in A. Moreover, and again as usual, one can model that the principal can always force the agent to take the outside option by including the reward "take your outside option" in Φ_s for every signal s. In Section 10, we show that the model extends immediately to an ex-ante individual rationality constraint of the form

$$\int_{R\times S\times A} u\left(r,s,a,t\right) d\kappa\left(r|s,a,t\right) dP\left(s|a,t\right) d\alpha(a|t) \geq u_0.$$

5.4 The Main Existence Result

Theorem 1 *If Assumptions 1-11 hold, and if the set M^* of incentive compatible mechanisms is nonempty,[27] then the principal's problem (2) possesses a solution.*

6 Examples and Discussion

Before we turn to proving Theorem 1, we provide some discussion and motivation for the various assumptions and definitions we have made.

6.1 Assumptions 1-4

The conditions on R, S, A and T are extremely mild. Simple examples of Polish spaces include $\mathbb{R}$, $[0,1]$, $\mathbb{Z}$, and any finite or countable product thereof. Hence, all spaces can include multiple dimensions, some of which may be discrete and some of which may be continuous. Only A is restricted to be compact. The definition of Φ allows considerable flexibility regarding the rewards that are feasible as a function of the observed signal. Thus for example, setting $\Phi_s = [m, \infty)$ for all s captures a limited liability constraint.

6.2 Assumption 5

The continuity assumptions on $u(\cdot)$ and $l(\cdot)$ are standard. The assumption that utility is bounded below is critical in ruling out the Mirrlees (1999) and Moroni-Swinkels (2013) examples. It and the assumption that losses are bounded below are substantive in some settings, but reasonable in many others. Further, as the next two examples illustrate, with some separability many cases with

[27] Section 6.7 provides weak conditions under which M^* is nonempty.

unbounded payoffs are covered by our formulation so long as expected payoffs are appropriately bounded.

Example 6 ***Unbounded Losses.*** *$S = (-\infty, \infty)$, $R = [0, \infty)$, and $\Phi_s = R$ for all s. The principal is risk neutral, receives revenue s, and pays compensation r. The principal's loss function, $l(r,s,a,t) = r - s$, is unbounded above and below. However, if $\zeta(a,t) = \int_S s dP(s|a,t)$ is continuous and bounded above by some $M < \infty$, then defining l instead by the nonnegative function $l(r,s,a,t) = r + M - \zeta(a,t)$ gives the principal the same incentives over expected losses.*

Example 7 ***Unbounded Utility.*** *$S = (0, \infty)$, $R = [0, \infty)$, and $\Phi_s = R$ for all s. If the agent's utility is $u(r,s,a,t) = \log(r+s) + w(a,t)$, where $w(a,t) \geq 0$, then utility is unbounded above and below. However, if $\zeta(a,t) = \int_S \log s dP(s|a,t)$ is continuous and bounded below by $-M$, then defining u instead by the nonnegative function $u(r,s,a,t) = \log(r+s) - \log s + \zeta(a,t) + M + w(a,t)$ gives the agent the same incentives.*

6.3 Assumptions 6 and 7

If, for example, the reward space is Euclidean, Assumption 6 implies that as rewards become unbounded, so does the loss to the principal.

Assumption 7 says that the losses to the principal per util provided are unbounded above when they provide the agent with arbitrarily high utility. If the agent's utility is bounded, Assumption 7 is satisfied trivially. The following example shows how to embed the utility and loss functions most often used in applied settings into our setting.

Example 8 ***Canonical Utility and Loss Functions.*** *Let S, R, Φ, and l be as in Example 6 and let $u(r,s,a) = v(r) - c(a)$ where $c(\cdot)$ is continuous and v is differentiable with $\lim_{r\to\infty} v'(r) = 0$ (as is true for risk-averse utility functions typically used in practice). Then, Assumption 7 is satisfied.*

If, in the previous example, $v(r) = r - \frac{1}{r}$, Assumption 7 would fail.

6.4 Assumptions 8-10

Assumption 8 is extremely mild. Assumption 9 (iii) is much weaker than assuming that P_x itself is absolutely continuous with respect to $P_{x'}$, since, for example, 9 (iii) permits P_x and $P_{x'}$ to be mutually singular.

To understand Assumption 10**,** note first that while Example 2 satisfies Assumption 8, it fails 10. In particular, let $a = 1$, $a' = 0$, and let $a'_n = \frac{1}{n}$. Any sequence $a_n \to 1$ is constant at $a_n = 1$ after some point. Fix any given $\hat{s} \in [0,1]$. For each n, and for $P_{\frac{1}{n}}$ a.e. $s \in S_{\frac{1}{n}}$, it must be that $\xi_{1/\frac{1}{n}}(s) = \frac{1}{2}$. Hence, for some s_n with binary expansion that agrees with $\hat{s}$ to the first $n-1$ digits,

$\xi_{1/\frac{1}{n}}(s_n) = \frac{1}{2}$. But then Assumption 10 is violated, since $\xi_{1/0}(s) = 1$ for almost all s, and $\hat{s}$ was arbitrary.

The following are some simple examples where Assumption 10 is satisfied.

Example 9 ***Example 1 revisited.*** *Assume S is Euclidean and $P_{a,t}$ admits a density $f(s|a,t)$ that, for each t, is jointly continuous in (s,a) at every (s_0,a_0) such that $f(s_0|a_0,t) > 0$. Then, setting $\xi_{a,t/a',t'}(s)$ equal to $\frac{f(s|a,t)}{f(s|a',t')}$ whenever $f(s|a',t') > 0$ and equal to zero otherwise, Assumption 10 is satisfied, because where $f(s|a',t') > 0$, $\xi_{a,t/a',t'}(s)$ is lower semicontinuous in (s,a,t,a',t').*

Example 10 ***Moving support.*** *As a particular instance of the previous example, let $A = [0,1]$, $S = [0,2]$, $T = \{t\}$, and let P_a be uniform on $[a, a+1]$. Set $S_a = (a, a+1)$, and set $\xi_{a/a'}(s) = 1$ for $s \in (a', a'+1) \cap (a, a+1)$, and 0 elsewhere.*

Note in the last two examples how useful it is that S_x need not be closed and so can differ from the support of P_x by a set having P_x measure zero.

Example 11 ***Moving atom.*** *Let $A = S = [0,1]$, $T = \{t\}$, and let P_a be the Dirac measure placing mass one on $s = a$. Then Assumption 10 is satisfied by taking $\xi_{x/x'} = 1$ when $x = x'$, and 0 when $x' \neq x$.*

Example 12 ***Discrete signal distributions.*** *Suppose that S is a finite set. Assumption 8 is equivalent to $P(s|a,t)$ being a continuous function of a for each $(s,t) \in S \times T$. Since for finite S, the Radon-Nikodym derivative at s is just the ratio of probabilities assigned to s, Assumption 10 holds by continuity since it is required to hold only at those signals s such that $P(s|a',t') > 0$.*

Each of these examples satisfies the more stringent (and often easier to check) condition that, $t, t' \in T$, $\xi_{a,t/a',t'}(s)$ is lower semi-continuous on $S \times A \times A$. The following example pushes Assumption 10 much harder.

Example 13 *The agent is equally likely to be one of two types t_1 or t_2 and, after observing his type, can choose any action, a from the set $\{out\} \cup [0,1]$. After an action is chosen, a signal from the set $[0,1] \cup \{out_1, out_2\}$ is generated as follows. When t_i chooses out, the signal is out_i, so the agent's type is revealed. When t_1 chooses $a \in (0,1]$ the signal is uniform on $[0,a]$ with constant density $\frac{1}{a}$. When t_2 chooses $a \in (0,1]$, the signal's support is again $[0,a]$ but the signal's density is $\frac{1}{a} + 1$ if $s \in [0, a/2]$ and is $\frac{1}{a} - 1$ if $s \in (a/2, a]$. Finally, when either type chooses action $a = 0$, the signal is zero.*

To see that this example satisfies Assumption 10, note that since densities are well-defined for $a > 0$, we can take the Radon-Nikodym derivative to be their ratio. When $a' = a = 0$ and $a'_n \to 0$ with $t' \neq t$ (the only non-obvious case), it is straightforward that setting $a_n = a'_n$ satisfies

Assumption 10.[28,29] If one instead takes $a_n = 2a'_n$, then $f(s_n|a_n, t_2)/f(s_n|a'_n, t_1) \to \frac{1}{2}$, but at the limit when $a = 0$, the Radon-Nikodym derivative jumps up to 1. So, in this example, it is important that Assumption 10 allows the choice of the sequence a_n to be tailored to the particular sequence a'_n.

6.5 Assumption 11

If the signal $s \notin S_{a',t'}$ occurs after the agent reports t' and is asked to take action a', then the principal is certain that the agent was either dishonest about his type or failed to take the recommend action. Assumption 11 allows the principal to be sure that in such a case, by rewarding the agent with $r_*(s)$, the agent receives his worst case utility $u_*(s, a, t)$ no matter what the true (a, t) is. Assumption 11 is most restrictive when S^0 is empty and is completely unrestrictive when $S_{a,t}$ is independent of $(a, t) \in A \times T$.[30]

In many cases of applied interest, Assumption 11 holds for the simple reason that there is a reward that is always available and that the agent uniformly finds worse than any other feasible reward.

Example 14 ***Worst Rewards.*** *Assume there is* $\underline{r}$ *such that for each* $(r, s, a, t) \in \Phi \times A \times T$, $\underline{r} \in \Phi_s$ *and* $u(\underline{r}, s, a, t) \leq u(r, s, a, t)$. *Then Assumption 11 holds by setting* $r_*(s) = \underline{r}$ *for every* $s \in S$.

The following is an example where there is no "natural" worst reward. In such a case, Assumption 11 still holds (since it becomes completely unrestrictive) if the support is not moving, i.e., if $S_{a,t}$ is constant.

Example 15 ***No Worst Reward.*** *Let* $T = \{t\}$, *let* $A = R$ *be a finite set, and for each* $s \in S$ *let* $u(r, s, a, t) = 1$ *if* $a = r$ *and zero otherwise.*

6.6 The Definition of a Mechanism

We arrive at our space of mechanisms via a straightforward application of the revelation principle (Myerson 1982). The only question might be why we do not permit the mechanism, after receiving the type-report t' and recommending the action a', to randomize over the contract and require the agent to choose his action before knowing which contract is realized. The answer is that this is

[28] Note that, for example, setting $a_n = (a'_n)^2$ would work equally well.

[29] Example 13 further illustrates why it is important to allow $S_{x'}$ to differ by a $P_{x'}$-measure 0 set from the support of $P_{x'}$. Assumption 9 (iii) would fail if $S_{a,t}$ were always defined to be the support $[0, a]$ of $P_{a,t}$ because $P_{a',t'}(s = 0) = 0$ when $a' > 0$ and $P_{a,t}(s = 0) = 1$ when $a = 0$. In contrast Assumption 9 (iii) can be satisfied in our setup by defining $S_{a,t} = (0, a]$ for $a > 0$.

[30] In this case, $S_{a,t} \backslash S^0 = \varnothing$, and so, since u is bounded below by the constant function zero we may set u_* identically to zero and, since r_* need not specify a feasible reward in this case, we may set r_* to be any constant function of the signal.

never required for optimality because the result of that randomization can always be mimicked by randomizing over rewards instead. For example, a fifty-fifty lottery over the contracts $\kappa_1(\cdot|\cdot, a', t')$ and $\kappa_2(\cdot|\cdot, a', t')$ is equivalent for both the principal and the agent (regardless of the agent's type) to the contract $\bar{\kappa}(\cdot|\cdot, a', t')$ defined by $\bar{\kappa}(\cdot|s, a', t') = \frac{1}{2}\kappa_1(\cdot|s, a', t') + \frac{1}{2}\kappa_2(\cdot|s, a', t')$ for every $s \in S$. Consequently, for a type t agent who reports t', any uncertainly about the contract $\kappa(\cdot|\cdot, a', t')$ that he will be assigned is captured entirely by the uncertainly over the recommended action a' that the mechanism will specify.

Insisting that a mechanism satisfy $\kappa(\cdot|s, a', t') = \delta_{r_*(s)}$ if $s \notin S_{a',t'}$ is innocuous because any κ that does not satisfy this restriction can (under the measurability conditions in Assumptions 9 and 11) be redefined so that it does and, under Assumption 11, the redefined mechanism gives the agent at least as strict incentives to be compliant, without affecting the utility of a compliant agent. The advantage of insisting that a mechanism satisfy $\kappa(\cdot|s, a', t') = \delta_{r_*(s)}$ if $s \notin S_{a',t'}$ is purely technical. It ensures the measurability of the integrand on the right-hand side of the inequality in Definition 2 above.[31]

It would have been essentially equivalent to define incentive compatibility in the weaker sense that, for some H-measure one set of types T^0, the inequality in Definition 2 holds only for every $t, t' \in T^0$. This is because one can, for any $t^0 \in T$, always redefine (α, κ) on $T \backslash T^0$ to be equal to its value at t^0 (i.e., any report t' outside T^0 is treated as if the report had been t^0 instead), thereby satisfying the inequality exactly as given in Definition 2, which we find more intuitive.

6.7 Statement of the Theorem

In order to apply Theorem 1, the set M^* of incentive compatible mechanisms must be nonempty. The following result covers a wide range of applications.

Proposition 1 *Suppose that Assumptions 1-5 and 8 hold. Then, M^* is nonempty if either one of the following conditions holds.*

(i) *There is a continuous function, $r_0 : S \to R$ such that $r_0(s) \in \Phi_s$ for every $s \in S$.*

(ii) *There is a compact subset C of R such that $\Phi_s \cap C$ is nonempty for every $s \in S$.*

Both (i) and (ii) hold in the common situation in which at least one fixed reward, r^* say, is always feasible, i.e., $\exists r^* \in R$ such that $r^* \in \Phi_s$ for every $s \in S$. But neither condition implies the other. For example, if $S = [0, \infty)$, then only (i) is satisfied if $\Phi_s = [s, s+1]$, while only (ii) is satisfied if $\Phi_s = 0$ for $s \leq 1$, $\Phi(1) = \{0, 1\}$ and $\Phi_s = 1$ for $s > 1$.

The proof of Proposition 1, which can be found in Appendix I, relies on results about when solutions to parameterized optimization problems can be chosen in a measurable way.

[31] See Footnote 42.

7 Existence - An Illustrative Example

In this section, we provide an example illustrating in a simple context many of the ideas behind the proof of Theorem 1.

Example 16 *Types, actions, and signals are as in Example 13.*[32] *Rewards are restricted to* $[0,2]$. *Type* t_1 *has utility* r *if he chooses* $a \in [0,1]$, *and utility* $1+r$ *from out. Type* t_2 *has utility* $\sqrt{r}$ *from any action. The principal loses* $a+r$ *when* t_1 *chooses* $a \in [0,1]$, *loses* r *when* t_2 *chooses out, and loses* $4+r$ *when either* t_1 *chooses out or when* t_2 *chooses* $a \in [0,1]$.

Suppose the principal wishes to implement the action $a = 1/n$. Let $\rho_n(\cdot)$ be the function from $S \times A \times T$ to R given by

$$\rho_n(s,a,t_1) = \begin{cases} 2 & \text{if } a = \frac{1}{n} \text{ and } s \in \left[\frac{1}{2n}, \frac{1}{n}\right] \\ 0 & \text{otherwise} \end{cases}$$

$$\rho_n(s,a,t_2) = \begin{cases} (1-\frac{1}{n})^2/2 & \text{if } a = out, \text{ and } s = out_2 \\ 0 & \text{otherwise} \end{cases}.$$

Consider the mechanism (α_n, κ_n) given by $\alpha_n(\cdot|t_1) = \delta_{1/n}$, $\alpha_n(\cdot|t_2) = \delta_{out}$, and $\kappa_n(\cdot|s,a,t) = \delta_{\rho_n(s,a,t)}$. That is, if the agent reports t_1 she is asked to take action $1/n$, if she reports t_2 she is asked to take action *out*, and when s is observed, the agent is paid according to the function $\rho_n(\cdot)$. It is easy to verify that (α_n, κ_n) is incentive compatible, and that indeed κ_n minimizes the principal's losses over incentive compatible mechanisms implementing α_n.[33] The principal's losses are $(3+1/n^2)/4$ and so strictly decrease with n.

Consider the sequence of mechanisms (α_n, κ_n) as $n \to \infty$. A very useful step toward existence would be to ensure that along some subsequence there is a limit mechanism that is incentive compatible and incentivizes t_1 to choose $a = 0$ and t_2 to choose *out*, and in which the principal's losses are no greater than the limit of his losses, namely $\lim_n (3+1/n^2)/4 = 3/4$.[34]

The Topology of Pointwise Convergence

In this example, each (α_n, κ_n) is tied down by the function ρ_n. So it is natural, especially given the previous literature, to begin with the topology of pointwise convergence applied to $\{\rho_n\}$. Now, $\rho_n(out_2, out, t_2) \to 1/2$ while $\rho_n(s,a,t_i) \to 0$ for every other (s,a,t_i). Hence, ρ_n converges pointwise to the limit function ρ^* in which the agent is paid $1/2$ when t_2 is reported, *out* is recommended, and out_2 is observed, and is otherwise paid 0. But given this reward function, t_1 strictly prefers

[32] Hence, this example violates the assumptions on information in Page (1987,1991) due to the manner in which a mass point at $s = 0$ arises in the signal distribution when the action tends to zero.

[33] To see that losses are minimized, observe that the proposed mechanism concentrates payments to t_1 on those signals that are less likely when t_2 deviates by announcing t_1 and choosing $a = 1/n$.

[34] It is an easy exercise to argue that the principal cannot do better than this.

out to $a = 0$. Thus, even though ρ_n has a perfectly sensible pointwise limit, the limit fails to be incentive compatible and (once the agent adjusts his behavior) can lead to a discontinuous upward jump in losses for the principal.[35] Thus, the topology of pointwise convergence is unsuitable.

Both because of examples like this, and because randomization can be a valuable tool for the principal, we conclude that pointwise convergence of functions is not the economically relevant notion of convergence for the problem at hand. We do not, in particular, see the economic motivation for the restriction used in previous papers that contracts, viewed as functions from signals to payments, come from a set that is compact in a topology under which the principal's and the agent's payoffs are continuous. To proceed, we must find another notion of convergence that, (i) yields a limit contract (at least along a subsequence), (ii) ensures that the limit contract implements the limit of the action profiles, and (iii) is flexible enough to deal with randomization in both rewards and actions.

The Weak* Topology and Randomization

For each t_i, consider the probability measure $\nu_{i,n}$ over reward/signal pairs (r, s) induced by ρ_n when t_i is honest and follows the recommended action.[36]. For example, $\nu_{1,n}$ is obtained by noting that, when the agent's type is t_1 and he chooses action $1/n$, there is a 1/2 chance that $r = 2$ and s is uniformly drawn from $[1/2n, 1/n]$, and a 1/2 chance that $r = 0$ and s is uniformly drawn from $[0, 1/2n)$. Consequently, $\nu_{1,n}$ converges (in the weak* topology) to the probability measure ν_1^* that assigns probability 1/2 to the reward/signal pair $(r, s) = (2, 0)$ and probability 1/2 to $(r, s) = (0, 0)$ – note that $s = 0$ with probability one. Similarly, $\nu_{2,n}$, the probability measure over reward/signal pairs that is induced by ρ_n when t_2 chooses *out*, places probability 1 on $(r, s) = \left((1 - 1/n)^2/2, out_2\right)$, and so converges to the probability measure ν_2^* that assigns probability 1 to $(r, s) = (1/2, out_2)$.

Thus, as $n \to \infty$, the given sequence of mechanisms yields well-defined limits – in the weak* topology – of the induced equilibrium distributions over rewards and signals. Of course, ν_1^* and ν_2^* do not themselves constitute a mechanism. But, they do generate a natural guess for one. In particular, ν_1^* suggests that, when the agent reports t_1 and is asked to take action $a = 0$, if the signal $s = 0$ is observed – the only signal given positive probability by v_1^* – the agent should be paid randomly, receiving $r = 0$ or $r = 2$ each with probability 1/2. That is, we are pointed in the direction of $\kappa^* (\cdot|0, 0, t_1) = \frac{1}{2}\delta_0 + \frac{1}{2}\delta_2$. Since ν_1^* places no weight on any (r, s) with $s \neq 0$, it does not pin down the conditional distribution of rewards when $s \neq 0$. However, since these signals don't occur when t_1 chooses $a = 0$ (i.e., they occur only if either the agent lied when reporting type t_1 or did not take the recommended action $a = 0$), setting rewards to their worst possible

[35] It is easy to construct examples in which a sequence of incentive compatible mechanisms has no pointwise convergent subsequence, which is another essential flaw with the topology of pointwise convergence.

[36] In the spirit of Migrom and Weber (1985), one might think of each $\nu_{i,n}$ as the *distributional contract* induced by the contract $\rho_{i,n}$.

level (in this case $r = 0$) for any such s is a natural choice, since the only role of rewards at such "out-of-equilibrium" signals is in discouraging non-compliant reports and actions.

Similarly, ν_2^* suggests that, when the agent reports t_2 and is asked to take action *out*, the agent's reward should be $r = 1/2$ if the signal is *out* and should be $r = 0$ if any other signal is observed.

But, in fact, the suggested mechanism – i.e., that which, when t_1 is announced and $s = 0$ is observed pays $r = 0$ and $r = 2$ equiprobably, and pays $r = 0$ if any other signal is observed, and when t_2 is announced and out_2 is observed pays $1/2$ and pays $r = 0$ if any other signal is observed – *is* the unique least cost way to implement the limit action profile $(0, out)$. Consequently, the weak* topology yields a well-defined limit – constructed by considering the weak* limits, type by type, of distributions on the equilibrium path together with worst possible punishments for out-of-equilibrium signals – and also yields lower semicontinuity of the principal's losses.

We take several lessons from this example. First, sequences of deterministic mechanisms, even those with "natural" deterministic limits, may have no economically relevant limit except in mixed contracts. Second, we have yet another example in which a randomized mechanism is necessary to achieve optimality. Third, when contracts are interpreted not as *functions* from signals to payments, but rather in terms of the equilibrium distributions they induce on signal/payment pairs, the mathematically natural mode of convergence is weak convergence of measures. In this example, weak convergence provides a limit distribution that generates a natural candidate for a limit contract, one that indeed implements the limiting action profile at the limit cost.

While simple in this setting, it is important to understand what has happened here. Starting from the economically natural way of thinking about a mechanism, we translated the problem into one of thinking about contracts as joint distributions ("distributional contracts") over rewards and signals for each type, and for the action recommended. Within the space of such distributions, we found that weak convergence of measures delivers compactness and a limit. Then, we translated that limit back into a mechanism.

The heart of the construction in this paper is to generalize this idea. Because, in general, the set of types is infinite, in order to arrive at a type by type (i.e., pointwise) convergent subsequence as in the two-type example above, we will have to consider averages of mechanisms (which, because randomization is allowed, always produces a well-defined mechanism) so that we can employ a "pointwise limit of averages" theorem due to Balder (1990).

Payoffs with a Compliant Agent

The measure $\nu_{1,n}$ was defined by the distribution over rewards and signals when t_1 compliantly chose $1/n$. Given this construction, it should not be a surprise that when we consider the limit v_1^* and back out of it a mechanism that specifies $\kappa^* (\cdot|0, 0, t_1) = \frac{1}{2}\delta_0 + \frac{1}{2}\delta_2$, then t_1's utility at the limit is the limit of his utility along the sequence. Nor should it be a surprise that the losses to the principal converge.

This argument becomes a touch more subtle in our general framework, since we allow for payoffs to be unbounded above (but not below). Using what is essentially the Portmanteau theorem, it is easy to show that the principal's losses are either continuous at the limit, or fall. For incentive compatibility to survive the limit, we need to rule out similar drops in the agent's utility at the limit. It is here that Assumptions 6 and 7 play their role. If one is dealing with a sequence of mechanisms whose cost is not exploding, then events that cost the principal more than any given c must have likelihood of order no greater than $1/c$. But then, because the ratio of the principal's losses to the agent's utility explodes, these events have vanishing effect on the utility calculations of the agent, and continuity at the limit is restored.

Continuity of Information

As we passed to the limit in our example both the principal and the compliant agent had payoff that converged appropriately. It does not obviously follow from this that t_2, who was willing not to imitate t_1 along the sequence, is also willing not to imitate t_1 at the limit (nor, in more elaborate examples, that t_1 does not want to report his true type and then choose a different action).

It is here that Assumptions 8-10 play their role. In particular, for any action a, consider the deviation for t_2 when facing (α_n, κ_n) of reporting t_1 and then choosing a. The key is to recognize that this deviation will generate a distribution over signals and types that can be obtained by updating the measure $\nu_{1,n}$ using the Radon-Nikodym derivative $\xi_{t_2,a/t_1,a'_n}(\cdot)$, where for ease of comparison with Assumption 10, we let $a'_n = 1/n$ denote t_1's recommended action. If $\xi_{t_2,a/t_1,a'_n}(\cdot)$ does not jump up at the limit, then it is plausible that t_2, who by construction did not want to announce t_1 and choose a when facing (α_n, κ_n), would also not want to do so at the limit (α^*, κ^*).

In fact, Assumption 10 is more permissive than this, requiring only that, for some sequence $a_n \to a$, $\xi_{t_2,a_n/t_1,a'_n}(\cdot)$ does not jump up at the limit, a condition that we have already verified for this example. Given such an $a_n \to a$, this condition ensures that t_2, facing (α^*, κ^*), does not want to lie by reporting t_1 and then take action a, because for each n, when facing (α_n, κ_n), t_2 could have claimed to be t_1 and could have taken action a_n. But by the incentive compatibility of (α_n, κ_n), t_2 could not have gained from doing so. And the fact that $\xi_{t_2,a_n/t_1,a'_n}(\cdot)$ can only jump down at the limit implies that, if there is any discontinuity at all, it becomes easier for the principal to detect the proposed deviation at the limit than along the sequence. This is enough to imply that the deviation is unattractive at the limit.

8 Proof of Theorem 1

Let $L(\alpha, \kappa) = \int_{R\times S\times A\times T} l\,(r, s, a, t)\, d\kappa\,(r|s, a, t)\, dP(s|a, t) d\alpha\,(a|t)\, dH\,(t)$, and let

$$\lambda^* = \inf_{(\alpha,\kappa)\in M^*} L(\alpha, \kappa),$$

which is well-defined since M^* is nonempty. If $\lambda^* = +\infty$, then any mechanism in M^* is optimal and we are done. If not, then since L is bounded below by zero, we have that $\lambda^* \in [0, \infty)$. It suffices to find $(\alpha^*, \kappa^*) \in M^*$ such that $L(\alpha^*, \kappa^*) = \lambda^*$.

Let (α_n, κ_n) be a sequence of mechanisms in M^* such that $\lambda^* = \lim_n L\,(\alpha_n, \kappa_n)$, and such that $L\,(\alpha_n, \kappa_n) < \infty$ for each n. For every n and every (a, t), let $\nu_n(\cdot|a, t) = P(\cdot|a, t) \otimes \kappa_n(\cdot|\cdot, a, t)$.[37,38] Then $\nu_n(\cdot|a, t)$ is an r.c.p. from $A\times T$ to $R\times S$ by Proposition 7.29 in Bertsekas and Shreve (1978), henceforth BS. Consequently, $\delta_{\nu_n(\cdot|a,t)}$ is measurable as a function from $A \times T$ into $\Delta(\Delta(R \times S))$, being the composition of the measurable (BS, Prop. 7.26) function $\nu_n(\cdot|\cdot) : A\times T \to \Delta(R\times S)$ and the continuous (Dirac) function, $\delta_v : \Delta(R \times S) \to \Delta(\Delta(R \times S))$. Hence, for each n, $\delta_{\nu_n(\cdot|a,t)}$ is an r.c.p. from $A\times T$ to $\Delta(R\times S)$ (BS, Prop. 7.26). For each $t \in T$, define $\mu_n(\cdot|t) = \alpha_n(\cdot|t)\otimes\delta_{\nu_n(\cdot|\cdot,t)}$. Then $\mu_n(\cdot|\cdot)$ is an r.c.p. from T to $\Delta(R \times S) \times A$ (BS, Prop. 7.29). So, in particular, $\mu_n(\cdot|t) \in \Delta(\Delta(R \times S) \times A)$ for every n and every $t \in T$.

For every n,

$$\begin{aligned} L\,(\alpha_n, \kappa_n) &= \int l(r, s, a, t) d\kappa_n(r|s, a, t) dP(s|a, t) d\alpha_n(a|t) dH(t) \\ &= \int_{A\times T} \int_{R\times S} l(r, s, a, t) dv_n(r, s|a, t) d\alpha_n(a|t) dH(t) \\ &= \int_T \left(\int_A \int_{\Delta(R\times S)} \left(\int_{R\times S} l(r, s, a, t) d\nu(r, s) \right) d\delta_{v_n(\cdot|a,t)}(\nu) d\alpha_n(a|t) \right) dH(t) \\ &= \int_T \left(\int_{\Delta(R\times S)\times A} \left(\int_{R\times S} l(r, s, a, t) d\nu(r, s) \right) d\mu_n(\nu, a|t) \right) dH(t). \end{aligned} \tag{3}$$

Since $L\,(\alpha_n, \kappa_n) \to_n \lambda^* < +\infty$, Lemmas 1 and 2 (see Appendix II) imply that there is an r.c.p. $\mu^*(\cdot|\cdot)$ from T to $\Delta(R\times S)\times A$ and a subsequence $\{n_j\}$ of $\{n\}$ such that the Cesaro mean[39] of each subsequence of $\{\mu_{n_j}(\cdot|t)\}$ converges, for H a.e. t, to $\mu^*(\cdot|t)$.[40] Without loss, we may assume that

[37] For $\eta \in \Delta(X)$, and γ an r.c.p. from X to Y, $\eta\otimes\gamma \in \Delta(X \times Y)$ is defined so that for each $B \in \mathcal{B}(X \times Y)$,

$$(\eta \otimes \gamma)\,(B) = \int_X \gamma(B_x|x) d\eta(x).$$

By Bertsekas and Shreve (1978), Corollary 7.26.1, $\gamma(B_x|x)$ is a measurable function of x. Being nonnegative, it is therefore integrable.

[38] As per our earlier discussion, one might call $\nu(\cdot|a, t)$ a *distributional contract* in the spirit of Milgrom and Weber (1985).

[39] The m-th Cesaro mean of a sequence $\{x_n\}$ is $\frac{1}{m}\sum_{n=1}^{m} x_n$.

[40] Lemma 1 is proved using a small variation of the proof of Theorem 2.1 of Balder (1990), which generalizes Komlos

the original sequence $\{\mu_n(\cdot|t)\}$ has this Cesaro-mean subsequence convergence property, so that

$$\bar{\mu}_m(\cdot|t) = \frac{1}{m}\sum_{n=1}^{m}\mu_n(\cdot|t) \to_m \mu^*(\cdot|t),\ H \text{ a.e. } t \in T. \tag{4}$$

Consequently,

$$\begin{aligned}
\lambda^* &= \lim_m \frac{1}{m}\sum_{n=1}^{m} L(\alpha_n,\kappa_n) \\
&= \lim_m \int_T \left(\int_{\Delta(R\times S)\times A}\left(\int_{R\times S} l(r,s,a,t)d\nu(r,s)\right)d\bar{\mu}_m(\nu,a|t)\right)dH(t) \\
&\geq \int_T \left(\underline{\lim}_m \int_{\Delta(R\times S)\times A}\left(\int_{R\times S} l(r,s,a,t)d\nu(r,s)\right)d\bar{\mu}_m(\nu,a|t)\right)dH(t) \\
&\geq \int_T \left(\int_{\Delta(R\times S)\times A}\left(\int_{R\times S} l(r,s,a,t)d\nu(r,s)\right)d\mu^*(\nu,a|t)\right)dH(t),
\end{aligned} \tag{5}$$

where the first inequality follows from Fatou's lemma and the second by BS, Proposition 7.31, since by the same proposition, $\int_{R\times S} l(r,s,a,t)d\nu(r,s)$ is a lower semicontinuous function of (ν,a).[41]

For any $\nu \in \Delta(R\times S)$, $a,a' \in A$, and $t,t' \in T$, let $U_*(\nu,a',t',a,t) = \int_S u_*(s,a,t)\,dP(s|a,t) + \int_{R\times S}[u(r,s,a,t) - u_*(s,a,t)]\xi(s,a,t,a',t')\,d\nu(r,s)$. Incentive compatibility of (α_n,κ_n) implies that for H a.e. t, and every t',

$$\begin{aligned}
&\int_{\Delta(R\times S)\times A}\left(\int_{R\times S} u(r,s,a,t)d\nu(r,s)\right)d\mu_n(\nu,a|t) \\
=& \int u(r,s,a,t)\,d\kappa_n(r|s,a,t)\,dP(s|a,t)\,d\alpha_n(a|t) \\
\geq& \int \left(\sup_{a\in A}\int u(r,s,a,t)\,d\kappa_n\left(r|s,a',t'\right)dP(s|a,t)\right)d\alpha_n(a'|t') \\
=& \int \sup_{a\in A}\left(\int_S u_*(s,a,t)\,dP(s|a,t)\right. \\
&\left.+\int_{R\times S}[u(r,s,a,t) - u_*(s,a,t)]\xi\left(s,a,t,a',t'\right)d\kappa_n\left(r|s,a',t'\right)dP\left(s|a',t'\right)\right)d\alpha_n(a'|t') \\
=& \int_{\Delta(R\times S)\times A}\sup_{a\in A}\left(\int_S u_*(s,a,t)\,dP(s|a,t)\right. \\
&\left.+\int_{R\times S}[u(r,s,a,t) - u_*(s,a,t)]\xi\left(s,a,t,a',t'\right)d\nu(r,s)\right)d\mu_n(\nu,a'|t') \\
=& \int_{\Delta(R\times S)\times A}\sup_{a\in A} U_*(\nu,a',t',a,t)d\mu_n(\nu,a'|t'),
\end{aligned} \tag{6}$$

(1967). Page (1991) was the first to apply Komlos' theorem to the principal-agent problem.

[41] BS Proposition 7.31 is equivalent to the statement that $\int f(x,y)\,dq(y)$ is l.s.c. in (x,q) whenever f is l.s.c. in (x,y).

where the first and third equalities follow by the definition of μ_n and the second equality follows because $\kappa_n(\cdot|s,a',t') = \delta_{r_*(s)}$ and $[u(r_*(s),s,a,t) - u_*(s,a,t)] = 0$ when $s \in S_{a,t} \backslash S_{a',t'}$ and because $\xi(s,a,t,a',t') = 0$ for $P(\cdot|a',t')$ a.e. $s \in S_{a',t'} \backslash S_{a,t}$.[42] Applying Lemmas 3 and 4 to the limits of the Cesaro means of the first and last terms in (6) (after taking an appropriate subsequence if necessary for Lemma 3) yields

$$\int_{\Delta(R\times S)\times A} \left(\int_{R\times S} u(r,s,a,t) d\nu(r,s) \right) d\mu^*(\nu, a|t) \geq \int_{\Delta(R\times S)\times A} \sup_{a\in A} U_*(\nu, a', t', a, t) d\mu^*(\nu, a'|t'), \tag{7}$$

for H a.e. $t, t' \in T$, and where both integrals are finite.

For each $t \in T$, define the closed set

$$W_t = \{(\nu, a) \in \Delta(R\times S) \times A : \nu(\Phi) = 1 \text{ and } \mathrm{marg}_S \nu = P_{a,t}\}.^{43}$$

By BS Corollary 7.27.1, there exists an r.c.p. α^* from T to A and an r.c.p. η^* from $A \times T$ to $\Delta(R\times S)$ such that for every $t \in T$, $\mu^*(\cdot|t) = \alpha^*(\cdot|t) \otimes \eta^*(\cdot|\cdot, t)$. Further, $\mu^*(W_t|t) = 1$ for H a.e $t \in T$, since $\bar{\mu}_m(W_t|t) = 1$ for every m. Hence, for H a.e. t, $\eta^*(\cdot|a,t)$ places probability 1 on $\{\nu \in \Delta(R\times S) : \nu(\Phi) = 1 \text{ and } \mathrm{marg}_S \nu = P_{a,t}\}$ for $\alpha^*(\cdot|t)$ a.e. $a \in A$.

For each (a,t) define $\nu^*(\cdot|a,t) \in \Delta(R\times S)$ so that for every measurable subset B of $R \times S$,

$$\nu^*(B|a,t) = \int_{\Delta(R\times S)} \nu(B) d\eta^*(\nu|a,t). \tag{8}$$

Hence, for H a.e. t, $\nu^*(\Phi|a,t) = 1$ and the marginal of $\nu^*(\cdot|a,t)$ on S is $P_{a,t}$ for $\alpha^*(\cdot|t)$ a.e. $a \in A$. Also, because for each measurable subset B of $R\times S$, $\nu(B)$ is a measurable real-valued function of ν on $\Delta(R\times S)$, BS Prop. 7.29 implies that $\nu^*(B|a,t)$ is a measurable real-valued function of (a,t) on $A \times T$. Hence, $\nu^*(\cdot|\cdot)$ is an r.c.p. from $A\times T$ to $R \times S$.

A consequence of the definition of v^* is that for every $f : R\times S\times A\times T \to [0,\infty)$ that is measurable in (r,s) for each (a,t),

$$\int_{\Delta(R\times S)\times A} \left(\int_{R\times S} f(r,s,a,t) d\nu(r,s) \right) d\eta^*(\nu|a,t) = \int_{R\times S} f(r,s,a,t) d\nu^*(r,s|a,t). \tag{9}$$

[42] The steps in (6) show that $\sup_{a\in A} \int_{S\times R} u(r,s,a,t)\, d\kappa(r|s,a',t') dP(s|a,t) = \sup_{a\in A} U_*(a,t,\nu_{a',t'},a',t')$, and the proof of Lemma 4 shows that $\sup_{a\in A} U_*(a,t,\nu,a',t')$ is lower semicontinuous in $(\nu, a') \in W_{t'}$. So, $\sup_{a\in A} U_*(a,t,\nu_{a',t'},a',t')$ is measurable in a', being the composition of measurable functions. But then also $\sup_{a\in A} \int_{S\times R} u(r,s,a,t)\, d\kappa(r|s,a',t') dP(s|a,t)$ is measurable in a', as claimed when defining incentive-compatibility in Section 5.3.

[43] W_t is the intersection of the two closed sets $\{\nu : \nu(\Phi) = 1\} \times A$ and $\{(\nu,a) : \mathrm{marg}_S \nu = P_{a,t}\}$. The first set is closed by the portmanteau theorem because Φ is closed and the second set is closed because $P_{a,t}$ is continuous in a.

In view of (9), applying $\mu^*(\cdot|t) = \alpha^*(\cdot|t) \otimes \eta^*(\cdot|\cdot,t)$ to (7) gives, for H a.e. $t, t' \in T$,

$$\begin{aligned}
& \int_A \left(\int_{R\times S} u(r,s,a,t) d\nu^*(r,s|a,t) \right) d\alpha^*(a|t) \\
\geq\ & \int_{\Delta(R\times S)\times A} \sup_{a\in A} \left(\int_S u_*\left(s,a,t\right) dP(s|a,t) \right. \\
& \left. + \int_{R\times S} [u\left(r,s,a,t\right) - u_*\left(s,a,t\right)] \xi\left(s,a,t,a',t'\right) d\nu\left(r,s\right) \right) d\eta^*(\nu|a',t') d\alpha^*(a'|t') \\
\geq\ & \int_A \sup_{a\in A} \left(\int_{\Delta(R\times S)} \left(\int_S u_*\left(s,a,t\right) dP(s|a,t) \right. \right. \\
& \left. \left. + \int_{R\times S} [u\left(r,s,a,t\right) - u_*\left(s,a,t\right)] \xi\left(s,a,t,a',t'\right) d\nu\left(r,s\right) \right) d\eta^*(\nu|a',t') \right) d\alpha^*(a'|t') \\
=\ & \int_A \sup_{a\in A} \left(\int_S u_*\left(s,a,t\right) dP(s|a,t) \right. \\
& \left. + \int_{R\times S} [u\left(r,s,a,t\right) - u_*\left(s,a,t\right)] \xi\left(s,a,t,a',t'\right) d\nu^*\left(r,s|a',t'\right) \right) d\alpha^*(a'|t'),
\end{aligned}$$

where the final equality follows from (9).

Again, by BS Corollary 7.27.1, there exists an r.c.p. Q from $A \times T$ to S and an r.c.p. κ^* from $S \times A \times T$ to R such that for all a', t', $\nu^*(\cdot|a',t') = Q(\cdot|a',t') \otimes \kappa^*(\cdot|\cdot,a',t')$. But because the marginal of $\nu^*(\cdot|a',t')$ on S is $P_{a',t'}$, we must have $Q(\cdot|a',t') = P_{a',t'}$ for every a', t', and so $\nu^*(\cdot|a',t') = P(\cdot|a',t') \otimes \kappa^*(\cdot|\cdot,a',t')$. Also, since $\nu^*(\Phi|a,t) = 1$ holds for $\alpha^*(\cdot|t)$ a.e. a and a.e t, $\kappa^*(\Phi_s|s,a,t) = 1$ for a measurable subset Z of $S \times A \times T$ such that $[H \otimes \alpha^* \otimes P](Z) = 1$.[44]

Fix some $(\alpha_0, \kappa_0) \in M$ and modify κ^* first so that $\kappa^*(\cdot|s,a,t) = \kappa_0(\cdot|s,a,t)$ if $(s,a,t) \notin Z$ and then so that $\kappa^*(\cdot|s,a',t') = \delta_{r_*(s)}$ if $s \notin S_{a',t'}$.[45] With these modifications, $(\alpha^*, \kappa^*) \in M$, and for H a.e. t and t',

$$\begin{aligned}
& \int_A \left(\int_{R\times S} u(r,s,a,t) d\kappa^*(r|s,a,t) dP(s|a,t) \right) d\alpha^*(a|t) \\
\geq\ & \int_A \sup_{a\in A} \left(\int_S u_*\left(s,a,t\right) dP(s|a,t) \right. \\
& \left. + \int_{R\times S} [u\left(r,s,a,t\right) - u_*\left(s,a,t\right)] \xi\left(s,a,t,a',t'\right) d\kappa^*\left(r|s,a',t'\right) dP(s|a',t') \right) d\alpha^*(a'|t') \\
=\ & \int_A \left(\sup_{a\in A} \int_{R\times S} u\left(r,s,a,t\right) d\kappa^*\left(r|s,a',t'\right) dP(s|a,t) \right) d\alpha^*(a'|t'),
\end{aligned}$$

and therefore (α^*, κ^*) is almost everywhere incentive compatible and so, if necessary, we modify it on a measure zero set of types so that it is incentive compatible (see Section 6.6, third paragraph). Hence, after the modification, $(\alpha^*, \kappa^*) \in M^*$.

[44] $\{(s,a,t) \in S \times A \times T : \kappa(\Phi_s|s,a,t) = 1\}$ is measurable. See footnote 37.

[45] The modified κ is still an r.c.p. owing to the measurability conditions in Assumptions (9) and (11).

Applying (9) to (5), we have

$$\begin{aligned}
\lambda^* &\geq \int_T \left(\int_{\Delta(R\times S)\times A} \left(\int_{R\times S} l(r,s,a,t) d\nu(r,s) \right) d\mu^*(\nu, a|t) \right) dH(t) \\
&= \int_T \int_{R\times S} l(r,s,a,t) d\nu^*(r,s|a,t) d\alpha^*(a|t) dH(t) \\
&= \int_T \int_{R\times S} l(r,s,a,t) d\kappa^*(r|s,a,t) dP(s|a,t) d\alpha^*(a|t) dH(t),
\end{aligned}$$

from which we conclude by the definition of λ^* that $(\alpha^*, \kappa^*) \in M^*$ achieves losses λ^* for the principal. ∎

9 A Metric on the Space of Mechanisms

The proof of Theorem 1 goes most of the way toward showing that, under our assumptions, there is a metric for the space of mechanisms under which payoffs are appropriately continuous and the search for an optimum can be restricted to a compact set. Indeed, let d_Δ be any metric for the weak* topology on $\Delta(R \times S \times A \times T)$. Define a metric, d_M, on the space of mechanisms, M, as follows.[46]

$$d_M((\alpha, \kappa), (\alpha', \kappa')) = d_\Delta(H \otimes \alpha \otimes P \otimes \kappa, H \otimes \alpha' \otimes P \otimes \kappa').$$

Under d_M, two mechanisms (α, κ), $(\alpha', \kappa') \in M$ are considered equivalent if for H a.e. t, (a) $\alpha(\cdot|t) = \alpha'(\cdot|t)$ and (b) for $\alpha(\cdot|t)$ a.e. a, $\kappa(\cdot|s,a,t) = \kappa'(\cdot|s,a,t)$ for $P(\cdot|a,t)$ a.e. $s \in S_{a,t}$. Recall that, by the definition of M, $\kappa(\cdot|s,a,t) = \kappa'(\cdot|s,a,t) = \delta_{r_*(s)}$ for all (s,a,t) such that $s \notin S_{a,t}$.

The metric d_M is similar to that which is induced on behavioral strategies by Milgrom and Weber's (1985) use of distributional strategies in Bayesian games. Thus, one might call any distribution in $\Delta(R \times S \times A \times T)$ of the form $H \otimes \alpha \otimes P \otimes \kappa$, a *distributional mechanism.*

For any $c \in \mathbb{R}$, let $M_c = \{(\alpha, \kappa) \in M : \int l(r,s,a,t) d\kappa(r|s,a,t) dP(s|a,t) d\alpha(a|t) dH(t) \leq c\}$. We wish to argue that under the metric d_M, both M_c and $M^* \cap M_c$ are compact for all c. If so, the principal's objective function is lower semicontinuous (since M_c, being compact, is closed for each c) and, when M^* is nonempty, either every incentive compatible mechanism yields infinite (and optimal) losses or the search for an optimum can be restricted to a nonempty compact set of the form $M^* \cap M_c$ for some c. Either way, the existence of an optimum is assured.

The proof that M_c is d_M-compact is similar to the proof of Lemma 2 (but using also that the singleton set $\{H\}$ is tight), and so we will not spell out the details.

The d_M-compactness of $M^* \cap M_c$ is shown as follows. Consider any sequence $(\alpha_n, \kappa_n) \in M^* \cap M_c$. Since we have just argued that M_c is d_M-compact we may assume in particular that $H \otimes \alpha_n \otimes P \otimes \kappa_n$

[46] $H \otimes \alpha \otimes P \otimes \kappa \in \Delta(R \times S \times A \times T)$ assigns, to any measurable subset E of $R \times S \times A \times T$, probability $\int \mathbf{1}_E(r,s,a,t) d\kappa(r|s,a,t) dP(s|a,t) d\alpha(a|t) dH(t)$.

converges. Consulting the proof of Theorem 1, and recalling the construction there of μ_n from (α_n, κ_n), we have that the Cesaro mean of a subsequence of $\mu_n(\cdot|t)$ converges to some $\mu^*(\cdot|t)$, H a.e. t. Without loss of generality, let us then suppose that the Cesaro mean of $\mu_n(\cdot|t)$ converges to $\mu^*(\cdot|t)$, H a.e. t.[47] Then, by the dominated convergence theorem, the Cesaro mean of $H \otimes \mu_n$ converges to $H \otimes \mu^*$.[48] Define (α^*, κ^*) from μ^* as in the proof Theorem 1. An immediate consequence (because integrals with respect to $H \otimes \alpha_n \otimes P \otimes \kappa_n$ can be equivalently written with respect to $H \otimes \mu_n$ — see, e.g., (3) — and integrals with respect to $H \otimes \alpha^* \otimes P \otimes \kappa^*$ can similarly be equivalently written with respect to $H \otimes \mu^*$) is that the Cesaro mean of $H \otimes \alpha_n \otimes P \otimes \kappa_n$ converges to $H \otimes \alpha^* \otimes P \otimes \kappa^*$, where, as the proof shows, $(\alpha^*, \kappa^*) \in M^* \cap M_c$. But then $H \otimes \alpha_n \otimes P \otimes \kappa_n$, which we know converges, must also converge to $H \otimes \alpha^* \otimes P \otimes \kappa^*$ (using the definition of weak convergence and that if a sequence of real numbers converges then its average converges to the same limit). Hence, (α_n, κ_n) d_M-converges to (α^*, κ^*).[49]

10 Additional Restrictions on the Mechanism

The function Φ allows considerable flexibility in ruling out certain rewards as a function of the signal, including various ex-post constraints that a court or law might place on what can be enforced within a contract. But the legal system might equally well look at a mechanism in its totality and rule that it is invalid. An example might be a law that prevents insurance policies from having more than some percentage gap between premiums and expected payouts. Our machinery is general enough that it can accommodate situations in which the underlying economics suggest such restrictions on the set of mechanisms. As we will show, such restrictions are flexible enough that they offer a convenient way to model situations in which the principal has decisions available at the interim stage.

Let $M' \subseteq M$ be any subset of the set of mechanisms. The principal's M'-restricted problem is to minimize $L(\alpha, \kappa)$ subject to $(\alpha, \kappa) \in M^* \cap M'$. Recall from Section 9 the metric d_M for M and that for any $c \in R$ the subset M_c of mechanisms that yield the principal expected losses no greater than c is d_M-compact.

Corollary 1 *Suppose that Assumptions 1-11 hold and that M' is a subset of M such that $M^* \cap M'$ is nonempty. If, under the metric d_M, either M' is closed or $M' \cap M_c$ is closed for every $c \in \mathbb{R}$, then the principal's M'-restricted problem possesses a solution.*

The proof follows immediately from the discussion in Section 9 using the facts that the principal's objective function is lower semicontinuous and that a closed subset of a compact set is compact.

[47] These "without loss of generality" convergence assumptions have formally restricted us to a subsequence of (α_n, κ_n).

[48] Since the Cesaro mean of $\int_T \left(\int f(\nu, a, t) d\mu_n(\nu, a|t) \right) dH(t)$ converges to $\int_T \left(\int f(\nu, a, t) d\mu^*(\nu, a|t) \right) dH(t)$ for every bounded continuous f.

[49] Formally, the subsequence of (α_n, κ_n) that we have retsricted attention to d_M-converges to (α^*, κ^*).

To illustrate how Corollary 1 applies, consider first a situation where a regulator puts an upper bound on expected profits conditional on a subset of types.

Example 17 *Suppose that $l(\cdot)$ is continuous and bounded and that*

$$M' = \{(\alpha, \kappa) \in M : \int_{R \times S \times A \times \hat{T}} l(r, s, a, t)\, d\kappa(r|s, a, t)\, dP(s|a, t)\, d\alpha(a|t)\, dH(t) \geq c\},$$

where $\hat{T}$ is some measurable subset of T. Then M' is d_M-closed by Lemma 3 of Billingsley (1968, p.39) and BS, Proposition 7.31.

Our next example includes situations where a regulator insists that the principal sign up a certain fraction of types, or induces a certain fraction of types to take a specific action.

Example 18 *Let $g : A \times T \to \mathbb{R}$ be upper semicontinuous and bounded above, and let $M_F = \{(\alpha, \kappa) \in M : \int g(a, t)\, d\alpha(a|t)\, dH(t) \geq c\}$. Then M_F is d_M-closed by BS, Proposition 7.31.*

For example, the principal could be an automobile company, a could be a vehicle type, and $g(a, t)$ could be the gas mileage of vehicle type a. The regulatory constraint could concern the average gas mileage of vehicles sold, where $\alpha(a|t)$ reflects the type of car that the automobile company sells to a customer of type t. Or, the principal could be a bank, and the constraint could be that a certain fraction of loans are made to a certain class of borrowers.

The next example features an ex-ante individual rationality constraint for the agent. If the mechanism does not provide the agent with sufficient ex-ante utility, the agent will choose not to participate.

Example 19 *Prior to learning his type, the agent can take an outside option and receive utility u_0 or he can choose to participate in the mechanism. Thus, the principal must choose a mechanism from $M' = \{(\alpha, \kappa) \in M : \int u(r, s, a, t)\, d\kappa(r|s, a, t)\, dP(s|a, t)\, d\alpha(a|t)\, dH(t) \geq u_0\}$ in order to get the agent to participate. M' is d_M-closed by BS, Proposition 7.31.*

Next, consider a regulation that certain outcomes be rare, as might occur in financial markets.

Example 20 *Let S^o be an open subset of S. Let $M' = \{(\alpha, \kappa) \in M : \int_{A \times T} P(S^o|a, t)\, d\alpha(a|t)\, dH(t) \leq c\}$. Then, because $P(S^o|a, t)$ is l.s.c. in a for each t, M' is d_M-closed by BS, Proposition 7.31.*

The next two examples illustrate how, with the aid of restrictions on the mechanism, our model can capture situations in which the principal and the agent have decisions that must be made at the interim stage, i.e., after the agent learns his type, but before the agent takes an action.

In the next example, the agent's reward, $r = (r_1, r_2)$, has two coordinates.[50] It is assumed that the first coordinate, r_1, can as usual be chosen by the principal after the signal is observed, but that

[50] Each coordinate can be multidimensional.

the second coordinate, r_2, must be chosen by the principal at the interim stage. So, r_1 should be interpreted as the reward specified by the contract, while r_2 should be interpreted as an additional decision(s) that the principal controls.

Example 21 *Suppose that $R = R_1 \times R_2$ and that, for every s, $\Phi_s = \Phi_{1s} \times R_2$. If the principal can choose r_1 after observing the signal s, but must choose r_2 before observing the signal s, then the mechanism, (α, κ), must be restricted to the set $M' = \{(\alpha, \kappa) \in M : H \otimes \alpha \otimes P \otimes \kappa$ can be written as $H \otimes \alpha_1 \otimes P \otimes \kappa_1$, where α_1 is an r.c.p. from T to $R_2 \times A$ and κ_1 is an r.c.p from $R_2 \times S \times A \times T$ to $R_1\}$. Several examples follow. In each example, the principal's choice of r_2 must occur at the interim stage and in examples (a) and (b) r_2 is unobservable to the agent.*

(a) r_2 is the effort exerted by the principal toward a joint project with the agent.

(b) r_2 is the principal's choice of how intensively to monitor the agent. In particular, suppose that the functions u and l can be written as $u(r_1, r_2, s, a, t) = u_1(r_1, d, s, a, t) f(s|r_2, a, t)$ and $l(r_1, r_2, s, a, t) = l_1(r_1, d, s, a, t) f(s|r_2, a, t)$, where $f \geq 0$ and $\int_S f(s|r_2, a, t)\, dP(s|a, t) = 1$ for every r_2, a, t. Then, we may interpret $dP(s|r_2, a, t) = f(s|r_2, a, t)\, dP(s|a, t)$ as the signaling technology that is determined in part by the principal's choice of r_2, and we may interpret u_1 and l_1 as the agent's and the principal's payoff functions.

(c) r_2 is the principal's decision regarding an interim outside option. In particular, suppose that $R_2 = \{r_O, r_A\}$ and $u(r_1, r_O, s, a, t) = u_O(t)$ and $l(r_1, r_O, s, a, t) = l_O(t)$. Then, r_O is interpreted as the principal's decision to take the outside option (effectively excluding the agent), and r_A is interpreted as the principal's decision to allow the agent to choose an action.

To see that M' is d_M-closed, suppose that the sequence $(\alpha_n, \kappa_n) \in M'$ d_M-converges to (α, κ). Then, for each (α_n, κ_n) there is a requisite $(\alpha_{1n}, \kappa_{1n})$, such that $H \otimes \alpha_n \otimes P \otimes \kappa_n = H \otimes \alpha_{1n} \otimes P \otimes \kappa_{1n}$. By Corollary 7.27.2 of Bertsekas and Shreve we may write $H \otimes \alpha \otimes P \otimes \kappa = H \otimes \alpha_1 \otimes \gamma \otimes \kappa_1$, where α_1 is an r.c.p. from T to $R_2 \times A$, γ is an r.c.p. from $R_2 \times A \times T$ to S, and κ_1 is an r.c.p. from $R_2 \times S \times A \times T$ to R_1. Since $(\alpha_n, \kappa_n) \in M'$ d_M-converges to (α, κ), we have $H \otimes \alpha_{1n} \otimes P \otimes \kappa_{1n} \to H \otimes \alpha_1 \otimes \gamma \otimes \kappa_1$ and so all of the various marginals also converge. In particular, $H \otimes \alpha_{1n} \otimes P \to H \otimes \alpha_1 \otimes \gamma$ and $H \otimes \alpha_{1n} \to H \otimes \alpha_1$. Consequently, for any continuous $f : R_2 \times S \times A \times T \to [0, 1]$, $\int f d(H \otimes \alpha_1 \otimes \gamma) = \lim_n \int f d(H \otimes \alpha_{1n} \otimes P) = \int f d(H \otimes \alpha_1 \otimes P)$, where the last equality follows from the continuity of $P_{a,t}$ in (a, t). Hence, since f is arbitrary, $H \otimes \alpha_1 \otimes \gamma = H \otimes \alpha_1 \otimes P$ and so $H \otimes \alpha \otimes P \otimes \kappa = H \otimes \alpha_1 \otimes P \otimes \kappa_1$, as desired.

The next example illustrates how our model can capture a situation in which both the principal and the agent have outside options that are available only at the interim stage.[51]

[51] We have already discussed in Section 5.3 how our model can capture outside options for the agent that are available at the time he chooses his action. Such outside options might be available in addition to the agent's interim outside option.

Example 22 *Model this as in part (c) of the previous example, where $l_O(t)$ and $u_O(t)$ are the payoffs to the principal and agent when either one of them takes their interim outside option and the agent's type is t. For the same reason as there, we must restrict the principal to IC mechanisms in the set M' defined there. But now, in addition, we must constrain the mechanism (α, κ) so that*

$$\int_{R\times S\times A} u\left(r,s,a,t\right) d\kappa\left(r|s,a,t\right) dP(s|a,t) d\alpha\left(a|t\right) \geq u_O\left(t\right), \ H \text{ a.e. } t\in T, \tag{10}$$

because the agent can opt out after learning his type. Note that the principal can satisfy this constraint by opting out at the interim stage. Call this additional constraint set M''. Thus, we are interested in the principal's $M'\cap M''$-restricted problem.

Let us now show that $M'\cap M''\cap M_c$ is d_M-closed for every $c\in\mathbb{R}$. Since M' is d_M-closed, it suffices to show that $M''\cap M_c$ is d_M-closed. But in fact, $M''\cap M_c$ is d_M-compact. Indeed, the proof of Theorem 1 establishes that if (α_n,κ_n) is any sequence in M_c, some subsequence converges to $(\alpha^*,\kappa^*)\in M_c$ such that, along the subsequence, the Cesaro mean of $\int u\left(r,s,a,t\right) d\kappa_n\left(r|s,a,t\right) dP(s|a,t) d\alpha_n\left(a|t\right)$ converges to $\int u\left(r,s,a,t\right) d\kappa^*\left(r|s,a,t\right) dP(s|a,t) d\alpha^*\left(a|t\right)$ for H a.e. t (see the left-hand sides of (6) and (7)). It follows that $M''\cap M_c$ is d_M-compact.

11 Randomization and Optimal Contracts

In this section, we consider when optimality can be achieved without randomization. Our result is not a very surprising one – a combination of risk aversion and separability makes pure contracts optimal, because whenever the principal is randomizing over rewards, he is better off replacing the lottery with its certainty equivalent. What is perhaps more worthy of note is that the conditions required for this are quite strong. In contrast, we are not aware of economically interesting conditions that rule out randomization over recommended actions, especially beyond the case of pure moral hazard with action-independent risk attitudes towards rewards.

Let $D=\{\delta_r|r\in R\}\subseteq\Delta\left(R\right)$ be the set of Dirac measures on R. Say that a mechanism (α,κ) has deterministic rewards if $\kappa\left(\cdot|s,a,t\right)\in D$ for $H\otimes\alpha\otimes P$ a.e. (s,a,t) in $S\times A\times T$.[52]

Proposition 2 *Let $e:\Delta\left(R\right)\times S\to R$ be measurable. Suppose that for all (s,a,t) and all $\rho\in\Delta\left(\Phi_s\right)$,*

1. $e\left(\rho,s\right)\in\Phi_s$
2. $\int u\left(r,s,a',t'\right) d\rho = u\left(e\left(\rho,s\right),s,a',t'\right)$ *for all* $(a',t')\in A\times T$, *and*
3. $\int l\left(s,r,a,t\right) d\rho \geq l\left(s,e\left(\rho,s\right),a,t\right).$

Then, an optimal mechanism with deterministic rewards exists. If $\int l\left(s,r,a,t\right) d\rho > l\left(s,e\left(\rho,s\right),a,t\right)$ for all $\rho\notin D$, then every optimal mechanism has deterministic rewards.

[52] If (α,κ) has deterministic rewards we may define a deterministic contract $\tau:S\times A\times T\to R$ by $\kappa\left(\cdot|s,a,t\right)=\delta_{\tau(s,a,t)}$. Measurability of τ follows directly from the fact that κ is an r.c.p. from $S\times A\times T$ to R.

The proof is simply to start from an optimal mechanism $\kappa(\cdot|\cdot), \alpha(\cdot)$, and, for each (s,a,t) define $\hat{\kappa}(\cdot|s,a,t) = \delta_{e(\kappa(\cdot|s,a,t),s)}$. The certainty-equivalence function $e(\cdot)$ leaves all utility calculations for the agent (compliant or otherwise) unaffected, and weakly lowers the expected cost to the principal. It does so strictly if on a positive $H \otimes \alpha \otimes P$-measure set of (s,a,t), $\kappa(\cdot|s,a,t) \notin D$, and if $\int l(s,r,a,t)\, d\rho > l(s,e(\rho,s),a,t)$ for all $\rho \notin D$.

Holmström's (1979) sufficient statistic result implies that in a pure moral hazard problem in which both the principal and the agent are risk averse with separable utilities, and where the payment space is convex – randomization over payments is never optimal. The following example, which is a simple implication of Proposition 2, generalizes Holmström's result to a setting that allows for adverse selection.

Example 23 *For each s let Φ_s be an interval of real numbers, and let $u(r,s,a,t) = v(r,s)\tau(s,a,t) + \theta(s,a,t)$, where $v(\cdot,s)$ is weakly concave and increasing in r for each s, and $l(s,r,a,t) = \varphi(r,s)\sigma(s,a,t) + \chi(s,a,t)$ where $\varphi(\cdot,s)$ is weakly convex and increasing in r for each s. Then, for each s and for each $\rho \in \Delta(R)$, define $e(\rho,s)$ by $v^{-1}\left(\int v(r,s)\, d\rho\right)$, the certainty equivalent to ρ from the point of view of the agent. Condition 1 is satisfied since v is increasing, and hence $e(\rho,s)$ is contained in the convex hull of the support of ρ. Condition 2 is satisfied by construction. Condition 3 is satisfied by Jensen's equality, with strict inequality if ρ is non-degenerate and either v is strictly concave or φ is strictly convex.*

12 Appendix I

Generalization of Example 4 Let us generalize Example 4 to a setting with convex type and action spaces. Call the mechanism derived in the main text $\mathbf{m}^*$. To see that the need for randomization here does not depend on the discrete type and action spaces, let us first extend the action space to the line segment joining a_1 and a_2, and extend payoffs by linear interpolation. While $\mathbf{m}^*$ remains optimal, there is also now precisely one deterministic IC mechanism that is optimal, namely that which always recommends the action midway between a_1 and a_2 and always pays $r = 1/2$ for compliance and pays $r = 0$ otherwise. We next adjust the example to render this and all other deterministic mechanisms suboptimal.

Add a third action, a_3, and a third type, t_3. All three types are equally likely to occur and payoffs (except two for agent t_3) are defined as follows.

	t_1	t_2	t_3
a_1	r r	$1+r$ $2+r$	$*$ 1
a_2	$1+r$ $2+r$	r r	$*$ 1
a_3	-1 0	-1 0	0 0

Note that the utility to t_1 and t_2 of a_3 is negative, so that t_1 and t_2 can never be asked to take action a_3. Also, the principal strictly prefers t_3 to take action a_3 since this leads to a loss of 0 while actions a_1 and a_2 both yield a loss of 1.[53]

Let $\mathbf{m}^{**}$ be the mechanism that recommends a_3 after a report of t_3 and that employs $\mathbf{m}^*$ when the report is t_1 or t_2. If $\mathbf{m}^{**}$ is IC, then it is clearly optimal since the principal can do no better than to get t_3 to choose a_3 and to get t_1 and t_2 to behave as before. We next define t_3's payoffs from actions a_1 and a_2 so that $\mathbf{m}^{**}$ is IC (and therefore optimal) and also so that no deterministic mechanism performs as well.

Let us first convexify the now larger discrete action space $\{a_1, a_2, a_3\}$ by extending it to the action-simplex with vertices a_1, a_2, and a_3. For each type t_1 and t_2, extend the payoffs (the first two columns of the 3x3 table) of the principal and the agent to the simplex by linear interpolation; for type t_3, do this for the principal only. When the agent's type is t_3, define his (jointly concave in action and reward) payoff function to be $u_{t_3}(\alpha_1, \alpha_2, \alpha_3, r) = -5(\alpha_1 - \frac{1}{2})^2 - 5(\alpha_2 - \frac{1}{2})^2 + \frac{5}{2}\alpha_3 + (\alpha_1 + \alpha_2)r$, where $(\alpha_1, \alpha_2, \alpha_3)$ is the point in the action-simplex placing weight α_i on vertex a_i.[54]

It is a simple matter to check that the mechanism $\mathbf{m}^{**}$ is IC and hence optimal. In particular, t_3 chooses action a_3.

Because $\mathbf{m}^{**}$ employs $\mathbf{m}^*$ after reports t_1 and t_2, our previous discussion when extending the action space to the line segment joining a_1 and a_2 implies that the only deterministic IC mechanism that can possibly be optimal is that which recommends action $(\alpha_1, \alpha_2, \alpha_3) = (1/2, 1/2, 0)$ to both t_1 and t_2 and pays $r = 1/2$ for compliance, and recommends a_3 to t_3. But then t_3 can obtain utility $1/2$ by imitating either t_1 or t_2 and so t_3 will not choose a_3 and the principal's loss will strictly exceed that from $\mathbf{m}^{**}$. Hence, no deterministic mechanism performs as well as $\mathbf{m}^{**}$ and so randomization over recommended actions is necessary for optimality here.

It is now a simple matter to extend the example also to a convex type space by placing a positive density everywhere on the type-simplex with vertices t_1, t_2 and t_3 such that almost all of

[53] That these new payoffs are independent of r is for simplicity only.

[54] Payoffs are not additively separable in actions and rewards, but this is for simplicity only. Additive separability is not the issue here.

the probability mass is equally distributed on sufficiently small neighborhoods of the three vertices. So long as the payoffs are extended to the type-simplex continuously (e.g., by linear interpolation), randomization will remain strictly optimal by a continuity argument.

Proof of Proposition 1 If (i) holds, then set $\hat{r} = r_0$ and apply Theorem 9.1 (ii) in Wagner (1977) to obtain a measurable function $\hat{a}(t)$ that, for each t, maximizes $\int_S u(\hat{r}(s), s, a, t) dP(s|a, t)$ among all $a \in A$. Then, the mechanism (κ, α) defined by $\kappa(\cdot|s, a, t) = \delta_{\hat{r}(s)}$ and $\alpha(\cdot|t) = \delta_{\hat{a}(t)}$ is incentive compatible and so M^* is nonempty.

The proof that (ii) suffices is more subtle. When (ii) holds, for every $(\nu, a, t) \in \Delta(C \times S) \times A \times T$ define

$$\mathbf{U}(\nu, a, t) = \begin{cases} \int_{C \times S} u(r, s, a, t) d\nu(r, s), & \text{if } \nu(\Phi) = 1 \text{ and } \mathrm{marg}_S \nu = P_{a,t}, \\ -1, & \text{otherwise.} \end{cases}$$

Since $u \geq 0$ and $F(t) = \{(\nu, a) \in \Delta(C \times S) \times A : \nu(\Phi) = 1 \text{ and } \mathrm{marg}_S \nu = P_{a,t}\}$ is closed for each t, we have that $\mathbf{U}(\nu, a, t)$ is upper semicontinuous in (ν, a) for each $t \in T$.

For each $t \in T$, $F(t)$ is compact[55] and nonempty,[56] and $\mathbf{U}(\nu, a, t)$ is continuous in (ν, a, t) on the graph of F, a measurable subset of $\Delta(C \times S) \times A \times T$.[57] Consequently, by Wagner (1977) Theorem 9.1 (ii), there are measurable functions $\hat{v}(t)$ and $\hat{a}(t)$ such that (a) $(\hat{v}(t), \hat{a}(t)) \in F(t)$ for every $t \in T$, and (b) $(\hat{v}(t), \hat{a}(t))$ maximizes $U(\nu, a, t)$ among all $(\nu, a) \in \Delta(C \times S) \times A$. For each t, Corollary 7.27.1 of Bertsekas and Shreve (1978) yields[58] $\hat{v}(t) = P_{\hat{a}(t),t} \otimes \hat{\mu}_t$ for some r.c.p. $\hat{\mu}_t$ from S to R. Define now a mechanism $(\hat{\alpha}, \hat{\kappa})$ by $\hat{\alpha}(\cdot|t) = \delta_{\hat{a}(t)}$ and $\hat{\kappa}(\cdot|s, a, t) = \hat{\mu}_t(\cdot|s)$,[59] for every (s, a, t). Because, for every t, $\hat{v}(t)$ and $\hat{a}(t)$ were jointly chosen from $F(t)$ to maximize agent t's expected utility, $(\hat{\alpha}, \hat{\kappa})$ is incentive compatible and so M^* is nonempty.

13 Appendix II

Recall that $W_t = \{(\nu, a) \in \Delta(R \times S) \times A : \nu(\Phi) = 1 \text{ and } \mathrm{marg}_S \nu = P_{a,t}\}$ is closed. Define $\mathcal{L}_0 : \Delta(R \times S) \times A \times T \to [0, \infty)$ by $\mathcal{L}_0(\nu, a, t) = \int_{R \times S} l(r, s, a, t) d\nu$ and define $\mathcal{L} : \Delta(\Delta(R \times S) \times A) \times T \to \mathbb{R}$

---

[55] Let (ν_n, a_n) be any sequence in $F(t)$ and suppose without loss that $a_n \to a$. Then $\nu_n(\Phi) = 1$ and (see fn. 37 for the meaning of "$\otimes$" and use Bertsekas and Shreve (1978) Corollary 7.27.1) $\nu_n = P_{a_n,t} \otimes \mu_n$ for some r.c.p. μ_n from S to C. Since $P_{a_n,t}$ converges to $P_{a,t}$, $\{P_{a_n,t}\}$ is tight. Since μ_n is concentrated on the compact set C, it follows that ν_n is tight and hence has a convergent subsequence. Hence, being closed, $F(t)$ is compact.

[56] Fix any $t \in T$. Since Φ is closed, the correspondence that for each s takes the value $\Phi_s \cap C$ is measurable. Since S is Polish, this correspondence is also weakly measurable (Wagner (1977), p. 863). Hence, Wagner (1977) Theorem 4.1 implies that there is a measurable $r_1 : S \to \bar{R}$ such that $r_1(s) \in \Phi_s \cap C$ for every $s \in S$. Then μ defined by $\mu(\cdot|s) = \delta_{r_1(s)}$ for each s, is an r.c.p. from S to R. For any $a \in A$, letting $\nu = P_{a,t} \otimes \mu$ (see fn. 37) gives $(\nu, a) \in F(t)$.

[57] To see measurability, note first that for any continuous function $f : S \to [0, 1]$, $\{(\nu, a, t) \in \Delta(C \times S) \times A \times T : \int_{C \times S} f(s) d\nu(r, s) = \int_S f(s) dP(s|a, t)\}$ is measurable. Hence, defining $f_{i,k}(s) = \frac{1}{1 + k\, dist(s, U_i)}$, where $\{U_i\}$ is a countable collection of open subsets of S generating the topology of S, the set $\cap_{i,k}\{(\nu, a, t) : \int_{C \times S} f_{i,k}(s) d\nu(r, s) = \int f_{i,k}(s) dP(s|a, t)\}$ is measurable. But this latter measurable set is precisely $\{(\nu, a, t) : \mathrm{marg}_S \nu = P_{a,t}\}$ which, when intersected with the closed set $\{(\nu, a, t) : \nu(\Phi) = 1\}$ yields the graph of F.

[58] See fn. 37.

[59] Propositions 7.26 and 7.29 of Bertsekas and Shreve can be used to show that $\hat{\alpha}$ and $\hat{\kappa}$ have the necessary measurability properties.

by $\mathcal{L}(\mu,t) = \int_{\Delta(R\times S)\times A} \mathcal{L}_0(\nu,a,t)d\mu(\nu,a)$. By BS, Proposition 7.31, $\mathcal{L}_0(\nu,a,t)$ is l.s.c. in (ν,a) and so, similarly, $\mathcal{L}(\mu,t)$ is l.s.c. in μ. For the first lemma below, let μ_n be as defined in the paragraph preceding (3), i.e., prior to any reindexing.

Lemma 1 *If for every $t \in T$, $\{\mu \in \Delta(W_t) : \mathcal{L}(\mu,t) \leq c\}$ is compact for every $c \in \mathbb{R}$, then there is an r.c.p. $\mu^*(\cdot|\cdot)$ from T to $\Delta(R\times S)\times A$ and a subsequence $\{n_j\}$ of $\{n\}$ such that for any subsequence $\{n_{j_k}\}$ of $\{n_j\}$, the Cesaro mean of $\{\mu_{n_{j_k}}(\cdot|t)\}$ converges to $\mu^*(\cdot|t)$ for H a.e. t.*[60]

Proof. Let $X = \Delta(R\times S)\times A$. Each $\mu_n(\cdot|\cdot)$, being an r.c.p., is a measurable function from T to $\Delta(X)$ by BS, Proposition 7.26. Therefore, if $\mu^*(\cdot|t)$ is the pointwise H-a.e. limit of Cesaro means of $\{\mu_n(\cdot|t)\}$, $\mu^*(\cdot|\cdot)$ will be similarly measurable and hence an r.c.p. by the same proposition. Thus, it suffices to prove the Cesaro limit result.

Since X is Polish, it contains countably many open subsets $U_1, U_2, ...$ that generate its topology.[61] Let d be any metric on X and for each $i,k \in \mathbb{N}$, let $\phi_{i,k} : X \to \mathbb{R}$ be the continuous function defined by $\phi_{i,k}(x) = \frac{1}{1+kd(x,U_i)}$, and define $\alpha_{i,k}(\mu) = \int_X \phi_{i,k}(x)d\mu(x)$ for each $\mu \in \Delta(X)$. Then $\mathcal{A} = \{\alpha_{i,k}\}$ is a countable set of affine continuous functions that countably separates $\Delta(X)$ as defined in Balder (1990). Since $\int \mathcal{L}(\mu_n(\cdot|t),t)dH(t) = \lambda_n \to \lambda^* < +\infty$, the Cesaro limit result is obtained by following the remaining steps in the proof of Theorem 2.1 in Balder (1990), setting $E = \Delta(X)$, $h = \mathcal{L}$, $\phi = C = 1$, and noting that steps in which Balder uses his inf-compactness condition go through here by our compactness hypothesis.[62] ■

Lemma 2 *For every $t \in T$, $\{\mu \in \Delta(W_t) : \mathcal{L}(\mu,t) \leq c\}$ is compact for every $c \in \mathbb{R}$.*

Proof. Let us first establish,

($*$) if $f : Z \to [0,+\infty)$ is lower semicontinuous on a Polish space Z, and $\{z \in Z : f(z) \leq c\}$ is compact for every $c \in \mathbb{R}$, then $\Gamma_c = \{\gamma \in \Delta(Z) : \int f(z)d\gamma \leq c\}$ is compact for every real $c \in \mathbb{R}$.

To see ($*$), note first that Γ_c is closed since $\int f d\gamma$ is a lower semicontinuous function of γ by BS, Proposition 7.31. It thus suffices by Prohorov's theorem to show that Γ_c is tight. Choose $\varepsilon > 0$. If $\gamma \in \Gamma_c$, then because $f \geq 0$, $\gamma\{z : f(z) > c/\varepsilon\} < \varepsilon$ (Kolmogorov's inequality). Hence, each γ in Γ_c places probability at least $1-\varepsilon$ on the compact set $\{z : f(z) \leq c/\varepsilon\}$, proving ($*$).

Fix $t \in T$. If $\{(\nu,a) \in W_t : \mathcal{L}_0(\nu,a,t) \leq c\}$ is compact for every $c \in \mathbb{R}$, then ($*$) implies that $\{\mu \in \Delta(W_t) : \mathcal{L}(\mu,t) \leq c\}$ is compact for every $c \in \mathbb{R}$ and the proof would be complete. Hence,

[60] The exceptional null set of t's can depend on the subsequence $\{n_{j_k}\}$ of $\{n_j\}$.

[61] i.e., every open subset of X is the union of some collection of the U_i.

[62] The reason that we cannot directly apply Balder's Theorem 2.1 is that for each t, $\{\mu \in \Delta(X) : \mathcal{L}(\mu,t) \leq c\}$ need not be compact, as Balder's Theorem 2.1 requires. However, the fact that for each t, $\mu_n(\cdot|t) \in \Delta(W_t)$ and our hypothesis that $\{\mu \in \Delta(W_t) : \mathcal{L}(\mu,t) \leq c\}$ is compact for every $c \in \mathbb{R}$ serves the same purpose at the relevant point in Balder's proof.

given $c \in \mathbb{R}$, it suffices to show that $C = \{(\nu, a) \in W_t : \mathcal{L}_0(\nu, a, t) \leq c\}$ is compact. Note that C is closed since W_t is closed and $L_0(\cdot, t)$ is lower semicontinuous.

Fix $\varepsilon > 0$. Since $P_{a,t}$ is continuous in a on the compact set A, $\{P_{a,t}\}_{a \in A}$ is compact and hence by Prohorov's theorem tight and so there is a compact subset Y of S such that $P(Y|a, t) > 1 - \varepsilon$ for every $a \in A$. Let $D = \{(r, s, a) \in \Phi \times A : l(r, s, a, t) \leq c/\varepsilon\}$ and note that $D \cap (R \times Y \times A)$ is compact by Assumption 6.

For any $(\nu, a) \in C$ we have $\nu(\Phi) = 1$ and so, $c \geq \int_\Phi l(r, s, a, t) d\nu(r, s) = \int_{\Phi \times A} l(r, s, a', t) d\nu(r, s) d\delta_a(a')$. Hence by Kolmogorov's inequality and $l \geq 0$,

$$(\nu \times \delta_a)(D) > 1 - \varepsilon. \tag{11}$$

Also,

$$\begin{aligned} (\nu \times \delta_a)(D \cap (R \times Y^c \times A) &\leq (\nu \times \delta_a)(R \times Y^c \times A) \\ &= \nu(R \times Y^c) = P(Y^c|a, t) \leq \varepsilon, \end{aligned} \tag{12}$$

where the second equality follows because $(\nu, a) \in C \subseteq W_t$ implies that the marginal of ν on S is $P(\cdot|a, t)$.

Together, (11) and (12) imply $(\nu \times \delta_a)(D \cap (R \times Y \times A)) > 1 - 2\varepsilon$. Thus $G = \{(\nu, \delta_a) \in \Delta(R \times S) \times \Delta(A) : (\nu, a) \in C\}$ is tight and hence relatively compact by Prohorov's theorem. Since C is the image of G under the continuous map $(\nu, \delta_a) \to (\nu, a)$, C too is relatively compact. Hence, being closed, C is compact. ■

For the remainder of this appendix, it is assumed, as in the main proof following (3), that the Cesaro mean of each subsequence of $\{\mu_n(\cdot|t)\}$ converges to $\mu^*(\cdot|t)$ for H a.e. t.

Lemma 3 *Let $\mathcal{U}(\nu, a, t) = \int_{R \times S} u(r, s, a, t) d\nu(r, s)$. There is a subsequence, $\{n_j\}$ of $\{n\}$ such that for H a.e. $t \in T$,*

$$f_n(t) = \int_{\Delta(R \times S) \times A} \mathcal{U}(\nu, a, t) d\hat{\mu}_n(\nu, a|t) \to_n f(t) = \int_{\Delta(R \times S) \times A} \mathcal{U}(\nu, a, t) d\mu^*(\nu, a|t),$$

where $\hat{\mu}_n$ denotes the n^{th} Cesaro mean of $\{\mu_{n_j}\}$. Moreover, $\int_T f_n(t) dH(t) \to_n \int_T f(t) dH(t) < +\infty$.

Proof. For $m \in \mathbb{N}$, let x_m be the supremum of $u(r, s, a, t)/l(r, s, a, t)$ over all $(r, s, a, t) \in \Phi \times A \times T$ such that $u(r, s, a, t) > m$ (and where the sup of the empty set is zero). By Assumption 7, x_m is finite for all m sufficiently large. Let $u_m(r, s, a, t) = \min(u(r, s, a, t), m)$.

Recall from the main proof that $\lim_n L(\alpha_n, \kappa_n) = \lambda^* < +\infty$, and so we may choose c such that $L(\alpha_n, \kappa_n) \leq c$ holds for all n.

Define $\gamma_n = H \otimes \mu_n \in \Delta(\Delta(R \times S) \times A \times T)$. Then,[63]

$$\int u(r,s,a,t) d\nu d\gamma_n = \int u_m(r,s,a,t) d\nu d\gamma_n + \int_{u(r,s,a,t)>m} (u(r,s,a,t) - m) d\nu d\gamma_n. \tag{13}$$

Consider the last term. Then, because $\nu(\Phi) = 1$ for γ_n a.e. $(\nu, a, t) \in \Delta(R \times S) \times A \times T$,

$$\begin{aligned}
\int_{u(r,s,a,t)>m} (u(r,s,a,t) - m) d\nu d\gamma_n &\leq \int_{u(r,s,a,t)>m} u(r,s,a,t) d\nu d\gamma_n \\
&= \int_{u(r,s,a,t)>m} \frac{u(r,s,a,t)}{l(r,s,a,t)} l(r,s,a,t) d\nu d\gamma_n \\
&\leq x_m \int l(r,s,a,t) d\nu d\gamma_n \\
&= x_m L(\alpha_n, \kappa_n) \\
&\leq x_m c,
\end{aligned}$$

where the second inequality follows because l is nonnegative and the last equality follows from (3). Hence,

$$\int u_m(r,s,a,t) d\nu d\gamma_n \leq \int u(r,s,a,t) d\nu d\gamma_n \leq \int u_m(r,s,a,t) d\nu d\gamma_n + x_m c, \tag{14}$$

where the second inequality follows from (13) .

To see that $\int u(r,s,a,t) d\nu d\gamma_n$ is bounded, pick any m large enough that x_m is finite, and note that (14) implies that for all n,

$$\begin{aligned}
\int u(r,s,a,t) d\nu d\gamma_n &\leq \int u_m(r,s,a,t) d\nu d\gamma_n + c x_m \\
&\leq m + c x_m.
\end{aligned}$$

Thus we have shown that $\sup_n \int \mathcal{U}(\nu, a, t) d\mu_n(\nu, a|t) dH(t) < +\infty$. Consequently, letting $f_n(t) = \int \mathcal{U}(\nu, a, t) d\mu_n(\nu, a|t)$, by Komlos (1967), there is a subsequence $\{n_j\}$ of $\{n\}$ and a measurable $f : T \to R$ such that the Cesaro mean of $\{f_{n_j}(t)\}$ converges to $f(t)$ for H a.e. $t \in T$.

For any m, because $\int_{R \times S} u_m(r,s,a,t) d\nu(r,s)$ is bounded and continuous in (ν, a), letting $\hat{\mu}_n$ be the n^{th} Cesaro mean of $\{\mu_{n_j}\}$, the Cesaro mean convergence property of $\{\mu_n\}$ implies

$$\int_{\Delta(R \times S) \times A} \left(\int_{R \times S} u_m(r,s,a,t) d\nu \right) d\hat{\mu}_n(\nu, a, |t) \to_n \int_{\Delta(R \times S) \times A} \left(\int_{R \times S} u_m(r,s,a,t) d\nu \right) d\mu^*(\nu, a|t) \text{ for } H \text{ a.e. } t,$$

and so by the dominated convergence theorem, letting $\hat{\gamma}_n = H \otimes \hat{\mu}_n$ and $\gamma^* = H \otimes \mu^*$,

$$\int_{\Delta(R \times S) \times A \times T} \left(\int_{R \times S} u_m(r,s,a,t) d\nu \right) d\hat{\gamma}_n(\nu, a, t) \to_n \int_{\Delta(R \times S) \times A \times T} \left(\int_{R \times S} u_m(r,s,a,t) d\nu \right) d\gamma^*(\nu, a, t).$$

[63] It is understood that $\int$, without delimiters, implies that integration is with respect to all variables. Also, $\int u(r,s,a,t) d\nu d\gamma_n$, for example, is shortand for $\int_{\Delta(R \times S) \times A \times T} \left(\int_{R \times S} u(r,s,a,t) d\nu(r,s) \right) d\gamma_n(\nu, a, t)$.

Hence, from (14) we have

$$\int u_m(r,s,a,t)d\nu d\gamma^* \leq \underline{\lim}_n \int u(r,s,a,t)d\nu d\hat{\gamma}_n \leq \overline{\lim}_n \int u(r,s,a,t)d\nu d\hat{\gamma}_n \leq \int u_m(r,s,a,t)d\nu d\gamma^* + cx_m.$$

But, $\lim_{m\to\infty} x_m = 0$ by Assumption 7, and $\lim_m \int u_m(r,s,a,t)d\nu d\gamma^* = \int u(r,s,a,t)d\nu d\gamma^*$ by the monotone convergence theorem, and so we conclude that

$$\lim_n \int \mathcal{U}(\nu,a,t)d\hat{\mu}_n(\nu,a,|t)dH(t) = \int \mathcal{U}(\nu,a,t)d\mu^*(\nu,a|t)dH(t) < +\infty, \tag{15}$$

proving the second part of the Lemma.

Now,

$$\begin{aligned}
\lim_n \int_T \int_{\Delta(R\times S)\times A} \mathcal{U}(\nu,a,t)d\hat{\mu}_n(\nu,a,|t)dH(t) &\geq \int_T \underline{\lim}_n \left(\int_{\Delta(R\times S)\times A} \mathcal{U}(\nu,a,t)d\hat{\mu}_n(\nu,a,|t)\right) dH(t) \\
&= \int_T \lim_n \left(\int_{\Delta(R\times S)\times A} \mathcal{U}(\nu,a,t)d\hat{\mu}_n(\nu,a,|t)\right) dH(t) \\
&\geq \int_T \left(\int_{\Delta(R\times S)\times A} \mathcal{U}(\nu,a,t)d\nu d\mu^*(\nu,a,|t)\right) dH(t),
\end{aligned}$$

where the first inequality is by Fatou's lemma, the equality is because the Cesaro mean of $f_{n_j}(t) = \int_{\Delta(R\times S)\times A} \mathcal{U}(\nu,a,t)d\mu_{n_j}(\nu,a,|t)$ converges for H a.e.t (say for all t in the H-measure one set T^0), and the second inequality follows by Proposition 7.31 of BS applied to the integral in parentheses since u is continuous and bounded below (so $\mathcal{U}$ is lower semicontinuous and bounded below). Both inequalities (in particular the last) must be equalities by (15). Hence, the (by Proposition 7.31 of BS) nonnegative function of $t \in T^0$,

$$\lim_n \left(\int_{\Delta(R\times S)\times A} \mathcal{U}(\nu,a,t)d\hat{\mu}_n(\nu,a,|t)\right) - \left(\int_{\Delta(R\times S)\times A} \mathcal{U}(\nu,a,t)d\nu d\mu^*(\nu,a,|t)\right), \tag{16}$$

integrates to zero with respect to H and so it must be zero for H a.e. t. ∎

Recall that $U_*(\nu,a',t',a,t) = \int_S u_*(s,a,t)\, dP(s|a,t) + \int_{R\times S}[u(r,s,a,t) - u_*(s,a,t)]\xi(s,a,t,a',t')\, d\nu(r,s)$.

Lemma 4 *Let $\{\hat{\mu}_n\}$ be as in Lemma 3. For H a.e. $t, t' \in T$,*

$$\underline{\lim}_n \int_{\Delta(R\times S)\times A} \sup_{a\in A} U_*(\nu,a',t',a,t)d\hat{\mu}_n(\nu,a'|t') \geq \int_{\Delta(R\times S)\times A} \sup_{a\in A} U_*(\nu,a',t',a,t)d\mu^*(\nu,a'|t').$$

Proof. Because $\hat{\mu}_n(W_{t'}|t') = \mu^*(W_{t'}|t') = 1$, for all n, it suffices by BS, Proposition 7.31, to show that for any $t, t' \in T$, $\sup_{a\in A} U_*(\nu,a',t',a,t)$ is lower semicontinuous in (ν,a') on the closed set $W_{t'}$. Let (ν_n,a'_n) be a sequence in $W_{t'}$ converging to (ν_0,a'_0). Fix any $\varepsilon > 0$ and choose $a_0 \in A$

such that $U_*(\nu_0, a_0', t', a_0, t) + \varepsilon \geq \sup_{a \in A} U_*(\nu_0, a_0', t', a, t)$. Choose a sequence a_n in A converging to a_0 according to Assumption 10. For any $x, x' \in A \times T$, recall that $\xi(s, x, x')$ denotes the Radon-Nikodym derivative of $P_x|_{S_{x'}}$ with respect to $P_{x'}$ evaluated at $s \in S$. Let $Z = \{1, \frac{1}{2}, \frac{1}{3}, ..., 0\}$ and define $g : S \times Z \to [0, +\infty]$, by

$$g(s, z) = \begin{cases} \xi(s, a_n, t, a_n', t'), & \text{if } z = 1/n \text{ and } s \in S_{a_n', t'}, \\ +\infty, & \text{otherwise} \end{cases},$$

and let $\underline{g}$ denote the lower envelope of g. That is, for any $(s, z) \in S \times Z$,

$$\underline{g}(s, z) = \inf_{(s_n, z_n) \to (s, z)} \underline{\lim}_n g(s_n, z_n).^{64}$$

By the definition of U_*, for each n,

$$U_*(\nu_n, a_n', t', a_n, t) = \int_S u_*(s, a_n, t)\, dP(s|a_n, t) + \int_{R \times S} [u(r, s, a_n, t) - u_*(s, a_n, t)] \xi\left(s, a_n, t, a_n', t'\right) d\nu_n(r, s). \tag{17}$$

Consider the first term on the right-hand side of (17) Because $u_*(s, a, t)$ is nonnegative and continuous in (s, a) for each t, BS, Proposition 7.31, implies

$$\underline{\lim}_n \int_S u_*(s, a_n, t)\, dP(s|a_n, t) \geq \int_S u_*(s, a, t)\, dP(s|a_0, t).$$

For the second term on the right-hand side of (17), letting $\bar{u}$ denote the continuous function $u - u_*$, we have

$$\begin{aligned} \underline{\lim}_n \int_{R \times S} \bar{u}(r, s, a_n, t)\, \xi\left(s, a_n, t, a_n', t'\right) d\nu_n(r, s) &= \underline{\lim}_n \int_{R \times S} \bar{u}(r, s, a_n, t)\, g(s, 1/n) d\nu_n(r, s) \\ &= \underline{\lim}_n \int_{R \times S \times Z} \bar{u}(r, s, a_n, t)\, g(s, z) d(\nu_n(r, s) \times \delta_{\frac{1}{n}}(z)) \\ &\geq \underline{\lim}_n \int_{R \times S \times Z} \bar{u}(r, s, a_n, t)\, \underline{g}(s, z) d(\nu_n(r, s) \times \delta_{\frac{1}{n}}(z)) \\ &\geq \int_{R \times S \times Z} \bar{u}(r, s, a_0, t)\, \underline{g}(s, z) d(\nu_0(r, s) \times \delta_0(z)) \\ &= \int_{R \times S} \bar{u}(r, s, a_0, t)\, \underline{g}(s, 0) d\nu_0(r, s) \\ &\geq \int_{R \times S} \bar{u}(r, s, a_0, t)\, \xi(s, a_0, t_0, a_0', t_0') d\nu_0(r, s). \end{aligned}$$

where the first inequality follows because $\bar{u}$ is nonnegative on $\Phi \times A \times T$ by Assumption 11, $\nu_n(\Phi) = 1$ since $(\nu_n, a_n') \in W_{t'}$, and $g \geq \underline{g}$; the second inequality follows by BS, Proposition 7.31, because

[64] Measurability of g follows from the measurability of $\xi(s, a_n, t, a_n', t')$ in s for each n, while $\underline{g}$ is measurable because it is lower semicontinuous.

$\bar{u}\,(r,s,a,t)\,\underline{g}(s,z)$ is lower semicontinuous and bounded below (by 0); and the third inequality follows because $(\nu_0, a_0') \in W_{t'}$ implies that the marginal of ν_0 on S is $P_{a_0',t'}$, and Assumption 10 implies that $\underline{g}(s,0) \geq \xi(s, a_0, t, a_0', t')$ holds for $P_{a_0',t'}$ a.e. $s \in S$. Hence,

$$\begin{aligned}
\underline{\lim}_n \sup_{a\in A} U_*(\nu_n, a_n', t', a, t) &\geq \underline{\lim}_n U_*(\nu_n, a_n', t', a_n, t) \\
&\geq U_*(\nu_0, a_0', t', a_0, t) \\
&\geq \sup_{a\in A} U_*(\nu_0, a_0', t', a, t) - \varepsilon.
\end{aligned}$$

Since $\varepsilon > 0$ is arbitrary, we are done. ■

References

[1] Arnott, R., and J. E. Stiglitz (1988): "Randomization with Asymmetric Information," *The Rand Journal of Economics,* 19, 344-362.

[2] Balder, E. (1990): "New Sequential Compactness Results for Spaces of Scalarly Integrable Functions," *Journal of Math. Anal. and Appl.*, 151, 1-16.

[3] Balder, E. (1996), "On the Existence of Optimal Contract Mechanisms for Incomplete Information Principal-Agent Models," *Journal of Economic Theory*, 68, 133-148.

[4] Bertsekas, D. P., and S. E. Shreve (1978): *Stochastic Optimal Control: The Discrete Time Case*, Academic press, London. Available for download at http://dspace.mit.edu/handle/1721.1/4852.

[5] Billingsley, P., (1968): *Convergence of Probability Measures*, Wiley Series in Probability and Mathematical Statistics, New York.

[6] Carlier, G., and R. A. Dana (2005): "Existence and Monotonicity of Solutions to Moral Hazard Problems," *Journal of Mathematical Economics,* 41, 826–843.

[7] Chen Y., and Frank H. Page (2012): "Optimal Contracts with the Inclusion of Observable Actions in Moral Hazard Problems," working paper, Indiana University.

[8] Conlon, J. (2009): "Two New Conditions Supporting the First-Order Approach to Multisignal Principal-Agent Problems," *Econometrica,* 77, 249–278.

[9] Gjesdal, F. (1982): "Information and Incentives: The Agency Information Problem," *The Review of Economic Studies*, 49, 373-390.

[10] Grossman, S., and O. Hart (1983): "An Analysis of the Principal Agent Problem," *Econometrica*, 51, 7-45.

[11] Holmström, B. (1979): "Moral Hazard and Observability," *Bell Journal of Economics,* 10, 74-91.

[12] Jewitt, I. (1988): "Justifying the First Order Approach to Principal-Agent Problems," *Econometrica,* 56, 1177-1190.

[13] Jewitt I., O. Kadan, and J. M. Swinkels (2008): "Moral Hazard with Bounded Payments," *Journal of Economic Theory,* 143, 59-82.

[14] Kadan, O, and J.M. Swinkels (2013): "On the Moral Hazard Problem without the First-Order Approach". *Journal of Economic Theory*, 148, 2313-2343.

[15] Komlos, J. (1967): "A generalization of a problem of Steinhaus," *Acta Math. Acad. Sci. Hungur.,* 18, 217-229.

[16] Laffont, J. - J., and D. Martimort (2002), "The Theory of Incentives - The Principal-Agent Model," Princeton University Press.

[17] Milgrom, P. R., and R. J. Weber (1985): "Distributional Strategies for Games with Incomplete Information," *Mathematics of Operations Research,* 10, 619-632.

[18] Mirrlees, J. (1976): "The Optimal Structure of Incentives and Authority within an Organization," *Bell Journal of Economics,* 7, 105-131.

[19] Mirrlees, J. (1999): "The Theory of Moral Hazard and Unobservable Behavior: Part I," *Review of Economic Studies,* 66, 3-21.

[20] Moroni, S., and J. M. Swinkels (2014): "Existence and Non-Existence in the Moral Hazard Problem" *Journal of Economic Theory,* 150, 668-682.

[21] Myerson, R. B. (1982): "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems," *Journal of Mathematical Economics*, 10, 67-81.

[22] Page, F. H. (1987): "The Existence of Optimal Contracts in the Principal-Agent Model," *Journal of Mathematical Economics,* 16, 157-167.

[23] Page, F. H. (1991) "Optimal Contract Mechanisms for Principal-Agent Problems with Moral Hazard and Adverse Selection," *Economic Theory,* 1, 323-338.

[24] Rogerson, W. P. (1985): "The First-Order Approach to Principal-Agent Problems," *Econometrica,* 53, 1357-1367.

[25] Sinclair-Desgagne, B. (1994): "The First-Order Approach to Multi-Signal Principal-Agent Problems," *Econometrica,* 62, 459-465.

[26] Wagner, D. H. (1977): "Survey of Measurable Selection Theorems," *Siam Journal of Control and Optimization,* 15, 859-903.