Existence of Optimal Mechanisms in Principal-Agent Problems*

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

First Draft, October 2009
This Draft, September 2016

Abstract

We provide general conditions under which principal-agent problems with either one or multiple agents admit mechanisms that are optimal for the principal. Our results cover 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 allow 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 a group of agents to behave optimally from her point of view. Each agent has private information summarized by his "type" and can take an action that the principal cannot directly observe. However, the principal can observe signals whose distribution depends on the agents' types and actions, and can choose (possibly randomly) rewards for the agents. This principal-agent setting is quite general, incorporating as special cases the case of a single agent, pure moral-hazard (one possible type), pure adverse selection (one possible action) and settings with both, as for example, a health insurance provider that

*We thank the editor, Joel Sobel, and the referees, as well as 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

worries not only about what the agents may know about their health but also about any actions the agents may take that affect their health.[1]

The purpose of this 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 agents each privately report their type. Given the reported types, the mechanism privately recommends (possibly randomly) an action for each agent to take. Then, simultaneously, the agents choose their actions and the mechanism specifies the agents' contracts - mappings from signals into (possibly randomized) rewards (the specified contracts can depend on the vectors of reported types and recommended actions). Signals are then generated with distribution determined by the true types and actions of the agents. Given the signals, the agents' reward is then generated according to the contracts.

Existence of an optimal mechanism in such settings is a significant question. Principal-agent problems are central to economics.[2] Moreover, 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). But, none of this is relevant unless one knows that an optimal solution exists.[3] So it is troubling that Mirrlees (1999) provides an example of a surprisingly simple economic setting (a single agent, pure moral hazard, logarithmic utility, normally distributed signals) in which an optimal mechanism does not in fact exist.[4]

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 the monotone likelihood ratio property 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 he is paid, and the desirability of his office. The utility of the agents and the principal can depend in a general manner (with appropriate continuity) on the types, actions, signal, and reward. The principal and agents can have common or opposing interests or anything in between. We permit the utility of the agents and the disutility (henceforth loss) of the principal to

[1] Previous versions of this paper considered only the case of a single agent. We are grateful to a referee for suggesting that we think harder about the multi-agent case.

[2] 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.

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

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

be unbounded above. Signal supports can vary with the agents' actions and types and can contain atoms.

For simplicity, we set up our baseline model for the case of a single agent and establish our main existence result (Theorem 4.11) for this case. We then show how the single-agent case extends naturally to a setting with multiple agents. So, from this point forward, we will, for the most part, couch all of our discussion within the single-agent context until we reach the multi-agent model in Section 13.

The key to our existence results is to associate with each incentive compatible mechanism the joint distribution that it induces on the space of rewards, signals, actions, and types. We call such a distribution a *distributional mechanism***.** Thus, two mechanisms are different precisely when they generate different joint distributions on the items of economic interest. The advantage of casting the problem in terms of distributional mechanisms is that, first, when endowed with the weak* topology, this space is metrizable.[5] Second, and more importantly, under this metric there are weak assumptions ensuring that the principal's loss is lower semicontinuous as a function of the mechanism and that the set of incentive compatible mechanisms that bound from above the principal's loss is compact. Consequently, an optimal mechanism exists so long as the set of incentive compatible mechanisms is nonempty, a condition that is typically trivial to verify in single-agent applications, but somewhat less so in the multi-agent case. We give general sufficient conditions for the existence of an incentive compatible mechanism in both cases (Propositions 8.4 and 13.6).

Our distributional-mechanism approach is inspired by the useful role of distributional strategies in Bayesian games. Distributional strategies were introduced by Milgrom and Weber (1985), who showed that strategies in Bayesian games, namely measurable functions from types into distributions over actions, could be usefully topologized when identified with the joint distributions on actions and types that they induce when combined with the given prior. Despite this important connection, our results are not a direct translation of Milgrom and Weber's. Most critically, in our setting, the distribution of signals, and in particular its support, may depend on the true type and action of the agent. Consequently, a distributional mechanism only pins down the contract part of the mechanism "on-path" i.e., only on signal events that occur with positive probability when the agent always truthfully reports his type and always takes the recommended action. It does not pin down the contract part of the mechanism on signal events that occur with positive probability only when the agent lies about his type or chooses not to take the recommended action. On these events, which are crucial to consider for incentive compatibility, we must construct the mechanism "by

[5] Recall that, according to the weak* topology, a sequence of probability measures μ_n converges to μ if and only if for each continuous and bounded function f, the sequence of expectations $\int f d\mu_n$ converges to $\int f d\mu$.

hand," and a number of new assumptions introduced here permit us to do this. In contrast, distributional strategies in Bayesian games, under Milgrom and Weber's assumptions, completely pin down the players' strategies up to irrelevant measure zero subsets of their type spaces.

Our first main result, Theorem 4.11, establishes the existence of an optimal mechanism in the case of a single agent under a number of assumptions, some of which are only technical in nature. But 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 the loss the principal can suffer.[6] Second, we require that if the utility of the agent can be made unboundedly high, it becomes arbitrarily 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 mild. In particular (see Section 6), it is satisfied whenever the signal distribution admits a density that is sufficiently continuous, which is common in applications. In addition, this new continuity of information condition covers important settings that previous models could not, for example, any setting in which one or more dimensions of the agent's multi-dimensional action in $[0,1]^k$ is observable. Fourth, we sometimes require that, if the principal observes a signal that is inconsistent with the reported type and the recommended action, then there is a way to punish the agent that does not depend on his true type and action. These assumptions apply naturally to the multi-agent case, allowing us to establish our second main result, Theorem 13.5, which extends our single-agent existence result to a multi-agent setting.[7]

A key methodological contribution is that the existence of an optimal mechanism is established without imposing any restrictions, beyond measurability, on the set of mechanisms. This is in contrast to the most general existence results to date (e.g., Page (1987), (1991), and Balder (1991)) that require a mechanism to employ only deterministic contracts (i.e., contracts that do not randomize over the agent's reward) from some prespecified compact set, even if this is not in the principal's best interest. As a result, a mechanism that would be considered "optimal" through the lens of this literature might not actually be fully optimal. There are two reasons for this. First, there may be deterministic contracts outside the prespecified compact set that are better for the principal. Second, randomization by the principal that can relax the agent's incentive constraints is not permitted. Because we

[6] With some separability, unbounded ex-post payoffs can be accommodated by instead placing bounds on expected payoffs.

[7] As we shall see, our multi-agent result requires somewhat stronger informational assumptions than our single-agent result owing to the nature of the multi-agent incentive constraints. This provides a good reason to treat the two models separately given the importance of the single-agent model in its own right.

allow randomization and we allow any measurable contract, the solutions whose existence we establish are fully optimal.

In some settings, the additional randomization that we allow is not used. For example, we show that if both the principal and the agent are risk averse, and the payoffs to the agent are sufficiently separable in reward, action, and type, then the optimal mechanism never requires randomization over the agent's reward. But, even under these conditions, proving the existence of an optimal mechanism is simpler and can be established with more generality by allowing the possibility of randomization from the start.

While much of our analysis places no *a priori* restrictions on the mechanism, restricting attention to deterministic contracts can sometimes be quite natural since, for example, they may be simpler to enforce. Consequently, the existence results in the literature certainly remain important. But there are many settings in which randomization arises naturally and so we should want models and results that include this possibility as well. For a more complete discussion of the role of randomization, see Section 8. More generally, the economic setting of interest may call for any number of restrictions 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. In Section 11, we show how our model and results readily adapt to such restrictions, and we also show how to model settings that include ex-ante participation constraints, settings in which the principal can choose to exclude the agent, and settings in which the principal has decisions to take beyond the choice of the agent's contract. But the fundamental starting point is that in which no auxiliary restrictions are placed on the mechanism, and it is our general existence result for this setting from which all of our other results follow.

The remainder of the paper proceeds as follows. Section 2 presents three examples to motivate our approach to information continuity and the role of randomization. Section 3 introduces the single-agent model. Section 4 presents our assumptions and main existence result for the single-agent case. The assumptions are discussed and illustrated in Section 5. In Section 6 we present a simplified case which is quite common in applications and explain how our results apply. Section 7 discusses relevant literature. In Section 8 we discuss the different roles of randomization in our setting and introduce the metric space of distributional mechanisms. Section 9 provides a sketch of the proof of the main theorem. The formal proof is presented in Section 10. Section 11 shows how to adapt our model to ex-ante and interim restrictions on the set of mechanisms. Section 12 discusses the circumstances under which optimal contracts can be chosen to be deterministic without loss to the principal, and discusses the question of existence if one wishes to foreclose the principal from one or both forms of randomization even though they might be beneficial. Finally, Section 13 shows how our existence result for the single-agent problem extends to the multi-agent case.

2 Three Examples

This section presents three examples. The first two highlight the need for an appropriate notion of continuity of information, without which an optimal mechanism can fail to exist. The third illustrates that randomized contracts may be required for existence, and that, by restricting attention to an a priori "simple" set of deterministic contracts one may actually dramatically increase the complexity of the optimal mechanism.

Example 2.1 *Let the type space be a singleton. 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 is equal to* $r + 1$ *if he takes action* $a = 1$.[8] *The principal's losses 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 any 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.[10] 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 a 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 Example 2.1. But, as the next example illustrates, this is not enough.

Example 2.2 *Modify Example 2.1 only in that* S_k *is the subset of signals* s *in* $[0, 1]$ *where the k-th digit in the binary expansion of* s *is* 0*. So, for each* k*,* $S_k = \cup_{m \geq 0 \text{ even}} [\frac{m}{2^k}, \frac{m+1}{2^k}] \cap [0, 1]$

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

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

[10] Facing such a contract, the agent gets utility 3 with certainty by taking action $1/k$, and earns some lottery over a utility of 3 and 0 by any other action except $a = 1$. The action $a = 1$ earns him utility $r + 1 = 4$ half of the time and utility 1 the rest of the time, for expect utility $2.5 < 3$.

still has (Lebesgue) measure $1/2$, *but is the union of non-adjacent closed intervals of length* $1/2^k$ *starting with* $[0, 1/2^k]$. *Now, as* $a \to 0$, *the signal distribution not only converges in the weak* topology to the uniform distribution, but 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 reshuffled in an irrelevant way. A satisfactory continuity-of-information 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 (in either example) on S_k, $\frac{f(s|1)}{f\left(s|\frac{1}{k}\right)} = \frac{1}{2}$. 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. In particular, each util given to the compliant agent choosing $a = \frac{1}{k}$ adds only half a util to the utility of a deviating agent choosing $a = 1$, and so it is easy to motivate the agent to choose action $a = \frac{1}{k}$ over action $a = 1$. But, $\frac{f(s|1)}{f(s|0)} = 1$ on $[0, 1]$, and hence this ratio jumps up at $a = 0$, and so it is impossible to motive the agent to choose $a = 0$.

The continuity-of-information assumptions we introduce in Section 4 (see in particular Assumption 4.8) rule out such upward jumps in how difficult it is to reward one action without making another action more attractive.

In our third example, a mechanism with a simple randomized contract is optimal. But when the principal is restricted to deterministic contracts, an optimal mechanism fails to exist. If the principal is restricted further to deterministic contracts with uniformly bounded variation, existence is restored, but a tight variational bound implies substantial losses for the principal. As the bound tends to infinity, the principal's payoff converges to the fully optimal payoff with randomization. However, in the limit, the contract oscillates wildly, becoming arbitrarily complex. So, in this example, the restriction to "simple" deterministic contracts either substantially reduces the principal's payoff or substantially increases the complexity of the mechanism.

Example 2.3 *The set of signals is* $S = [0, 1]$. *The set of types is* $T = \{\pm 1, \pm 2, ...\}$, *where* $H(\{t\}) = \frac{1}{2^{|t|+1}}$ *for each* $t \in T$. *The set of actions is* $A = \{-1, +1\}$ *and the set of feasible rewards is* $R = [0, \infty)$. *Given reward* r, *signal* s, *action* a, *and type* t, *the principal's loss is*

$$l(r, s, a, t) = \begin{cases} 0, & \text{if } ta > 0 \\ 1, & \text{if } ta < 0 \end{cases},$$

and the agent's utility is

$$u(r,s,a,t)=\begin{cases}1, & \text{if } a=+1\\ 2-2e^{-r}, & \text{if } a=-1 \text{ and } t>0\\ r, & \text{if } a=-1 \text{ and } t<0\end{cases}.$$

The signal distribution depends only on the agent's type. If $t<0$, the signal is uniform on $[0,1]$. If $t>0$, the signal is uniform on $[a_t,b_t]$, where $[a_1,b_1],[a_2,b_2],...$ is a list of all nondegenerate closed subintervals of $[0,1]$ that have rational endpoints. This captures the idea that the principal is concerned that, for any interval $[a,b]\subseteq[0,1]$, if high rewards are given to the agent when the signal falls into that interval, some undesirable type of the agent will generate a signal that virtually guarantees himself these high rewards.

The following simple mechanism yields the principal her minimum possible expected losses of zero, and so is optimal. No matter what type the agent reports, offer the contract that, regardless of the signal, randomizes with equal probabilities over a reward of 0 and a reward of 3. With this mechanism, all negative (risk neutral) types have a strict incentive to take action $a=-1$, and all positive (risk-averse) types have a strict incentive to take action $a=1$.[11]

Now suppose that the principal is restricted to deterministic contracts. Under this restriction, an optimal mechanism no longer exists. Indeed, the principal can no longer achieve losses of zero,[12] but he can get arbitrarily close, by, for example, choosing n sufficiently large and rewarding the agent with $r=0$ when the n-th digit in the binary expansion of the signal is even and rewarding the agent with $r=3$ when it is odd. As n grows large, expected losses approach zero, with the contract oscillating arbitrarily often between 0 and 3 as the signal moves from 0 to 1.

To restore existence, let us follow the literature (e.g., Holmstrom (1979)) and restrict the contract space further to the (compact) set containing all deterministic contracts whose total variation is less than some uniform bound $B>0$. By Page (1991), an optimal mechanism exists within this restricted class of mechanisms. But, if B is small, the principal's losses can be substantial, while as discussed above, as B tends to infinity, the contract becomes arbitrarily complex as it attempts to synthesize as much "randomization" as possible.

[11] The example is robust to a variety of perturbations, including rewards that are costly to the principal.

[12] Suppose, to the contrary, that some such mechansim yields zero expected losses. Fix some $\hat{t}<0$, and let $c:S\to R$ be the deterministic contract that is offered when the agent reports type $\hat{t}$. For $\hat{t}$ to want to choose $a=-1$, $\hat{t}$ must not prefer $a=+1$. Hence, we must have $E(c(s)|s\sim U(0,1))\geq 1$. For all positive types to want to choose $a=+1$, they must not prefer to report $\hat{t}$, and take action $a=-1$. Hence, we must have $E(e^{-c(s)}|s\sim U(a_t,b_t))\geq\frac{1}{2}$, for every $t\in\{1,2,...\}$. This implies that $e^{-c(s)}\geq 1/2$ for (Lebesgue) almost every $s\in[0,1]$, contradicting $E(c(s)|s\sim U(0,1))\geq 1$.

Finally, let us remark that the situation can be worse still. If, when $at < 0$, the loss to the principal is $2^{|t|}$ instead of 1 (a specification permitted by our model), then every deterministic contract (with bounded variation or not) results in infinite losses for the principal, while the simple randomized contract still yields the minimum possible losses of zero.

3 A Principal-Agent Model

In this section, we present the (single-agent) model in its most basic form, making only enough assumptions so as to permit a statement of the problem.

The sets of rewards R, signals S, actions A, and types T, are measurable spaces, i.e., each set is endowed with its own sigma-algebra of measurable subsets. All product sets are endowed with their product sigma-algebras. The type space, T, is also endowed with a probability measure H, i.e., a prior, on its measurable sets. Typical elements of R are denoted by lower case letters r, r', r'', and typical sequences in R are denoted by r_n etc., and similarly for typical elements of and sequences in S, A, and T. Thus for example, whenever convenient, we may abbreviate statements such as "for all $s \in S$" by "for all s" etc.

The set of feasible rewards can depend on the signal and is captured by Φ, a measurable subset of $R \times S$. For each signal s, the set of feasible rewards given the signal s is $\Phi_s = \{r : (r, s) \in \Phi\}$, which we assume to be nonempty.

For any measurable space X, let $\Delta(X)$ denote the set of probability measures on the measurable subsets of X. If Y is any other measurable space, a *transition probability* is a mapping, γ say, from Y into $\Delta(X)$ such that for every measurable $E \subseteq X$, $\gamma(E|y)$ is a measurable function of $y \in Y$.

The *signal technology* is given by a transition probability, P, from $A \times T$ into $\Delta(S)$. We write $P_{a,t}$ or $P(\cdot|a, t)$ for the value in $\Delta(S)$ of P at any action and type pair (a, t). That is, if the agent's type is t and he takes action a, then the signal is generated according to the probability measure $P_{a,t}$.

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 measurable.

By a straightforward application of the revelation principle (Myerson 1982), our space of mechanisms is as follows.

Definition 3.1 *A mechanism is any (α, κ) such that $\alpha : T \to \Delta(A)$ and $\kappa : S \times A \times T \to \Delta(R)$ are transition probabilities satisfying $\kappa(\Phi_s|s, a, t) = 1$ for every $(s, a, t) \in S \times A \times T$.*

A mechanism (α, κ) works as follows. Nature draws the agent's type from T according to H. After learning his type, t, the agent reports a type, t', to the mechanism. The mechanism then recommends to the agent an action a' that is generated by the probability measure $\alpha(\cdot|t') \in \Delta(A)$. After learning the recommended action a', the agent chooses an action a from A. Finally, given the signal s generated by $P_{a,t}$, the mechanism generates the agent's reward $r \in \Phi_s$ according to the probability measure $\kappa(\cdot|s, a', t') \in \Delta(R)$. So, in particular, signals are generated according to the true type and action of the agent, while rewards depend on the reported type and recommended action.

Given (a', t'), we interpret the transition probability $\kappa_{a',t'} : S \to \Delta(R)$ as a "contract" in which the principal may randomize over the rewards offered to the agent as a function of the observed signal. In this interpretation, the agent knows the (randomized) contract before choosing his action.[13] Henceforth, a *contract* is a transition probability mapping S into $\Delta(R)$.

Using the revelation principle once more, we may restrict attention to mechanisms that are incentive compatible.

Definition 3.2 *A mechanism* (α, κ) *is incentive compatible if for* H*-almost every type* t*, and for every type* t'*,*

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

$$\geq \int_A \left(\sup_{a\in A} \int_{R\times S} u(r, s, a, t)\, d\kappa(r|s, a', t')\, dP(s|a, t) \right) d\alpha(a'|t').^{14}$$

Denote the set of all incentive-compatible mechanisms by M*.*

The left-hand side of the inequality in Definition 3.2 is the utility to the agent of type t

[13] A second interpretation (see. e.g., Myerson, 1982) is that $\kappa_{a',t'} : S \to \Delta(R)$ represents a randomization over a collection of deterministic contracts, each of which is a function mapping the observed signal into a reward (e.g., for each measurable set E in R, $\kappa(E|s, a', t') = \lambda(\{\omega \in [0, 1] : r(\omega, s, a', t') \in E\})$, where λ is Lebesgue measure and $\{r(\omega, \cdot, a', t')\}_{\omega\in[0,1]}$ is a collection of deterministic contracts, one of which is chosen randomly through the choice of ω). In this second interpretation, given (a', t'), the agent must choose his action *without knowing the deterministic contract*, $r(\omega, \cdot, a', t') : S \to R$, that he will eventually face (since he does not know ω). Either interpretation is acceptable, and the agent's set of optimal actions is the same regardless of which interpretation is chosen.

[14] For any $a, a' \in A$, let $g(a, a') = \int_{R\times S} u(r, s, a, t)\, d\kappa(r|s, a', t')\, dP(s|a, t)$. Then, g is Borel measurable, but $G(a') = \sup_a g(a, a')$, which is the function of a' in parentheses on the righthand side of the displayed inequality in Definition 3.2, might not be. However, for any real c, the set $\{a' : G(a') > c\}$ is analytic because it is the projection onto the first coordinate of the Borel set $\{(a, a') : g(a, a') > c\}$. Consequently, letting $\mathcal{B}(A)$ denote the Borel subsets of A, G is measurable with respect to the completion of the measure space $(A, \mathcal{B}(A), \alpha(\cdot|t'))$ and it is with respect to this completion that the outermost integral over A on the righthand side is to be understood. That said, under the assumptions that we shall make, there is an optimal mechanism in which G is in fact Borel measurable and so, at the optimum, there is no need for this technical caveat (see fn. 44).

from reporting his true type to the mechanism and taking the recommended action, while the right-hand side is the utility to the agent of type t from reporting that his type is t', and then choosing an optimal action when the mechanism recommends action a'.[15]

For any incentive compatible mechanism $(\alpha, \kappa) \in M$, let

$$L(\alpha, \kappa) \equiv \int\limits_{R\times S\times A\times T} l\left(r, s, a, t\right) d\kappa\left(r|s, a, t\right) dP(s|a, t) d\alpha\left(a|t\right) dH\left(t\right), \tag{3.1}$$

be the principal's expected loss when the agent reports honestly and takes the recommended action. Then, the principal's problem is

$$\min_{(\alpha,\kappa)\in M} L(\alpha, \kappa). \tag{3.2}$$

Remark 3.3 *The above specification of the principal's problem does not explicitly include participation constraints for the agent. An outside option that is always available to the agent can be modeled by simply including it as an action in A. The model also captures settings in which the principal can 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 11, we extend the model and our existence result to settings that include an ex-ante participation constraint for the agent of the form*

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

We next provide general conditions under which a solution to problem (3.2) exists.

4 Assumptions and the Main Result

In this section we state our assumptions and our main result on the existence of an optimal mechanism. Section 5 provides a discussion of the assumptions and also includes a variety of examples.

Recall that a Polish space is a separable topological space that can be metrized by means of a complete metric. The measurable subsets of any Polish space X will always be the Borel sets $\mathcal{B}(X)$, and $\Delta(X)$ will denote the space of probability measures on $\mathcal{B}(X)$ endowed with

[15] It would have been equivalent to define incentive compatibility in the weaker sense that, for some H-measure one set of types T^0, the inequality in Definition 3.2 holds only for every $t, t' \in T^0$. This is because one can, for any $t^0 \in T^0$, 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 3.2, which we find more intuitive.

the topology of weak convergence (the weak* topology).

Our assumptions are as follows.

Assumption 4.1 *R, S, A, and T are nonempty Polish spaces, and A is compact.*

Assumption 4.2 *Φ is a closed subset of $R \times S$, where for each s, $\Phi_s = \{r : (r, s) \in \Phi\}$ is non-empty.*

Assumption 4.3 *$P : A \times T \to \Delta(S)$ is a transition probability such that $P_{a,t}$ is continuous in a for each t.*

Assumption 4.4 *$u : R \times S \times A \times T \to \mathbb{R}$ and $l : R \times S \times A \times T \to \mathbb{R}$ are measurable and bounded below, without loss of generality by 0, and, for every $(a, t) \in A \times T$ and for $P_{a,t}$ a.e. $s \in S$, $u(\cdot, t)$ is continuous at (r, s, a) and $l(\cdot, t)$ is lower semicontinuous at (r, s, a) for all $r \in R$.*

Assumption 4.5 *For any type t, for any $c \in \mathbb{R}$, and for any compact subset Y of S, the closure of $\{(r, s, a) \in \Phi \times A : s \in Y \text{ and } l(r, s, a, t) \leq c\}$ is compact.*

Assumption 4.6 *For any type t, and for any sequence (r_n, s_n, a_n) in $\Phi \times A$, if $u(r_n, s_n, a_n, t) \to \infty$, then $l(r_n, s_n, a_n, t)/u(r_n, s_n, a_n, t) \to \infty$.*

Assumption 4.7 *There is a collection $\{S_{a,t}\}_{(a,t)\in A\times T}$ of subsets of S such that $\{(s, a, t) : s \in S_{a,t}\}$ is a measurable subset of $S \times A \times T$; $P_{a,t}(S_{a,t}) = 1$ for all a, t; and $P_{a',t'}(E) = 0 \Rightarrow P_{a,t}(E) = 0$ for all a, t, a', t' and for all measurable $E \subset S_{a',t'}$.*[16]

Under Assumption 4.7, $P_{a,t}$ is absolutely continuous with respect to $P_{a',t'}$ when both measures are restricted to the measurable subsets of $S_{a',t'}$. Equivalently, the unrestricted measure $P_{a,t}(\cdot \cap S_{a',t'})$ on the measurable subsets of S is absolutely continuous with respect to the unrestricted measure $P_{a',t'}$. Applying the Radon-Nikodym theorem to the unrestricted measures, a Radon-Nikodym derivative defined on all of S exists and so we denote by $g_{a,t/a',t'} : S \to [0, +\infty]$ any version of the Radon-Nikodym derivative of $P_{a,t}(\cdot \cap S_{a',t'})$ with respect to $P_{a',t'}$.[17]

Assumption 4.8 *For all a, t, a', t', and for every sequence $a'_n \to a'$, there is $a_n \to a$ and there are versions of the Radon-Nikodym derivatives $g_{a_n,t/a'_n,t'}$ and $g_{a,t/a',t'}$ such that,*

$$\underline{\lim}_n g_{a_n,t/a'_n,t'}(s_n) \geq g_{a,t/a',t'}(s) \ \forall s \in S \text{ and } \forall s_n \to s. \tag{4.1}$$

[16] The measurability of $\{(s, a, t) : s \in S_{a,t}\}$ implies the measurability of each of its slices $S_{a,t}$. See the discussion in Section 5 for why it is useful to allow the $S_{a,t}$ to differ from the support of $P_{a,t}$.

[17] Any two versions of the Radon-Nikodym derivative are $P_{a',t'}$ almost everywhere equal. They can also be chosen to be $P_{a',t'}$ almost everywhere finite, but it is useful to allow values of $+\infty$ on measure zero sets. See Remark 5.7.

Assumption 4.9 *If the collection* $\{S_{a,t}\}$ *in Assumption 4.7 is such that* $P_{a,t}(S_{a',t'}) < 1$ *for some* a, t, a', t', *then there is* $r_* \in R$ *such that* $r_* \in \Phi_s$ *for all* $s \in S$, *and* $u(r_*, s, a, t) \leq u(r, s, a, t)$ *for all* $(r, s, a, t) \in \Phi \times A \times T$.

Remark 4.10 *For numerous applications, it is possible to verify all nine of our assumptions by recasting one's model so that only Assumptions 4.1-4.6 need be checked. One such leading case is when all spaces are Euclidean and the signal distribution admits a density that is continuous in* s, a, *and* t.[18] *See Corollary 6.1 and Remark 6.2 in Section 6 for this and more general such cases.*

We can now state our main result.

Theorem 4.11 *If Assumptions 4.1-4.9 are satisfied, then, provided that there is at least one incentive-compatible mechanism, the principal's problem (3.2) has a solution.*

Theorem 4.11 follows directly from Theorem 8.3 in Section 8.1, according to which the set of incentive compatible mechanisms can be metrized so that (i) for any constant, the subset of incentive compatible mechanisms that bounds the principal's expected losses weakly below that constant is compact, and (ii) the principal's loss function is lower semicontinuous. The existence of a loss-minimizing incentive compatible mechanism then follows immediately if the set of incentive compatible mechanisms is nonempty. While in applications it is often easy to establish the existence of at least one incentive compatible mechanism, Proposition 8.4 provides general conditions under which an incentive compatible mechanism exists.

5 Discussion of the Assumptions and Examples

Simple examples of Polish spaces include $\mathbb{R}$, $[0, 1]$, $\mathbb{Z}$, and a variety of function spaces (e.g., any $\mathcal{L}_p$ space, the space of continuous functions on any compact metric space, and others),[19] as well as any finite or countable products of them. Hence, all of the spaces R, S, A, and T can include multiple dimensions, some of which may be discrete and some of which may be continuous. Only A is required to be compact.

The definition of Φ allows considerable flexibility regarding the rewards that are feasible as a function of the observed signal. Of course, the set of feasible rewards need not depend on s, in which case Φ_s is a fixed closed subset of R. For example, for $m \geq 0$ setting $\Phi_s = [m, \infty)$ for all s captures a minimum payment constraint, with the case $m = 0$ corresponding to a limited liability constraint.

[18] Note that there is no requirement here of a constant support.

[19] Such function spaces can be useful, for example, when modeling settings in which the agent may be rewarded with stock options, which are real-valued functions of some future state variable.

Assumption 4.3 imposes the minimal requirement that as a varies for any given t, $P_{a,t}$ moves continuously in the weak* topology. This is automatic if $P(\cdot|a,t)$ can be represented by a density that is continuous in a. Example 2.1 fails this assumption.

The continuity assumptions on $u(\cdot)$ and $l(\cdot)$ are slightly more permissive than standard assumptions since we allow some discontinuities. See Corollary 6.1 and Remark 6.2 for how this additional generality can be helpful in applications. 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. With enough separability, payoffs that are not bounded below can nevertheless be handled. We give two examples.

Example 5.1 ***Unbounded Losses.*** *$S = (-\infty, \infty)$, $R = [0, \infty)$, $T = \{t_0\}$ 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 5.2 ***Unbounded Utility.*** *$S = (0, \infty)$, $R = [0, \infty)$, $T = \{t_0\}$ 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.*

When all spaces are Euclidean, Assumption 4.5 says that, for any fixed $t \in T$, the principal's losses are unbounded along any sequence in $R \times S \times A$ in which only the component corresponding to the agent's reward is unbounded. This allows that unbounded rewards to the agent might not lead to unbounded losses when, in addition, the signal is unbounded along the sequence since, for example, the signal may indicate the firm's revenue.

Assumption 4.6 says that the losses to the principal per unit of utility provided to the agent are unbounded above when the agent's utility is arbitrarily high. If the agent's utility is bounded, this is satisfied trivially. A typical applied setting has $R = [0, \infty)$, $l(r, s, a) = r - s$, and $u(r, s, a) = v(r) - c(a)$, where c is continuous and v is differentiable. If $\lim_{r\to\infty} v'(r) = 0$, then Assumption 4.6 is satisfied. On the other hand, if, for example, $v(r) = r - \frac{1}{r}$, then Assumption 4.6 fails.

Next, let us consider Assumption 4.7. For any a and t, the set $S_{a,t}$ can be interpreted as the set of signals that are considered possible when the agent's type is t and he takes action a. Under Assumption 4.7, the specification of rewards on any subset of $S_{a',t'}$ that has measure zero (and so is irrelevant) for a compliant agent with true action/type (a', t') is also irrelevant for any non-compliant agent with true action/type $(a, t) \neq (a', t')$. Notice that Assumption 4.7 is automatic if S is Euclidean and for each a, t, $P(\cdot|a, t)$ can be represented by a density that is positive on $S_{a,t}$.

Remark 5.3 *It can always be assumed that $S_{a,t}$ is contained in the support of $P_{a,t}$ (supp$P_{a,t}$) since if the collection $\{S_{a,t}\}$ satisfies Assumption 4.7 then so does the collection $\{S_{a,t} \cap \text{supp} P_{a,t}\}$.*[20]

Our next example illustrates that it can be useful to allow $S_{a,t}$ to be a proper subset of the support of $P_{a,t}$

Example 5.4 *Let $S = A = [0, 1]$, and for every measurable $E \subseteq S$, let*

$$P(E|a, t) = a\mathbf{1}_E(0) + (1 - a)\int_E f(s|t)\, ds,$$

where $f(\cdot|t)$ is a positive density on $[0, 1]$ for each t. If we insist that $S_{a,t}$ be the support of $P_{a,t}$, then $S_{a,t} = [0, 1]$ for all $a < 1$, $S_{a,t} = \{0\}$ for $a = 1$, and the absolute continuity part of Assumption 4.7 will fail for $a' = 0$, $a = 1$, and $E = \{0\}$. However, if for $a < 1$ we instead set $S_{a,t} = (0, 1]$, then all of our informational assumptions are satisfied.

Our main continuity of information condition is Assumption 4.8. This assumption prevents upward jumps in how difficult it is to reward any given type for a particular action without making a different action more attractive for that type or for some other type. In common applications the Radon-Nikodym derivative that is featured in this assumption simply reflects the ratio between the densities on signals induced by the different actions and types, and, in these applications the assumption is satisfied if the densities are bounded away from zero, and are, for each t, continuous in a and s.

A simple sufficient condition for Assumption 4.8 is the following.

Assumption 4.8′ For every a, t, and t' the Radon-Nikodym derivative $g_{a,t/a',t'}(s)$ is lower semicontinuous as a function of (s, a') on $S \times A$.

Remark 5.5 *In a pure screening problem (i.e., A is a singleton set) Assumption 4.8 reduces to the condition that $g_{t/t'}(s)$ is lower semicontinuous in s and so Assumptions 4.8 and Assumption 4.8′ become equivalent.*

[20] The requisite measurability condition is satisfied because the set $\{(s, a, t) : s \in \text{supp} P_{a,t}\}$ is measurable, being the complement of the union of measurable sets of the form $V \times \{(a, t) : P(V|a, t) = 0\}$, where the union is taken over any countable basis of open sets V for the topology on the Polish space S.

To better understand Assumption 4.8, note that while Example 2.2 satisfies Assumption 4.3, it fails to satisfy Assumption 4.8. 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 $g_{1/\frac{1}{n}}(s) = \frac{1}{2}$. Hence, for some s_n with binary expansion that agrees with $\hat{s}$ to the first $n-1$ digits, $g_{1/\frac{1}{n}}(s_n) = \frac{1}{2}$. But then Assumption 4.8 is violated, since $g_{1/0}(s) = 1$ for almost all s, and $\hat{s}$ was arbitrary.

The next two examples provide some indication of the range of Assumption 4.8.

Example 5.6 ***Observable Actions.*** *Let $S = A = [0, 1]$ and let $P_{a,t}$ be the Dirac measure placing mass one on $s = a$.*

Example 5.6 fails to satisfy the informational assumptions in Page (1987, 1991), since, for example, $0 = P_{1/n,t}(\{0\}) \nrightarrow_n P_{0,t}(\{0\}) = 1$, but it satisfies Assumption 4.8. To see this, let $g_{a,t/a',t'} = 1$ if $a = a'$ and $g_{a,t/a',t'} = 0$ otherwise and consider any a, a' and $a'_n \to a'$. If $a' = a$, then choosing $a_n = a'_n$ will satisfy inequality (4.1), while if $a' \neq a$, then choosing $a_n = a$ will satisfy it. It is important to this example that Assumption 4.8 allows the choice of the sequence a_n to be tailored to the particular sequence a'_n. In particular, Example 5.6 fails to satisfy Assumption 4.8′.[21]

Remark 5.7 *When verifying Assumption 4.8, the versions of all of the Radon-Nikodym derivatives that are employed can depend on (a, t) and (a', t') as well as on the sequences $\{a_n\}$ and $\{a'_n\}$. So, in particular, it is sufficient for inequality (4.1) to hold only for $P_{a',t'}$ a.e. $s \in S$ since, by choosing a version of $g_{a,t/a',t'}$ that is zero on the remaining $P_{a',t'}$-measure zero set of signals, (4.1) will hold for all $s \in S$. Similarly, for any sequence of sets S_n whose complements have $P_{a'_n,t'}$ measure zero, it suffices for (4.1) to hold only for all $s_n \to s$ such that $s_n \in S_n$ for each n because, for each n, we may choose a version of $g_{a_n,t/a'_n,t'}(s)$ that is equal to $+\infty$ on the complement of S_n.*[22]

This brings us to Assumption 4.9. To see what it is saying, suppose that $P_{a,t}(S_{a',t'}) < 1$. Then, if type t takes action a but reported type t' and was recommended action a', there is a positive probability that the signal observed by the principal will be in $S_{a,t} \backslash S_{a',t'}$, i.e., will be considered possible for (a, t) but impossible for (a', t'). Hence, there is positive probability that the principal will infer that the agent either lied or took the wrong action. Assumption 4.9 ensures that the principal has a feasible "worst reward" that is sure to penalize the agent

[21] Any version of the requisite Radon-Nikodym derivative must satisfy $g_{a,t/a',t'}(a') = 1$ if $a = a'$ and $g_{a,t/a',t'}(a') = 0$ otherwise, and so for any sequence of distinct actions $a'_n \to a$, $g_{a,t/a'_n,t'}(a'_n) = 0$ for every n, but $g_{a,t/a,t'}(a) = 1$. Hence, $g_{a,t/a',t'}(s)$ is not lower semicontinuous in (s, a').

[22] Recall that redefining the values of a nonnegative Lebesgue integrable function to be $+\infty$ on a measure zero set does not affect the value of its Lebesgue integral.

in such cases. Notice that Assumption 4.9 is trivially satisfied when $S_{a,t}$ is independent of a, t (since $P_{a,t}(S_{a',t'}) = 1$ for all a, t, a', t'). Also, when Assumption 4.9 fails to hold, the agent's utility function can sometimes be harmlessly modified so that Assumption 4.9 is satisfied.[23]

6 A Standard Case

It is useful to point out how Theorem 4.11 applies to the standard case in the literature in which there is $Q \in \Delta(S)$ such that $P_{a,t}$ is absolutely continuous with respect to Q for every $(a, t) \in A \times T$.[24] In particular, let us suppose that there is a measurable $f : S \times A \times T \to [0, \infty)$ such that for every measurable $E \subseteq S$,

$$P(E|a, t) = \int_E f(s|a, t)dQ(s). \tag{6.1}$$

In this case, the agent's expected payoff from any IC mechanism (α, κ) is

$$\int u(r, s, a, t)f(s|a, t)d\kappa(r|s, a, t)dQ(s)d\alpha(a|t)dH(t).$$

Consequently, the situation for the agent is equivalent to one in which his utility function is uf and the signal is drawn according to Q regardless of his action a and type t. Analogously, the situation for the principal is equivalent to one in which her loss function is lf and the signal is always drawn according to Q.

Since the model with signals always drawn according to Q trivially satisfies Assumptions 4.3-4.9, in this special case an optimal mechanism exists under weak additional assumptions. Indeed, we have the following immediate corollary of Theorem 4.11, which can be quite useful in applications.

Corollary 6.1 *Suppose that Assumptions 4.1-4.6 hold when the agent's utility function is uf, the principal's loss function is lf, and the signal is always drawn according to $Q \in \Delta(S)$. Then, provided that at least one incentive-compatible mechanism exists, an optimal mechanism exists when the agent's utility function is u, the principal's loss function is l, and the signal is drawn according to $P_{a,t}$ given by (6.1).*

Remark 6.2 *(a) If u and l satisfy Assumptions 4.4 and 4.6, then uf and lf also satisfy Assumptions 4.4 and 4.6 if f is bounded, and, for every $(a, t) \in A \times T$, $f(\cdot|\cdot, t)$ is continuous*

[23] Suppose, for example, that $R = [0, \infty)$, $\Phi_s = [s, s + 1]$, and u is increasing in r. Then, there is no fixed r^* in R that satisfies Assumption 4.9. In this case, simply redefine R by appending to it an isolated point r^*, and extend u by defining $u(r^*, s, a, t) = u(s, s, a, t)$ for all (s, a, t). Assumption 4.9 is now satisfied.

[24] See, for example, Page (1987, 1991).

at (s,a) for Q a.e. $s \in S$ (the exceptional set of signals can depend on a and t).[25] *(b) If l satisfies Assumption 4.5, then lf also satisfies Assumption 4.5 if either f is positive and bounded away from zero, or if $\Phi \cap (R \times Y)$ is compact for every compact subset Y of S.*

We give three examples to illustrate Corollary 6.1. In each of them we assume that $0 \in \Phi_s$ for all $s \in S$, that Φ satisfies Assumption 4.2, and that u and l satisfy Assumptions 4.4-4.6. By applying Remark 6.1, each example can be shown to satisfy the hypotheses of Corollary 6.1 and so each example admits an optimal mechanism.[26]

In the first two examples, it should be noted that there are no $S_{a,t}$ sets satisfying Assumption 4.7 that are independent of a and t. Consequently, if we try to apply Theorem 4.11 directly, we would, in particular, need to know that there is a worst reward in order to satisfy Assumption 4.9. By instead using Corollary 6.1, we can avoid Assumption 4.9 entirely.

Example 6.3 *Let $R = A = [0,1]$, $S = [0,2]$, $T = \{t_0\}$, and let P_{a,t_0} be uniform on $[a, a+1]$. Then, we may define Q to be the uniform distribution on $[0,2]$, and we may define $f(s|a,t) = 2$ if $s \in [a, a+1]$ and $f(s|a,t) = 0$ otherwise.*

Example 6.4 ***Discrete signal distributions.*** *Let $R = A = [0,1]$, let S be a finite set and let $P(s|a,t)$ be a continuous function of a for each $(s,t) \in S \times T$. Then we may define $Q(s) = 1/|S|$ for each s, and we may define $f(s|a,t) = P(s|a,t)/Q(s)$. (Notice that f may sometimes be zero.)*

Example 6.5 *Let $R = A = [0,1]$, $S = (-\infty, \infty)$, $T = (1, \infty)$, and let $P_{a,t}$ be a normal distribution with mean a and standard deviation t. The natural candidate for the carrying measure Q here is Lebesgue measure, but it is not a probability measure on S. Instead, define Q to be the probability measure on $\mathbb{R}$ with density $q(s) = e^{-2|s|}$, and define $f(s|a,t) = \frac{e^{2|s|}}{t\sqrt{2\pi}} e^{-\frac{(s-a)^2}{2t^2}}$.*

7 Related Literature and Applications

Our paper is the first to offer a general existence result without imposing onerous restrictions on either the primitives of the model or on the set of allowed mechanisms. As such, it opens

[25] This condition on f (and its implications for $P_{a,t}$) is more restrictive than Page's (1987, 1991) assumption that $P_{a,t}(E)$ is continuous in a for each closed $E \subseteq S$. However, in this absolutely continuous case, some such stronger assumption is unavoidable since, as Example 2.2 shows, Page's more permissive assumption does not suffice for the existence of an optimal mechanism when the contract space is unrestricted, as it is here.

[26] In each example, an incentive-compatible mechanism exists, e.g., the mechanism that, if the agent reports type t, recommends an action that maximizes $\int u(0,s,a,t)dP(s|a,t)$ and that always gives the agent reward $r = 0$.

the door for a range of applications. Previous papers in this area belong to one of the following two groups:

1. **Papers imposing restrictions on the primitives of the model.** Grossman and Hart (1983), establish existence of an optimum in a pure moral hazard problem. 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 by assuming that effort is one-dimensional, that likelihood ratios are monotone and bounded, and that the first-order approach is valid.[27] All three papers require signals and rewards to be one-dimensional (as in Holmström (1979)) and the principal's losses to depend on an additively separable function of them, and none permits the agent to possess private information. Kahn (1993) establishes existence in a pure adverse selection problem, relying on restrictions on the set of types, and on the distributions and utilities considered.

2. **Papers imposing restrictions on the allowed mechanisms.** Holmström (1979), Page (1987,1991), and Balder (1996, Section 3.2) all require the mechanism to employ deterministic contracts that are observed by the agent prior to his action decision and that are contained in some fixed function space that is compact in the topology of pointwise convergence. As some of these papers note, the needed compactness can be obtained by restricting contracts to be of uniformly bounded variation. The actual bound is left unspecified even though it can significantly affect the optimal solution.

In contrast, the approach taken here permits a significant weakening of the standard restrictions on the primitives of the model and dispenses with any restrictions on the set of mechanisms or contracts. This is accomplished by permitting (but not requiring) randomization over the agent's reward and by introducing new continuity of information conditions. As a result, a large set of potential applications is covered and a fully optimal solution to the principal's problem can be studied.

The success of our approach hinges on a key insight due to Page (1991). He showed that a powerful sequential compactness result due to Komlos (1967) and its important generalization by Balder (1990) could be used to establish the existence of a mechanism that is optimal, at least within the restricted class of mechanisms described in item #2 above. Balder (1996) generalizes Page's existence result within the same restricted class of mechanisms.

[27] Conditions facilitating the first order approach are typically quite demanding. See Mirrlees (1976), Rogerson (1984), Jewitt (1988), Sinclair-Desgagne (1994), Conlon (2009)), and Chade and Swinkels (2016).

Our existence result has several immediate applications to the literature seeking to go beyond the restrictive first order approach (FOA). Chaigneau, Edmans, and Gottlieb (2014) study Holmstrom's (1979) informativeness principle in the pure moral hazard problem when the FOA fails. Kierkegaard (2014) studies a moral hazard problem in which the FOA does not necessarily hold using a spanning condition. Kadan and Swinkels (2013) study properties of optimal contracts in the moral hazard problem when the FOA may fail and provide several comparative statics results. Renner and Schmedders (2015) present a computational method for providing an approximate solution to moral hazard problems not relying on the FOA. Ke and Ryan (2015) present a general methodology for solving moral hazard problems without assuming the FOA. All of these papers require, but do not include, an existence result. Our paper supplies the missing result.

Several papers in the literature study specialized principal-agent models with moral-hazard and/or adverse-selection and allow various forms of randomization, including randomization over the contract and randomization over rewards. These papers do not, however, establish the existence of an optimal mechanism. Examples are Gjesdal (1982), Fellingham, Kwon, and Newman (1984), Arnott, and Stiglitz (1988), and Strausz (2006). In the next section, we discuss several examples in which randomized mechanisms are economically natural.

Page (1994) and Balder (1996) contain existence results for the problem of a principal interacting with multiple agents. They are able to establish the existence of a mechanism that is optimal for the principal among all dominant strategy mechanisms. However, their techniques do not extend to the problem of Bayesian incentive compatible mechanisms. In Section 13, we show how our existence result for the single-agent model implies the existence of a Bayesian incentive compatible (BIC) mechanism that is optimal for the principal among all BIC mechanisms even when there are multiple agents.

8 Randomization

The principal in our model 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 will typically depend on this recommendation, will have the effect of randomizing over the agent's contract.[28] Second, after the agent chooses his action and a signal is generated, the contract itself may specify a randomization over the agent's reward. As noted, the most recent and closely related literature on the existence of an

[28] In fact, randomization over the recommended action that does not induce randomization over the contract is never necessary for optimality. So, whenever randomization over the recommended action is mentioned, the reader should keep in mind that this effectively means randomizaiton over the contract.

optimal mechanism (Page (1991) and Balder (1996)) has embraced the first opportunity for randomization but has not permitted the second, i.e., they have restricted attention to deterministic contracts. In this section, we discuss the importance, both practical and theoretical, of permitting both forms of randomization. But first, a clarification is in order.

Whether the contract that is offered to the agent is deterministic (i.e., is a function from the signal into a deterministic reward) or is randomized (i.e., is a function from the signal to a probability distribution over rewards), we always assume that the agent is fully informed of the contract (i.e., the function) prior to taking his action. This assumption is without loss of generality in our setting that allows randomized contracts because any uncertainty over the contract is equivalent to a known contract that randomizes over the reward. On the other hand, this assumption, which Page (1991) and Balder (1996) both impose, is not without loss of generality when the contract is restricted to being deterministic. In general, the principal can strictly gain by not informing the agent of which deterministic contract is in effect until after the agent has chosen his action.

Both opportunities for randomization can strictly benefit the principal.

As already well-recognized in the literature (e.g., Gjesdal 1982 and Arnott and Stiglitz 1988) randomization can strictly benefit the principal. We have already seen (Example 2.3) that differences in risk aversion across agent types can make randomization over rewards beneficial for the principal. We next provide an example showing how randomization over the recommended action (i.e., randomization over the contract) can be beneficial. In this example, the best way to dissuade the type 1 agent from announcing that he is type 2 is to occasionally ask an agent who announces type 2 to take an observable action that, while suboptimal for the principal, is sufficiently unpleasant for the agent when he is type 1.

Example 8.1 ***Randomization over recommended actions.*** *There are two equiprobable types t_1 and t_2, and three actions a_1, a_2 and a_3. 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 principal's loss if the given type takes the given action and the reward is r.*

	a_1	a_2	a_3
t_1	r / r	$1+r$ / 10	1 / 10
t_2	r / 10	r / r	2 / 10

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 $r_2 \geq r_1$.[29]

Consequently, in any *IC* mechanism that does not randomize over recommended actions, either type t_1 chooses action a_2 or a_3, or type t_2 chooses action a_1 or a_3. This implies that, without randomization, the principal's expected loss is at least $10/2 = 5$.

But consider the following incentive compatible mechanism that involves randomization over recommended actions. If t_1 is reported, then the mechanism recommends action a_1 and selects the contract that pays $r = 2$ if a_1 is observed and pays $r = 0$ if any other action is observed. If t_2 is reported then a lottery occurs. One-half of the time the mechanism recommends action a_2 and selects the contract that pays the agent $r = 2$ if a_2 is observed and pays $r = 0$ otherwise. The other half of the time the mechanism recommends action a_3 and selects the contract that pays the agent $r = 0$ regardless of the action that is observed. With this mechanism, the principal's expected losses improve to $4 < 5$.[30]

Randomization is economically natural in many (but not all) real-world economic settings.

Another important reason for allowing randomization within the theory is based on the amount of randomization actually used in the economy. For example, in academia, at tenure time there are essentially only two rewards, "tenure" or "fire." Yet, over an intermediate range of performance, the outcome of the case - the reward - is stochastic, depending on non-contractable features such as who is present at the meetings, who reads the case, who writes the evaluation letters, etc. While much of this randomness is intrinsic, it is not clear that the university would be better off if it was eliminated, since the effect might be that many people who knew that their case was far enough from the line, on either side, would simply stop exerting effort.

Good audit strategies are typically stochastic. A given action (say a report of income) generates a signal to the tax authority that is met with the reward of either an accepted tax return and the associated tax liability, or the reward of an audit, in which case the agent suffers a penalty if his true type is different than the action taken. A buyer making a reservation on Priceline has a choice between simply reserving a room at their favorite hotel in the relevant area, or getting a lower rate if they agree to learn the name of the hotel only after they commit to purchase. This randomness presumably helps to sort types who differ

[29] 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.

[30] With some effort, one can show that the above randomized mechanism is optimal, and that the example can be modified to a setting with convex type and action spaces.

in how finicky they are about the specific hotel.

Whether the randomness seen by the buyer reflects actual randomization deep within a Priceline computer, or a deterministic function from inventory levels to hotel assignments, where neither the current inventory levels nor the function is well understood by the buyer, seems inconsequential. Ultimately, Priceline's design requires its users to either choose an outcome that is random from their perspective, or pay extra to ensure their preferred choice. This deliberate design allows Priceline to separate types with different willingnesses-to-pay to avoid the uncertainty. It seems the same to us whether the randomness seen by the buyer is achieved by Priceline refusing to release information it has (on, say, available inventory) or by performing a coin flip, and we have little doubt that they are up to the task in either event.

The above are examples of randomization over rewards, the kind of randomization that we permit but that the recent literature has not. For an example of randomization over both contracts and rewards, consider a starting employee washing and parking cars at Enterprise Rent-A-Car Corporation, which recruits management on a promote-from-within basis. Rewards in the entry-level job are poor compared to the outside option of the typical (college graduate) hiree. One reason that the employee is willing to take the job is that there is a reasonable chance that they will soon get promoted into a higher level position, one with more aggressive performance pay. Whether that promotion occurs depends on any number of factors that are random from the point of view of the new employee, such as whether the immediate superior to the employee is likely to be terminated or promoted, or whether a new branch is likely to be opened in the area.[31] Rather than resolve any of this uncertainty, Enterprise may be better off to leave it unresolved, and use the lottery as a way to satisfy the participation constraint of a broader pool of entry-level workers.[32]

We are not arguing that every economically interesting environment is one where randomization is natural. In many realistic settings the underlying economic or social context precludes randomization at one or both relevant stages. For example, in academia, there is randomization over rewards at intermediate performance levels, but a university is likely to feel constrained not to randomize over the contracts that junior faculty receive. It would probably be viewed as repugnant if an insurance company randomized over whether an applicant who reported a specific health history received liver transplant coverage. However, on the occasion that a liver transplant is actually needed, whether one is available is indeed random, and priority in the transplant queue does depend on the health history. In each

[31] If effort during the time spent in the entry-level position is a critical part of the picture, then one might wish to model this as a multi-period mechansim, which moves beyond the scope of this paper.

[32] Note that in settings of this sort, randomization over contracts can be valuable even without a type space.

of these two cases, the right modeling choice may be to assume that randomization over contracts is not allowed, but that contracts can randomize over rewards. In other settings, it may be that it is reasonable to randomize over contracts, but not reasonable that the contract randomizes over rewards for any given outcome. Our model and main results assume that both kinds of randomization are possible. Nevertheless, in Section 12.2, we are also able to shed light on the existence question when one or both kinds of randomization are not permitted.

The restriction to deterministic contracts can lead to complex solutions that obscure the underlying economics.

In a pure moral hazard context, Harris and Raviv (1976) and Holmström (1979) restrict attention to deterministic contracts. This is perfectly justified in the context of their particular models because the principal and the agent are risk averse, and payoffs are additively separable in effort and rewards, and hence randomization over the agent's reward is never helpful (Holmström (1979), Proposition 3). Even so, because the set of deterministic contracts is not compact in a useful topology, proving the existence of an optimal mechanism is not possible without further assumptions. To address this difficulty, Holmström (1979) restricts the set of contracts further by requiring them to be functions of uniformly bounded variation, remarking (fn. 10) that for a sufficiently large bound, this restricted set of contracts "...will contain all functions of practical relevance."

When randomization hurts the principal, there is no harm in including it since it will never be used. But in settings in which randomization strictly benefits the principal, the restriction to contracts with uniformly bounded variation may in fact exclude contracts of practical relevance. In Example 2.3, for any given bound, the optimal contract uses all the available variation, oscillating wildly in an attempt to replicate as much of the randomness as possible. All the while, a simple randomized contract exists.

Randomization permits an economically relevant and a mathematically useful topology on mechanisms.

To establish the existence of an optimal mechanism, we take the standard "continuity and compactness" approach. This requires that we find a topology under which the space of incentive compatible mechanisms is compact and the principal's loss function is lower semicontinuous. But how should this topology be chosen? Consider the following example.

Example 8.2 *An agent of a single type takes an action a from the set $A = S = R = [0, 2]$. When he chooses $a \in (0, 2]$ the signal s is uniform on $[0, a]$ with density $\frac{1}{a}$. When he chooses $a = 0$ the signal s is zero with probability 1, i.e., his action is revealed. The agent's utility is $u = a + 2r$ and the principal's losses are $l = a + r$.*

This example has a simple optimal solution,[33] yet it is rich enough to show that both the economics and the mathematics of the problem help to inform the choice of the topology, and that randomization plays a central role in this choice.

Consider a sequence of incentive compatible mechanisms where in mechanism n, the principal recommends the action $a = 2/n$, and pays the agent according to the contract $c_n : S \to R$ given by

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

The principal's expected loss from mechanism n is $1 + \frac{2}{n}$. If our continuity and compactness approach is to be successful, this sequence of mechanisms must have a subsequence that admits a limit mechanism that is incentive compatible and in which the principal's losses are no greater than 1, the limit of her losses along the subsequence. Under what topology is this the case?

Given the previous literature (e.g., Page (1987,1991)), it is natural to begin with the topology of pointwise convergence applied to $\{c_n\}$. Since $c_n(s) \to_n 0$ for every s, c_n converges pointwise to the contract $c^* \equiv 0$ in which the agent is paid zero regardless of the signal. But given the contract c^*, the agent's uniquely optimal action is $a = 2$ leading to losses of 2 for the principal. Thus, even though the sequence of contracts c_n has a pointwise limit, the limit contract c^* creates a discontinuous upward jump in the principal's losses. So, the topology of pointwise convergence applied to the contract space does not work.[34]

Instead, for each n, let us consider the (ex-ante) joint distribution, β_n, over rewards, signals and actions, (r, s, a), that is induced by c_n. Under c_n, the agent takes action $a = 2/n$ and the signal is equally likely to be in $[0, 1/n)$, in which case $r = 0$, as it is to be in $[1/n, 2/n]$, in which case $r = 2$. So the joint distribution, β_n, on $R \times S \times A$ gives probability $1/2$ to $\{(r, s, a)\} = \{0\} \times [0, 1/n) \times \{2/n\}$ and gives probability $1/2$ to $\{(r, s, a)\} = \{2\} \times [1/n, 2/n] \times \{2/n\}$. In the spirit of Migrom and Weber (1985), we may call β_n a distributional mechanism.[35] Then, β_n converges in the weak* topology to the probability measure $\beta^* \in \Delta(R \times S \times A)$ that assigns probability $1/2$ to $(r, s, a) = (0, 0, 0)$ and probability $1/2$ to $(r, s, a) = (2, 0, 0)$. In particular, $(s, a) = (0, 0)$ with probability one at the limit.

Thus, under the weak* topology, the given sequence of distributional mechanisms yields a well-defined limit distribution β^* over rewards, signals and actions. While β^* does not itself constitute a mechanism, it does generate a natural guess for one. The principal should

[33] The prinicpal should ask the agent to take action $a = 0$ and the principal should give the agent an expected reward of 1 if the signal $s = 0$ occurs and a reward of 0 otherwise.

[34] Things can be even worse for the topology of pointwise convergence since some sequences of contracts (e.g., such as in Example 2.3) have no pointwise convergent subsequence at all.

[35] For simplicity, the present example is one of pure moral hazard. In general a distributional mechanism will also include the distribution over types.

recommend the action $a = 0$, and, if the signal $s = 0$ is observed – the only signal given positive probability by β^* – the agent should be paid randomly, receiving $r = 0$ or $r = 2$ each with probability 1/2. Since β^* places all weight on $s = 0$, it does not pin down rewards when $s \neq 0$. However, since these signals don't occur when the agent chooses $a = 0$, setting rewards to $r = 0$ for any such s is a natural choice, since the only role of rewards at such "out-of-equilibrium" signals is to discourage non-compliant actions. The mechanism so constructed from β^* is incentive compatible and generates expected losses of 1 for the principal, as desired.

We take several lessons from this example. First, sequences of deterministic mechanisms, even those with seemingly natural deterministic limits, may have no economically relevant deterministic limit, but do have economically relevant limits that may involve randomization. Second, when mechanisms are interpreted in terms of the joint distributions that they induce on rewards, signals and actions, the mathematically natural mode of convergence is weak* convergence of measures. In this example, weak* convergence not only provides a limit distribution that generates a natural candidate for a limit mechanism, this limit mechanism has the crucial properties that it is incentive compatible and does not cause the principal's losses to discontinuously jump up, precisely as needed for the efficacy of the continuity and compactness approach to existence.

The heart of the construction in this paper is to generalize this example by showing that, under weak assumptions, and with the weak* topology on the joint distribution over rewards, signals, actions, and types induced by the mechanism, the space of incentive compatible mechanisms is compact and the principal's loss function is lower semicontinuous.

8.1 The Metric Space of Distributional Mechanisms

The topology of weak* convergence of measures has two key advantages. First, closed subsets of probability measures on compact sets are compact, and second, expectations of bounded continuous functions are continuous in the underlying measure. It is therefore natural to define a metric on the space of mechanisms in terms of the measures they induce on the ambient space $R \times S \times A \times T$.

If X and Y are measurable spaces, if $\eta \in \Delta(X)$, and if $\gamma : X \to \Delta(Y)$ is a transition probability, define $\eta \otimes \gamma \in \Delta(Y \times X)$ such that for all measurable sets $B \subseteq X$ and $C \subseteq Y$,

$$(\eta \otimes \gamma)(B \times C) = \int_B \gamma(C|x) d\eta(x).$$

That is, $\eta \otimes \gamma$ is the joint probability measure on $Y \times X$ induced by the marginal probability measure $\eta(\cdot) \in \Delta(X)$ and the "conditional" probability measure $\gamma(\cdot|x) \in \Delta(Y)$ for each

$x \in X$.

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 incentive-compatible (IC) mechanisms, M, by

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

That is, the distance between two mechanisms is determined by the distance between the probability measures that the two mechanisms induce on $R \times S \times A \times T$.

Under d_M, two IC mechanisms (α, κ), $(\alpha', \kappa') \in M$ are considered equivalent if $\alpha(\cdot|t) = \alpha'(\cdot|t)$ for H a.e. $t \in T$ and $\kappa(\cdot|s, a, t) = \kappa'(\cdot|s, a, t)$ for $H \otimes \alpha \otimes P$ a.e. $(s, a, t) \in S \times A \times T$. That is, d_M considers equivalent any two mechanisms whose distributions on $R \times S \times A \times T$ are the same when the agent always reports his type honestly and takes the recommended action, even though the two mechanisms might be different when the agent falsely reports his type or fails to take the recommended action.[36] Both the principal and the agent are indifferent between any two IC mechanisms that are equivalent in this sense.

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. As such, for any $(\alpha, \kappa) \in M$, let us call any probability measure in $\Delta(R \times S \times A \times T)$ of the form $H \otimes \alpha \otimes P \otimes \kappa$, a *distributional mechanism.*

Suppose that the set M of IC mechanisms is nonempty. If every mechanism in M yields the principal infinite losses, then all the mechanisms in M are (trivially) optimal. Otherwise, the search for an optimal mechanism can be restricted to a nonempty subset of M,

$$M_c \equiv \{(\alpha, \kappa) \in M : L(\alpha, \kappa) \leq c\}$$

for some non-negative c.

Our approach to proving existence is to show that under the metric d_M, $L(\alpha, \kappa)$ is lower semicontinuous on M and, for all $c \geq 0$, M_c is a compact subset of M. The existence of a loss-minimizing incentive compatible mechanism then follows immediately. Thus, the central underlying result of this paper is the following.

Theorem 8.3 *Endow M with the d_M-metric. Under Assumptions 4.1-4.9, M_c is a compact subset of M for all $c \geq 0$, and the principal's loss function $L : M \to [0, +\infty]$ defined in (3.1) is lower semicontinous. Thus, if M is nonempty, then the principal's problem (3.2) possesses a solution.*

[36] Thus, as when considering functions in L_p spaces in analysis, a "point" (α, κ) in the metric space (M, d_M) is the equivalence class of mechanisms (α', κ') such that $H \otimes \alpha' \otimes P \otimes \kappa' = H \otimes \alpha \otimes P \otimes \kappa$.

Before turning to the proof we provide sufficient conditions for the set of incentive compatible mechanisms, M, to be nonempty. These conditions cover a variety of common applications.

Proposition 8.4 *Suppose that Assumptions 4.1, 4.3, and 4.4 hold and that there is a measurable* $\phi : S \to R$ *such that* $\phi(s) \in \Phi_s$ *for every* s*, and such that,*

$$\int_S u(\phi(s), s, a, t) dP(s|a, t) \text{ is continuous in } a \text{ for each } t.$$

Then an incentive compatible mechanism exists, i.e., M *is nonempty.*

The proof of Proposition 8.4, which can be found in the Appendix, uses results on measurable selections of solutions to parameterized optimization problems.

Remark 8.5 *Given the other assumptions, the displayed continuity requirement in Proposition 8.4 is satisfied if either (a) some fixed reward is always feasible and the agent's utility given that reward does not depend on the signal, or (b) the funciton* ϕ *is continuous and, for each* t, $\int_{\{s:u(\phi(s),s,a,t)>n\}} u(\phi(s), s, a, t) dP(s|a, t) \to_n 0$ *uniformly in* $a \in A$.

One implication of Theorem 8.3 is that mechanisms that come close to giving the principal his minimum loss are close, in the d_M-metric, to an optimal mechanism. Hence, studying the optimal mechanism does not lose relevance if one is considering a principal who, for whatever reason, is only approximately optimizing.

9 A Sketch of the Proof

The heart of the proof of Theorem 8.3 is to show that $M_c = \{(\alpha, \kappa) \in M : L(\alpha, \kappa) \leq c\}$ is d_M-compact for every $c \in \mathbb{R}$ since this immediately implies that $L : M \to [0, \infty]$ is lower semicontinuous. Suppose then that (α_n, κ_n) is a sequence of mechanisms in M_c. We must show that there is a subsequence $(\alpha_{n_j}, \kappa_{n_j})$ and there is $(\alpha^*, \kappa^*) \in M_c$ such that $(\alpha_{n_j}, \kappa_{n_j})$ d_M-converges to (α^*, κ^*).

For any mechanism, let us define a participant's *equilibrium* payoff to be that obtained when the agent always reports truthfully and always takes the recommended action.

To convey the essence of the proof, we shall make several simplifying assumptions here. In particular, let us start by assuming that all spaces, R, S, A, and T, are compact, and that the payoff functions $u(r, s, a, t)$ and $l(r, s, a, t)$ are continuous. Then, finding a subsequence $(\alpha_{n_j}, \kappa_{n_j})$ and a limit mechanism $(\alpha^*, \kappa^*) \in M_c$ such that the principal's and the agent's equilibrium payoffs along the subsequence converge to their equilibrium payoffs at the limit is straightforward.

Indeed, given (α_n, κ_n), we may consider its distributional mechanism $H \otimes \alpha_n \otimes P \otimes \kappa_n \in \Delta(R \times S \times A \times T)$. Given our simplifying compactness assumption, $\Delta(R \times S \times A \times T)$ is compact in the weak* topology and so there is a subsequence $H \otimes \alpha_{n_j} \otimes P \otimes \kappa_{n_j}$ and a probability measure $\beta^* \in \Delta(R \times S \times A \times T)$ to which $H \otimes \alpha_{n_j} \otimes P \otimes \kappa_{n_j}$ weak* converges. Moreover, we show with the aid of standard results that β^* can be decomposed as $\beta^* = H \otimes \alpha^* \otimes P \otimes \kappa^*$ for some mechanism (α^*, κ^*). Notice then that $(\alpha_{n_j}, \kappa_{n_j})$ d_M-converges to (α^*, κ^*).

The principal's equilibrium payoff along the subsequence is

$$\int l(r, s, a, t) d\kappa_{n_j}(r|s, a, t) dP(s|a, t) d\alpha_{n_j}(a|t) dH(t) \leq c,$$

and, by the definition of weak* convergence (and our simplifying assumption that $l(\cdot)$ is continuous) this payoff sequence converges to

$$\int l(r, s, a, t) d\kappa^*(r|s, a, t) dP(s|a, t) d\alpha^*(a|t) dH(t) \leq c,$$

which is the principal's equilibrium payoff from the mechanism (α^*, κ^*). Similarly, the agent's equilibrium payoff converges. So, if (α^*, κ^*) is incentive compatible we would have $(\alpha^*, \kappa^*) \in M_c$ and we would be done.

But there is a difficulty. The difficulty is that weak*-convergence does not pin down (α^*, κ^*) at out-of-equilibrium reports and actions that are available to the agent and so it is entirely possible that (α^*, κ^*) is not incentive compatible. The "hard" part of the proof is to refine the convergence of the subsequence so that incentive compatibility is maintained. We next sketch how this is done.

The mechanism $(\alpha_{n_j}, \kappa_{n_j})$ is incentive compatible if (see fn. 15) there is a subset of types T^0 such that $H(T^0) = 1$ and for every $t, t' \in T^0$,

$$\int_{R \times S \times A} u(r, s, a, t) \, d\kappa_{n_j}(r|s, a, t) \, dP(s|a, t) \, d\alpha_{n_j}(a|t)$$

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

So, incentive compatibility will be maintained at the limit so long as for every $t, t' \in T^0$ the agent's equilibrium payoff from $(\alpha_{n_j}, \kappa_{n_j})$ when his type is t converges to his equilibrium payoff at the limit (α^*, κ^*), and his payoff from any false report of t' does not jump up at the limit.

Since the needed limit results must hold for all $t, t' \in T^0$, what we require is a pointwise convergence result in the agent's type. When there are finitely many or even countably many types, this poses no particular difficulty over and above the single-type case. However, when there are a continuum of possible types, as occurs in many applications, it may be impossible to find a subsequence along which convergence occurs pointwise in the agent's true and announced types t and t'. It is here where we use the insight of Page (1987,1991), who noted that one can make use of a powerful result due to Komlos (1967) and its generalization by Balder (1990), to obtain the necessary pointwise convergence. But some preparation is needed.

First, let's strengthen one of our informational assumptions, Assumption 4.7, and suppose that for all a, t, a', t', the measures $P_{a,t}$ and $P_{a',t'}$ are mutually absolutely continuous. Then, given the Radon-Nikodym derivative $g_{a,t/a',t'}$, we may write

$$dP\left(s|a,t\right) = g_{a,t/a',t'}(s)dP(s|a',t').$$

Defining $\nu_{n_j}(\cdot|a',t') = P(\cdot|a',t') \otimes \kappa_{n_j}(\cdot|\cdot,a',t') \in \Delta(R\times S)$, we may then write

$$d\kappa_{n_j}\left(r|s,a,t\right) dP\left(s|a,t\right) = g_{a,t/a',t'}(s)d\nu_{n_j}(r,s|a',t').$$

Consequently, and using that $g_{a,t/a,t}(s) = 1$ for $P_{a,t}$ a.e. s, the IC inequality can be written as,

$$\int\limits_{R\times S\times A} u\left(r,s,a,t\right) d\nu_{n_j}(r,s|a,t)d\alpha_{n_j}(a|t)$$

$$\geq \int_A \left(\sup_{a\in A} \int\limits_{R\times S} u\left(r,s,a,t\right) g_{a,t/a',t'}(s)d\nu_{n_j}(r,s|a',t') \right) d\alpha_{n_j}(a'|t').$$

The interpretation suggested by this notation is that after the agent reports t', the mechanism uses $\alpha_{n_j}(\cdot|t')$ to randomly recommend to the agent an action a', which then determines $\nu_{n_j}(\cdot|a',t')$, a joint distribution over $R\times S$. The agent's action choice now only affects what we may regard as his surrogate utility function, $u(r,s,a,t)g_{a,t/a',t'}(s)$.

So, given t', we can view the mechanism as jointly choosing $a' \in A$ and $\nu_{n_j}(\cdot|a',t') \in \Delta(R\times S)$ and revealing these to the agent before the agent chooses his action. With this point of view in mind, define $\mu_{n_j}(\cdot|t') = \alpha_{n_j}(\cdot|t') \otimes \delta_{\nu_{n_j}(\cdot|\cdot,t')}(\cdot)$, where $\delta_{\nu_{n_j}(\cdot|a',t')}(\cdot) \in \Delta(\Delta(R\times S))$ is the Dirac measure that puts probability one on $\nu_{n_j}(\cdot|a',t') \in \Delta(R\times S)$ for any a' chosen by $\alpha_{n_j}(\cdot|t')$. Then we can write the IC inequality as

$$\int\limits_{\Delta(R\times S)\times A}\left(\int\limits_{R\times S} u\left(r,s,a,t\right)d\nu(r,s)\right)d\mu_{n_j}(\nu,a|t)$$

$$\geq \int\limits_{\Delta(R\times S)\times A}\left(\sup_{a\in A}\int\limits_{R\times S} u\left(r,s,a,t\right)g_{a,t/a',t'}(s)d\nu(r,s)\right)d\mu_{n_j}(\nu,a'|t').$$

and we are finally in a position to apply the Komlos-Balder pointwise convergence result.

The Komlos-Balder result permits us to show (see Lemma A.4) that the subsequence n_j can be chosen so that for all $t\in T^0$ (more precisely, for all t in a set that differs from T^0 by an H-measure zero set of types) the Cesaro mean of the sequence of probability measures $\mu_{n_j}(\cdot|t)\in\Delta(\Delta(R\times S)\times A)$ weak* converges to some $\mu^*(\cdot|t)\in\Delta(\Delta(R\times S)\times A)$.[37] So, even though the subsequence itself need not converge pointwise in the agent's type, its Cesaro mean will, and as we are about to see, this turns out to suffice because the average of incentive compatible mechanisms is itself an incentive compatible mechanism.

The next step is to show that $\sup_{a\in A}\int_{R\times S}u(r,s,a,t)g_{a,t/a',t'}(s)d\nu$ is lower semicontinuous in (ν,a'). It is here that another of our informational assumptions, Assumption 4.8 plays its key role. In particular, note that under the stronger assumption that $g_{a,t/a',t'}(s)$ itself is lower semi-continuous in (s,a'), for each a,t, and t', the integral is lower-semicontinous in (ν,a'), and so the supremum over a is as well. We prove this for the general case in Lemma A.6. We may thus conclude that for every $t,t'\in T^0$,

$$\int\limits_{\Delta(R\times S)\times A}\left(\int\limits_{R\times S} u\left(r,s,a,t\right)d\nu(r,s)\right)d\mu^*(\nu,a|t)$$

$$\geq \int\limits_{\Delta(R\times S)\times A}\left(\sup_{a\in A}\int\limits_{R\times S} u\left(r,s,a,t\right)g_{a,t/a',t'}(s)d\nu(r,s)\right)d\mu^*(\nu,a'|t').$$

Finally, it is shown that, by applying to μ^* the reverse of the construction that created μ_{n_j}, one can obtain the limit mechanism (α^*,κ^*) that, when substituted into the above inequality, is seen to be incentive compatible.

[37] i.e., $\frac{1}{N}\sum_{j=1}^{N}\mu_{n_j}(\cdot|t)$ weak* converges to $\mu^*(\cdot|t)$ as $N\to\infty$, for H a.e. $t\in T^0$.

10 Proof of Theorem 8.3

If $M_c = \{(\alpha, \kappa) \in M : L(\alpha, \kappa) \leq c\}$ is d_M-compact, then it is d_M-closed. Consequently, since $L : M \to [0, \infty]$ is nonnegative, this would establish that L is lower semicontinuous. Hence, it suffices to show that M_c is d_M-compact for every $c \in \mathbb{R}$.

Fix any $c \in \mathbb{R}$, and consider any sequence $(\alpha_n, \kappa_n) \in M_c$. If $P_{a,t}(S_{a',t'}) < 1$ for some a, t, a', t', then without loss of generality, we may assume that, for every n, $\kappa_n(\{r_*\}|s, a, t) = 1$ for all $(s, a, t) \in S \times A \times T$ such that $s \notin S_{a,t}$ (a measurable subset of $S \times A \times T$ by Assumption 4.7) That is, we may assume without loss of generality that if the signal indicates that the agent either lied about his type or did not take the recommended action, then the agent is assigned a worst possible reward. Since (by Assumption 4.9) the new κ_n, when it differs from the old, makes the punishment for not complying as severe as possible, the new mechanism remains incentive compatible. Moreover, the new mechanism is d_M-equivalent to the old because the mechanism is unchanged on a set having probability one when the agent is truthful and takes the recommended action.

We will show that the sequence $(\alpha_n, \kappa_n) \in M_c$ has a subsequence that d_M-converges to a limit in M_c. We proceed in several steps.

Step 1: *Associate with each mechanism (α_n, κ_n) a transition probability $\mu_n : T \to \Delta(\Delta(R \times S) \times A)$.*

For each n, define $\nu_n : A \times T \to \Delta(R \times S)$ as follows. For every $(a, t) \in A \times T$ and for every $E \in \mathcal{B}(R \times S)$,

$$\nu_n(E|a, t) = \int_{R \times S} \mathbf{1}_E(r, s) d\kappa_n(r|s, a, t) dP(s|a, t).$$

Then $\nu_n : A \times T \to \Delta(R \times S)$ is a transition probability by Proposition 7.29 in Bertsekas and Shreve (1978), henceforth BS. Define $\gamma_n : A \times T \to \Delta(\Delta(R \times S))$ by $\gamma_n(\{\nu_n(\cdot|a, t)\}|a, t) = 1$ for every $(a, t) \in A \times T$. Then γ_n is measurable because it is the composition of the measurable (BS, Prop. 7.26) function $\nu_n : A \times T \to \Delta(R \times S)$ and the continuous function that maps any $\nu \in \Delta(R \times S)$ into the Dirac measure $\delta_\nu(\cdot) \in \Delta(\Delta(R \times S))$ that puts probability one on ν. Hence, being measurable, $\gamma_n : A \times T \to \Delta(\Delta(R \times S))$ is a transition probability (BS, Prop. 7.26).

For each $t \in T$, define $\mu_n : T \to \Delta(\Delta(R \times S) \times A)$ as follows. For every $t \in T$ and for every $E \in \mathcal{B}(\Delta(R \times S) \times A)$,

$$\mu_n(E|t) = \int_{\Delta(R \times S) \times A} \mathbf{1}_E(\nu, a) d\gamma_n(\nu|a, t) d\alpha_n(a|t).$$

Then $\mu_n : T \to \Delta(\Delta(R \times S) \times A)$ is a transition probability (BS, Prop. 7.29). In particular, $\mu_n(\cdot|t) \in \Delta(\Delta(R \times S) \times A)$ for every $t \in T$.

Step 2: *Establish the equivalence of expectations using $H \otimes \alpha_n \otimes P \otimes \kappa_n$ or using $H \otimes \mu_n$.*

For any measurable function $\zeta : R \times S \times A \times T \to \mathbb{R}$ and for any n, we can calculate the expectation of ζ by using either $H \otimes \alpha_n \otimes P \otimes \kappa_n$ or using $H \otimes \mu_n$ as follows, where $\nu_n : A \times T \to \Delta(R \times S)$ and $\gamma_n : A \times T \to \Delta(\Delta(R \times S))$ are as defined in Step 1.

$$\begin{aligned}
& \int \zeta(r,s,a,t) d\kappa_n(r|s,a,t) dP(s|a,t) d\alpha_n(a|t) dH(t) \\
= & \int_{A\times T} \left(\int_{R\times S} \zeta(r,s,a,t) dv_n(r,s|a,t) \right) d\alpha_n(a|t) dH(t) \\
= & \int_T \int_A \left(\int_{\Delta(R\times S)} \left(\int_{R\times S} \zeta(r,s,a,t) d\nu(r,s) \right) d\gamma_n(\nu|a,t) \right) d\alpha_n(a|t) dH(t) \\
= & \int_T \left(\int_A \int_{\Delta(R\times S)} \left(\int_{R\times S} \zeta(r,s,a,t) d\nu(r,s) \right) d\gamma_n(\nu|a,t) d\alpha_n(a|t) \right) dH(t) \\
= & \int_T \left(\int_{\Delta(R\times S)\times A} \left(\int_{R\times S} \zeta(r,s,a,t) d\nu(r,s) \right) d\mu_n(\nu,a|t) \right) dH(t). \qquad (10.1)
\end{aligned}$$

Step 3: *For each n, find a subsequence $\{n_j\}$ of $\{n\}$ such that $\left\{H \otimes \mu_{n_j}\right\}$ and the Cesaro mean of $\{\mu_{n_j}(\cdot|t)\}$ converge.*

We will show in particular that there exists a transition probability $\mu^* : T \to \Delta(\Delta(R \times S) \times A)$ and a subsequence $\{n_j\}$ of $\{n\}$ such that $\left\{H \otimes \mu_{n_j}\right\}$ converges to $H \otimes \mu^*$ and such that the Cesaro mean[38] of $\{\mu_{n_j}(\cdot|t)\}$ converges to $\mu^*(\cdot|t)$ for H almost every $t \in T$.

This step follows from Lemma A.4 in the Appendix, which relies heavily on Balder (1990). Without loss, we may assume that the original sequence $\{n\}$ has these properties. Then, in particular, letting $\bar{\mu}_m(\cdot|t)$ denote the m-th Cesaro mean of $\{\mu_n(\cdot|t)\}$, we have

$$\bar{\mu}_m(\cdot|t) = \frac{1}{m} \sum_{n=1}^{m} \mu_n(\cdot|t) \to_m \mu^*(\cdot|t),\ H \text{ a.e. } t \in T. \qquad (10.2)$$

Step 4: *Use μ^* to construct a candidate limit mechanism (α^*, κ^*).*

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

$$W_t = \{(\nu, a) \in \Delta(R \times S) \times A : \nu(\Phi) = 1 \text{ and } \mathrm{marg}_S \nu = P_{a,t}\}.^{39}$$

[38]The Cesaro mean of a sequence $\{x_n\}$ is the sequence whose m-th term (sometimes called the m-th Cesaro mean of $\{x_n\}$) is $\frac{1}{m}\sum_{n=1}^{m} x_n$.

[39]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

Construct α^*.

By BS Corollary 7.27.1, there is a transition probability $\alpha^* : T \to \Delta(A)$ and a transition probability $\eta^* : A \times T \to \Delta(\Delta(R \times S))$ 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 $\{\nu \in \Delta(R \times S) : \nu(\Phi) = 1 \text{ and } \text{marg}_S\nu = P_{a,t}\}$ for $\alpha^*(\cdot|t)$ a.e. $a \in A$.

Collapse a Lottery.

Note that η^* takes each element of $A \times T$ to a probability measure on $\Delta(R \times S)$, rather than to an element of $\Delta(R \times S)$. For each (a,t) define $\nu^*(\cdot|a,t) \in \Delta(R \times S)$ so that for every $E \in B(R \times S)$,

$$\nu^*(E|a,t) = \int_{\Delta(R\times S)} \nu(E) d\eta^*(\nu|a,t). \tag{10.3}$$

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 $E \in B(R \times S)$, $\nu(E)$ is a measurable real-valued function of ν on $\Delta(R \times S)$, BS Prop. 7.29 implies that $\nu^*(E|a,t)$ is a measurable real-valued function of (a,t) on $A \times T$. Hence, $\nu^* : A \times T \to \Delta(R \times S)$ is a transition probability.

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

$$\int_{\Delta(R\times S)} \left(\int_{R\times S} \zeta(r,s,a,t) d\nu(r,s) \right) d\eta^*(\nu|a,t) = \int_{R\times S} \zeta(r,s,a,t) d\nu^*(r,s|a,t). \tag{10.4}$$

Construct κ^*.

Again, by BS Corollary 7.27.1, there exists a transition probability $\tilde{P} : A\times T \to \Delta(S)$ and a transition probability $\kappa^* : S \times A \times T \to \Delta(R)$ such that for all a,t, $\nu^*(\cdot|a,t) = \tilde{P}(\cdot|a,t) \otimes \kappa^*(\cdot|\cdot, a, t)$. But because the marginal of $\nu^*(\cdot|a,t)$ on S is $P(\cdot|a,t)$ for $H \otimes \alpha^*$ a.e. (a,t) we must have $\tilde{P}(\cdot|a,t) = P(\cdot|a,t)$ for $H \otimes \alpha^*$ a.e. (a,t), and so $\nu^*(\cdot|a,t) = P(\cdot|a,t) \otimes \kappa^*(\cdot|\cdot, a, t)$ for $H \otimes \alpha^*$ a.e. (a,t). Also, since $\nu^*(\Phi|a,t) = 1$ holds for $H \otimes \alpha^*$ a.e. (a,t), $\kappa^*(\Phi_s|s,a,t) = 1$ for all (s,a,t) in a measurable subset Z of $S \times A \times T$ such that $[H \otimes \alpha^* \otimes P](Z) = 1$.[40] To satisfy the formal definition of a mechanism, modify $\kappa^*(\Phi_s|s,a,t)$ on $(S \times A \times T)\backslash Z$ so that $\kappa^*(\Phi_s|s,a,t) = 1$, e.g. by setting $\kappa^*(\cdot|s,a,t) = \kappa_1(\cdot|s,a,t)$ for all $(s,a,t) \in (S \times A \times T)\backslash Z$.[41] Finally, modify κ^* so that $\kappa^*(\{r_*(s,a,t)\}|s,a,t) = 1$ if $s \notin S_{a,t}$.[42]

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.

[40] $Z = \{(s,a,t) \in S \times A \times T : \kappa(\Phi_s|s,a,t) = 1\}$ is measurable since κ is a transition probability and Φ_s is the slice of a measurable (indeed closed) subset of $R \times S$.

[41] κ_1 is the second coordinate of the first term in our original sequence of mechanisms $(\alpha_1, \kappa_1), (\alpha_2, \kappa_2)$... .

[42] The modified κ^* is still a transition probability since $\{(s,a,t) : s \in S_{a,t}\}$ is measurable by Assumption 4.7, and r_* is a measurable function by 4.9.

The mechanism $(\alpha^*, \kappa^*) \in M$ is our candidate limit mechanism. By reasoning as in Step 2 combined with (10.4), for any measurable $\zeta : R \times S \times A \times T \to R$, we have,

$$\begin{aligned} & \int \zeta(r,s,a,t) d\kappa^*(r|s,a,t) dP(s|a,t) d\alpha^*(a|t) dH(t) \\ = & \int_T \left(\int_{\Delta(R\times S)\times A} \left(\int_{R\times S} \zeta(r,s,a,t) d\nu(r,s) \right) d\mu^*(\nu, a|t) \right) dH(t). \end{aligned} \tag{10.5}$$

Step 5: *Rewrite the utility to a deviating agent.*

If the sets $S_{a,t}$ in Assumption 4.7 are such that $P_{a,t}(S_{a',t'}) < 1$ for some a, t, a', t', then let $r_* \in R$ be as in Assumption 4.9 and define $u_*(s,a,t) = u(r_*, s, a, t)$ for all (s,a,t). Otherwise, define $u_*(s,a,t) = 0$ for all (s,a,t). Hence, in either case, $u(r,s,a,t) - u_*(s,a,t) \geq 0$ for all (r,s,a,t). For any $\nu \in \Delta(R \times S)$, $a, a' \in A$, and $t, t' \in T$, define

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

We will show that for each n, the utility to an agent of type t who reports type t' and then best responds conditional on whatever action a' is recommended, can be written as

$$\int_{\Delta(R\times S)\times A} \sup_{a\in A} U_*(\nu, a, t, a', t') d\mu_n(\nu, a'|t').^{43}$$

To see this, note that, in the mechanism (α_n, κ_n), the utility to an agent of type t who reports type t' and then optimally chooses an action after receiving a recommended action is as follows, where $\nu_n : A \times T \to \Delta(R \times S)$ and $\gamma_n : A \times T \to \Delta(\Delta(R \times S))$ are as defined in Step 1:

$$\begin{aligned} & \int_A \left(\sup_{a\in A} \int_{R\times S} u(r,s,a,t)\, d\kappa_n(r|s,a',t')\, dP(s|a,t) \right) d\alpha_n(a'|t') \\ = & \int_A \sup_{a\in A} \left(\int_{R\times S} u_*(s,a,t)\, d\kappa_n(r|s,a',t')\, dP(s|a,t) \right. \\ & \left. + \int_{R\times S} [u(r,s,a,t) - u_*(s,a,t)] d\kappa_n(r|s,a',t')\, dP(s|a,t) \right) d\alpha_n(a'|t') \\ = & \int_A \sup_{a\in A} \left(\int_S u_*(s,a,t)\, dP(s|a,t) \right. \\ & \left. + \int_{R\times S} [u(r,s,a,t) - u_*(s,a,t)] \mathbf{1}_{S_{a',t'}}(s) d\kappa_n(r|s,a',t')\, dP(s|a,t) \right) d\alpha_n(a'|t') \end{aligned}$$

[43] The inegral is well-defined because, as shown in the proof of Lemma A.6 in Appendix II, the integrand, $\sup_{a\in A} U_*(\nu, a, t, a', t')$, is lower semicontinuous, and hence measurable, as a function of $(\nu, a') \in \Delta(R\times S)\times A$.

$$
\begin{aligned}
&= \int_A \sup_{a\in A}\left(\int_S u_*\left(s,a,t\right)dP(s|a,t)\right.\\
&\quad\left.+\int_{R\times S}[u\left(r,s,a,t\right)-u_*\left(s,a,t\right)]g_{a,t/a',t'}(s)d\kappa_n\left(r|s,a',t'\right)dP\left(s|a',t'\right)\right)d\alpha_n(a'|t')\\
&= \int_{\Delta(R\times S)\times A}\sup_{a\in A}\left(\int_S u_*\left(s,a,t\right)dP(s|a,t)\right.\\
&\quad\left.+\int_{R\times S}[u\left(r,s,a,t\right)-u_*\left(s,a,t\right)]g_{a,t/a',t'}(s)d\nu\left(r,s\right)\right)d\mu_n(\nu,a'|t')\\
&= \int_{\Delta(R\times S)\times A}\sup_{a\in A}U_*(\nu,a,t,a',t')d\mu_n(\nu,a'|t'), \qquad (10.6)
\end{aligned}
$$

where the first equality follows by adding and subtracting $u_*\left(s,a,t\right)$; the second equality follows because (first term) $u_*\left(s,a,t\right)$ does not depend on r and (second term) either because $P_{a,t}(S_{a',t'})=1$, or, because $P_{a,t}(S_{a',t'})<1$ and $\kappa_n\left(\{r_*\}|s,a',t'\right)=1$ for every $s\in S_{a,t}\backslash S_{a',t'}$ and so (by the definition of $u_*(s,a,t)$) the quantity in square brackets is zero when $s\in S_{a,t}\backslash S_{a',t'}$; the third equality follows by Assumptions 4.7 and 4.8; the fourth equality follows because

$$
\begin{aligned}
&\left(\int_{R\times S}[u\left(r,s,a,t\right)-u_*\left(s,a,t\right)]g_{a,t/a',t'}(s)d\kappa_n\left(r|s,a',t'\right)dP\left(s|a',t'\right)\right)\\
=\ &\left(\int_{R\times S}[u\left(r,s,a,t\right)-u_*\left(s,a,t\right)]g_{a,t/a',t'}(s)d\nu_n\left(r,s|a',t'\right)\right)\\
=\ &\left(\int_{R\times S}[u\left(r,s,a,t\right)-u_*\left(s,a,t\right)]g_{a,t/a',t'}(s)d\nu\left(r,s\right)\right)d\gamma_n\left(\nu|a',t'\right),
\end{aligned}
$$

and because $\mu_n(\cdot|t)=\alpha_n(\cdot|t)\otimes\gamma_n(\cdot|\cdot,t)$; and where the fifth equality follows by the definition of U_*.

Step 6: *Show that (α^*,κ^*) is incentive compatible.*

Incentive compatibility of (α_n,κ_n) along with Step 5 imply that for H a.e. t, and every t',

$$
\begin{aligned}
\int_{\Delta(R\times S)\times A}\left(\int_{R\times S}u(r,s,a,t)d\nu(r,s)\right)d\mu_n(\nu,a|t) &= \int u\left(r,s,a,t\right)d\kappa_n\left(r|s,a,t\right)dP\left(s|a,t\right)d\alpha_n(a|t)\\
&\geq \int_{\Delta(R\times S)\times A}\sup_{a\in A}U_*(\nu,a,t,a',t')d\mu_n(\nu,a'|t'),
\end{aligned}
\qquad (10.7)
$$

where the equality follows from Step 2 with $\zeta=u$.

Show incentive compatibility in terms of μ^*.

Applying Lemmas A.5 and A.6 to the limits of the Cesaro means of the first and last terms in (10.7) implies that there is a measurable subset T^1 of T with $H(T^1) = 1$ such that for all $t, t' \in T^1$,

$$\int_{\Delta(R\times S)\times A} \left(\int_{R\times S} u(r,s,a,t) d\nu(r,s) \right) d\mu^*(\nu, a|t) \geq \int_{\Delta(R\times S)\times A} \sup_{a\in A} U_*(\nu, a, t, a', t') d\mu^*(\nu, a'|t'), \tag{10.8}$$

and where both integrals are finite.

Use incentive compatibility in terms of μ^* *to show incentive compatibility in terms of* (α^*, ν^*).

Recall that for every $a \in A$, $\eta^*(\cdot|a,t)$ is an element of $\Delta(\Delta(R \times S))$, i.e., $\eta^*(\cdot|a,t)$ is a *lottery* over elements of $\Delta(R\times S)$, and that we defined $\nu^*(\cdot|a,t) \in \Delta(R\times S)$ to collapse that lottery. In this step, we show that doing so does not affect incentive compatibility. Indeed, for all $t, t' \in T^1$,

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

where the first equality follows from (10.4), the first inequality follows by applying $\mu^*(\cdot|t) = \alpha^*(\cdot|t) \otimes \eta^*(\cdot|\cdot,t)$ to (10.8), the second inequality follows because the agent is no longer able to condition his action on the outcome of the lottery $\eta^*(\cdot|a',t')$ over $\nu \in \Delta(R\times S)$, and the final equality follows from (10.4).

Show incentive compatibility of (α^*, κ^*).

Since $\nu^*_{a,t} = P_{a,t} \otimes \kappa^*_{a,t}$ for $H \otimes \alpha^*$ a.e. (a,t), it follows that there is an H-measure 1 subset T^0 of T^1 such that for all $t, t' \in T^0$,

$$
\begin{aligned}
& \int_A \left(\int_{R\times S} u(r,s,a,t) d\kappa^*(r|s,a,t) dP(s|a,t) \right) d\alpha^*(a|t) \\
\geq \; & \int_A \sup_{a\in A} \left(\int_S u_*\left(s,a,t\right) dP(s|a,t) \right. \\
& \left. + \int_{R\times S} [u\left(r,s,a,t\right) - u_*\left(s,a,t\right)] g_{a,t/a',t'}\left(s\right) d\kappa^*\left(r|s,a',t'\right) dP(s|a',t') \right) d\alpha^*(a'|t') \\
= \; & \int_A \sup_{a\in A} \left(\int_S u_*\left(s,a,t\right) dP(s|a,t) \right. \\
& \left. + \int_{R\times S} [u\left(r,s,a,t\right) - u_*\left(s,a,t\right)] \mathbf{1}_{S_{a',t'}}(s) d\kappa^*\left(r|s,a',t'\right) dP(s|a,t) \right) d\alpha^*(a'|t') \\
= \; & \int_A \sup_{a\in A} \left(\int_S u_*\left(s,a,t\right) dP(s|a,t) \right. \\
& \left. + \int_{R\times S} [u\left(r,s,a,t\right) - u_*\left(s,a,t\right)] d\kappa^*\left(r|s,a',t'\right) dP(s|a,t) \right) d\alpha^*(a'|t') \\
= \; & \int_A \left(\sup_{a\in A} \int_{R\times S} u\left(r,s,a,t\right) d\kappa^*\left(r|s,a',t'\right) dP(s|a,t) \right) d\alpha^*(a'|t'),
\end{aligned}
$$

where the first equality follows by Assumptions 4.7 and 4.8, and the second equality follows either because $P_{a,t}(S_{a',t'}) = 1$ or because $P_{a,t}(S_{a',t'}) < 1$ and $\kappa_n\left(\{r_*\}|s,a',t'\right) = 1$ when $s \in S_{a,t} \backslash S_{a',t'}$ and so (by the definition of $u_*(s,a,t)$ in Step 5) the quantity in square brackets is zero when $s \in S_{a,t} \backslash S_{a',t'}$. 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 footnote 15). Hence, after the modification, $(\alpha^*, \kappa^*) \in M$.[44]

Step 7: *Show that* $(\alpha_n, \kappa_n) \to (\alpha^*, \kappa^*)$ *in the* d_M*-metric.*

We show that $H \otimes \alpha_n \otimes P \otimes \kappa_n$ converges to $H \otimes \alpha^* \otimes P \otimes \kappa^*$ in the weak* topology. Let $\zeta : R \times S \times A \times T \to [0,1]$ be continuous. Then,

$$
\begin{aligned}
& \int \zeta(r,s,a,t) d\kappa_n(r|s,a,t) dP(s|a,t) d\alpha_n(a|t) dH(t) \\
= \; & \int_T \int_{\Delta(R\times S)\times A} \left(\int_{R\times S} \zeta(r,s,a,t) d\nu(r,s) \right) d\mu_n(\nu,a|t) dH(t) \\
\to \; & \int_T \int_{\Delta(R\times S)\times A} \int_{R\times S} \zeta(r,s,a,t) d\nu(r,s) d\mu^*(\nu,a|t) dH(t) \\
= \; & \int \zeta(r,s,a,t) d\kappa^*(r|s,a,t) dP(s|a,t) d\alpha^*(a|t) dH(t),
\end{aligned}
$$

[44] Define $\bar{U}(\nu,t,a',t') = \sup_{a\in A} U_*(\nu,a,t,a',t')$. The last three equalities in the previous display show that $\sup_{a\in A} \int_{R\times S} u\left(r,s,a,t\right) d\kappa^*\left(r|s,a',t'\right) dP(s|a,t) = \bar{U}(P_{a',t'} \otimes \kappa^*_{a',t'}, t, a', t')$. By Lemma A.6, $\bar{U}(\nu,t,a',t')$ is u.s.c., and so measurable, in (ν, a'). Consequently, $\bar{U}(P_{a',t'} \otimes \kappa^*_{a',t'}, t, a', t')$ is measurable in a', being the composition of measurable functions. This verifies the claim stated in footnote 14.

where the first equality is from Step 2, the limit follows since $H \otimes \mu_n \to H \otimes \mu^*$ by Step 3, and the second equality follows from (10.5).

Step 8: *Show that* $L(\alpha^*, \kappa^*) \leq c$.

For any $(\nu, a, t) \in \Delta(R \times S) \times A \times T$, let $\mathcal{L}(\nu, a, t) = \int_{R\times S} l(r, s, a, t) d\nu(r, s)$. Since $(\alpha_n, \kappa_n) \in M_c$ for all n, we have $L(\alpha_n, \kappa_n) \leq c$, and so

$$\begin{aligned}
c &\geq \underline{\lim}_m \frac{1}{m} \sum_{n=1}^{m} L(\alpha_n, \kappa_n) \\
&= \underline{\lim}_m \frac{1}{m} \sum_{n=1}^{m} \int_T \left(\int_{\Delta(R\times S)\times A} \left(\int_{R\times S} l(r, s, a, t) d\nu(r, s) \right) d\mu_n(\nu, a|t) \right) dH(t) \\
&= \underline{\lim}_m \int_T \left(\int_{\Delta(R\times S)\times A} \mathcal{L}(\nu, a, t) d\bar{\mu}_m(\nu, a|t) \right) dH(t) \\
&\geq \int_T \left(\underline{\lim}_m \int_{\Delta(R\times S)\times A} \mathcal{L}(\nu, a, t) d\bar{\mu}_m(\nu, a|t) \right) dH(t) \\
&\geq \int_T \left(\int_{\Delta(R\times S)\times A} \mathcal{L}(\nu, a, t) d\mu^*(\nu, a|t) \right) dH(t) \\
&= L(\alpha^*, \kappa^*), \qquad (10.9)
\end{aligned}$$

where the first and last equalities follow from (10.1) and (10.5) respectively, and the second inequality follows from Fatou's lemma. It remains only to justify the third inequality.

By Lemma A.2, $\mathcal{L}(\nu, a, t)$ is nonnegative and lower semicontinuous on W_t. Therefore, since $\mu^*(W_t|t) = 1$ for H a.e. $t \in T$ (see Step 4), (10.2) and Lemma A.1 yield the third inequality.

Steps 6-8 together imply that $(\alpha_n, \kappa_n) \to (\alpha^*, \kappa^*) \in M_c$, completing the proof that M_c is compact. Q.E.D.

11 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 such as lower and upper bounds on payments or other restrictions that a court or law might place on what can be enforced within a contract. But the legal system or economic forces might equally well constrain the mechanism in its totality. One simple case is a participation constraint asking that the agent be given a minimal utility at the ex-ante stage. Another 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 such restrictions on the set of mechanisms. At the end of this section, we illustrate how the results here yield an existence result when the model includes some simple dynamics involving outside options.

Let $M' \subseteq M$ be any subset of the set of incentive compatible mechanisms. The principal's M'-restricted problem is

$$\min_{(\alpha,\kappa)\in M'} L(\alpha,\kappa).$$

Let

$$M_c = \{(\alpha,\kappa) \in M : L(\alpha,\kappa) \leq c\},$$

be the set of incentive compatible mechanisms that yield the principal expected losses no greater than c.

Corollary 11.1 *Suppose that Assumptions 4.1-4.9 hold and that 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 Theorem 8.3 because a closed subset of a compact set is compact.

Below we illustrate the value of Corollary 11.1 with a number of useful applications. We assume throughout that Assumptions 4.1-4.9 hold. Before getting to the applications, we first identify a collection of subsets M' of M that are closed.

Lemma 11.2 *Suppose that $g : R \times S \times A \times T \to \mathbb{R}$ and $\xi : T \to \mathbb{R}$ are measurable functions such that: $g(r,s,a,t) \geq \xi(t)$ for every (r,s,a,t), $g(r,s,a,t)$ is lower semicontinuous in (r,s,a) for each t, and $\int |\xi(t)|\, dH(t) < \infty$. Then, for any $b \in \mathbb{R}$, the set*

$$M' = \{(\alpha,\kappa) \in M : \int g(r,s,a,t)\, d\kappa(r|s,a,t)\, dP(s|a,t)\, d\alpha(a|t)\, dH(t) \leq b\}$$

is d_M-closed.

Proof. Suppose that (α_n, κ_n) in M' d_M-converges to $(\hat{\alpha}, \hat{\kappa}) \in M$, i.e., that $H \otimes \alpha_n \otimes P \otimes \kappa_n$ weak* converges to $H \otimes \hat{\alpha} \otimes P \otimes \kappa^*$. Then, by Lemma A.7 part (ii) (with $X = R \times S \times A$), $b \geq \underline{\lim}_n \int gd(H \otimes \alpha_n \otimes P \otimes \kappa_n) \geq \int gd(H \otimes \hat{\alpha} \otimes P \otimes \hat{\kappa})$. Hence, $(\hat{\alpha}, \hat{\kappa}) \in M'$ and so M' is d_M-closed. Q.E.D.

The first application of Corollary 11.1 considers a situation in which a regulator places an upper bound, π_0, on the principal's expected profits conditional on a subset of types.

Example 11.3 *Suppose that the principal's measurable profit function is $\pi : R \times S \times A \times T \to \mathbb{R}$, which might differ from $-l$, the negative of the principal's disutility (there may be non-monetary aspects of the principal's loss function that the regulator does not care about).*

Suppose that there is a measurable $\xi : T \to \mathbb{R}$ such that $\pi(r,s,a,t) \geq \xi(t)$ for all $t \in T$ and $\int |\xi(t)|\, dH(t) < \infty$. Suppose also that $\pi(r,s,a,t)$ is lower semicontinuous in (r,s,a) for each t. Then, for any measurable subset $\hat{T}$ of T,

$$M' = \{(\alpha,\kappa) \in M : \int \mathbf{1}_{\hat{T}}(t)\pi(r,s,a,t)\, d\kappa(r|s,a,t)\, dP(s|a,t)\, d\alpha(a|t)\, dH(t) \leq \pi_0\},$$

is d_M-closed by Lemma 11.2 (with $g = 1_{\hat{T}}\pi$).

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. For example, 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.

Example 11.4 *Suppose that $\zeta : A \times T \to \mathbb{R}$ is measurable and that for each $t \in T$, $\zeta(a,t)$ is upper semicontinuous in a. Suppose also that there is a nonnegative measurable $\xi : T \to [0,\infty)$ such that $\zeta(a,t) \leq \xi(t)$ for all $t \in T$ and $\int \xi(t) dH(t) < \infty$. Let $M' = \{(\alpha,\kappa) \in M : \int \zeta(a,t)\, d\alpha(a|t)\, dH(t) \geq b\}$. Then, by Lemma 11.2 (with $g = -\zeta$), M' is d_M-closed.*

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

Example 11.5 *Let S^o be an open subset of S. Let $M' = \{(\alpha,\kappa) \in M : \int_{A\times T} P(S^o|a,t)\, d\alpha(a|t)\, dH(t) \leq b\}$. Then, because $P(S^o|\cdot)$ is bounded and $P(S^o|a,t)$ is l.s.c. in a for each t, Lemma 11.2 (with $g(\cdot) = P(S^o|\cdot)$), implies that M' is d_M-closed.*

The next examples illustrate how, with appropriate restrictions on the space of mechanisms, 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. Our model up to now includes only one such decision, i.e., the principal's choice of the contract.

In the first set of such examples, the agent's reward, $r = (r_1, r_2)$, has two coordinates.[45] It is assumed that the first coordinate, r_1, can as usual be chosen by the principal after the signal is observed, but that the second coordinate, r_2, must be chosen by the principal at the interim stage, immediately after the agent reports his type. 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.

[45] Each coordinate can be multidimensional.

Example 11.6 *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* $\alpha_1 : T \to \Delta(R_2 \times A)$ *and* $\kappa_1 : R_2 \times S \times A \times T \to \Delta(R_1)$ *are transition probabilities*$\}$. *The proof that* M' *is closed is in the Appendix. 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.*

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.[46]

Example 11.7 *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*

$$\int_{R \times S \times A} u(r, s, a, t)\, d\kappa(r|s, a, t)\, dP(s|a, t) d\alpha(a|t) \geq u_O(t), \quad H \text{ a.e. } t \in T, \tag{11.1}$$

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. In the Appendix we show that* $M' \cap M'' \cap M_c$ *is* d_M*-closed for every* $c \in \mathbb{R}$, *implying that Corollary 11.1 applies.*

[46] We have already discussed in Section 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 to the agent in addition to the availability of an interim outside option.

Our final example adds to the principal's problem an additional constraint requiring that the agent receives some minimum ex-ante expected utility. A similar example would add a constraint requiring that the principal's ex-ante expected loss is below some give bound.

Example 11.8 *Suppose that there is a measurable $\xi : T \to [0, \infty)$ such that $u(r, s, a, t) \leq \xi(t)$ for all $t \in T$ and $\int \xi(t)dH(t) < \infty$. Prior to learning his type, the agent can take an outside option and receive utility u_O or he can choose to participate in the mechanism. Thus, the principal must choose a mechanism from*

$$M' = \{(\alpha, \kappa) \in M : \int u(r, s, a, t)\, d\kappa(r|s, a, t)\, dP(s|a, t)\, d\alpha(a|t)\, dH(t) \geq u_O\}$$

in order to get the agent to participate. In the Appendix we show that $M' \cap M_c$ is d_M-closed for every $c \in \mathbb{R}$, implying that Corollary 11.1 applies.

Taken together, Examples 11.6, 11.7, and 11.8 show that our Assumptions 4.1-4.9 yield the existence of an optimal mechanism for the principal in the following simple dynamic setting. At date 1, both the principal and the agent can choose to quit or continue. If either quits, payoffs are realized and the game ends. If they both continue, then at date 2, the agent learns his type and can decide whether to quit or continue. If he quits payoffs are realized and the game ends. If he continues, then at stage 3, the agent reports his type to the principal and the principal can decide whether to quit or to continue. If she quits, payoffs are realized and the game ends. If she continues, then at date 4, the principal makes all decisions that are under her control, one of which is the choice of the agent's contract. At date 5, the principal recommends an action for the agent and the agent chooses any action from those that are available. Payoffs are realized and the game ends.

12 Random vs. Deterministic Mechanisms

An important feature of our setup is that it allows for randomized mechanisms. However, depending on the setting, the optimal mechanism may or may not require randomization over rewards. Moreover, in various economic situations randomization may be precluded despite being beneficial to the principal. In this section we study these two separate issues. First, we provide sufficient conditions under which full optimality can be achieved without randomization over rewards. Second, we consider restricted settings in which only deterministic mechanisms are permitted.

12.1 Sufficient Conditions for the Optimality of Deterministic Contracts

In this section we show that risk aversion and separability imply that mechanisms with deterministic contracts are fully optimal. This intuitive result follows because any nontrivial randomization over the agent's reward is strictly worse for the principal than the agent's certainty equivalent. It is noteworthy that the conditions required for this result are quite strong. 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 $D = \{\delta_r | r \in R\} \subseteq \Delta(R)$ be the set of Dirac measures on R. Say that a mechanism (α, κ) has *deterministic rewards* if $\kappa(\cdot|s,a,t) \in D$ for $H \otimes \alpha \otimes P$ a.e. (s,a,t) in $S \times A \times T$.[47]

Proposition 12.1 *Let $e : \Delta(R) \times S \to R$ be measurable. Suppose that for all (s,a,t) and all $\rho \in \Delta(\Phi_s)$,*

1. $e(\rho, s) \in \Phi_s$
2. $\int u(r,s,a,t)\, d\rho = u(e(\rho,s), s, a, t)$, *and*
3. $\int l(r,s,a,t)\, d\rho \geq l(e(\rho,s), s, a, t)$.

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

The proof is simply to start from an optimal mechanism (α, κ), 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 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 12.1, generalizes Holmström's result to a setting that allows for adverse selection.

Example 12.2 *Let $R = [0, \infty)$ and 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 for each s, $v(r,s)$ is continuous, concave, and strictly increasing in r, and $l(r,s,a,t) = \varphi(r,s)\sigma(s,a,t) + \xi(s,a,t)$ where for each s, $\varphi(r,s)$ is convex and strictly increasing in r, and where the functions τ and σ are positive.*

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

For each s and for each $\rho \in \Delta(\Phi_s)$, because $v(\cdot, s)$ is continuous, we may define $e(\rho, s)$ so that $v(e(\rho, s), s) = \int v(r, s)\, d\rho(r)$, i.e., so that $e(\rho, s)$ is the agent's certainty equivalent to ρ. Then, condition 1 is satisfied since $v(\cdot, s)$ is strictly increasing and since ρ puts probability 1 on the interval Φ_s. Condition 2 is satisfied by construction. Condition 3 is satisfied since $e(\rho, s) \leq \int r d\rho$ by Jensen's inequality, and the inequality in 3 is strict if ρ is non-degenerate and, either $v(r, s)$ is strictly concave in r for all s (so that $e(\rho, s) < \int r d\rho$) or $\varphi(r, s)$ is strictly convex in r for all s.

12.2 Restriction to Deterministic Mechanisms

As discussed in Section 8, some applications may preclude randomization over recommended actions or rewards. In this section we provide conditions under which our main result extends to such setups. The key to such results is to consider a subset of mechanisms M' in which randomization is precluded. If M' is closed under the d_M–metric, then Corollary 11.1 applies.

Consider first the case in which the set of types is at most countable, where it is without loss of generality that all types have positive probability. If (α_n, κ_n) is any sequence of mechanisms such that $\alpha_n(\cdot|t)$ puts probability one on some action for every t, and (α_n, κ_n) d_M-converges to (α, κ), then, in particular, $H \otimes \alpha_n$ weak* converges to $H \otimes \alpha$. But this means that $\alpha_n(\cdot|t)$ weak* converges to $\alpha(\cdot|t)$ for every t, and so $\alpha(\cdot|t)$ also puts probability one on some action. Hence, the subset M' of mechanisms that do not allow randomization over contracts is d_M-closed and so we have the following result.

Proposition 12.3 *Suppose that Assumptions 4.1-4.9 hold and that there are at most countably many agent types. Then, the principal's problem, when restricted to mechanisms that do not randomize over recommended actions, possesses a solution provided that at least one such incentive-compatible mechanism exists.*

By essentially the same reasoning, a similar result holds when the principal is, in addition, restricted to deterministic contracts (i.e., contracts that do not randomize over rewards) if we add the restriction that the signal space is at most countable and has the discrete topology (the agent's action set can still be a continuum). Call a mechanism *deterministic* if, conditional on the reported type, it does not randomize over recommended actions, and, conditional on the reported type, recommended action, and observed signal, it does not randomize over rewards. Corollary 11.1 yields the following result.

Proposition 12.4 *Suppose that Assumptions 4.1-4.9 hold, that there are at most countably many types and countably many signals, and that the signal space has the discrete topology. Then, the principal's problem, when restricted to deterministic mechanisms, possesses a solution provided that at least one such incentive-compatible mechanism exists.*

Corollary 11.1 also allows us to provide an existence result for deterministic mechanisms when the countability assumptions in the previous propositions fail. This result makes use of the stronger informational assumptions considered in Section 6.[48] In particular, we assume here that $P_{a,t}(E) = \int_E f(s|a,t)dQ(s)$ for every measurable $E \subseteq S$.

Let G_1 be a set of measurable functions from T to A and let G_2 be a set of measurable functions from $S \times T$ to R. Any mapping in G_1 specifies, for any type of the agent, the action from A that he is expected to take. Any mapping from G_2 specifies, for any type of the agent, the contract (i.e., the function from signals to rewards) that governs his compensation. In applications, the sets G_1 and G_2 should ideally arise from economically meaningful restrictions on the space of mechanisms and contracts that are permitted. The set of feasible deterministic mechanisms here is then $G = G_1 \times G_2$, whose typical element is a function from $S \times T$ into $A \times R$. We endow this set with the topology of $H \times Q$ a.e. pointwise convergence. Corollary 11.1 yields the following result.

Proposition 12.5 *Suppose that the conditions in Corollary 6.1 are satisfied and that G is sequentially compact. If G contains at least one incentive-compatible mechanism, then the principal's problem restricted to mechanisms in G possesses a solution.*

13 Multiple Agents

In this section, we show how Theorem 8.3 leads to an existence result for a model in which a single principal interacts with multiple agents. The notation will be as in the single agent case except for the presence of a subscript $i \in \{1, ..., I\}$ for each of the I agents. So, T_i and A_i are agent i's type and action spaces, and $T = \times_i T_i$, $A = \times_i A_i$. We shall assume that the agents' types are drawn from T according to $H \in \Delta(T)$. Agent i's utility function is $u_i : R \times S \times A \times T \to \mathbb{R}$. A mechanism here is (α, κ) specifying transition probabilities $\kappa : S \times A \times T \to \Delta(R)$ and $\alpha : T \to \Delta(A)$ and where $\kappa(\Phi_s|s, a, t) = 1$ for all $(s, a, t) \in S \times A \times T$.

The mechanism (α, κ) works as follows. Nature draws a vector of types t according H. Each agent i simultaneously learns his type t_i, and then privately reports his type to the mechanism. For any vector of reported types $t' \in T$, the mechanism chooses a vector of actions $a' \in A$ according to $\alpha(\cdot|t')$. Each agent i is then privately informed of his recommended action a'_i. Then all agents simultaneously choose an action, with agent i choosing an action from A_i. For any vector of chosen actions $a \in A$, a signal $s \in S$ is drawn according to

[48] For this result it is possible to relax the informational assumptions in Section 6 to some extent by allowing Q to depend on type.

$P_{a,t} \in \Delta(S)$. Finally, a reward $r \in R$ is drawn according $\kappa(\cdot|s, a', t')$. Agent i receives utility $u_i(r, s, a, t)$ and the principal receives loss $l(r, s, a, t)$.

13.1 Absolute Continuity: Types and Signals

For every agent i, let H_i denote the marginal of H on T_i and let $\bar{H} = \times_i H_i$ denote the product of the marginals. We shall assume from now on that H is absolutely continuous with respect to the product of its marginals $\bar{H}$.[49] Consequently, by the Radon-Nikodym theorem, there is a measurable $h : T \to [0, \infty)$ such that for every measurable $C \subseteq T$,

$$H(C) = \int_C h(t) dH_1(t_1)...dH_I(t_I) = \int_C h(t) d\bar{H}(t). \tag{13.1}$$

We shall also assume from now on that there is $Q \in \Delta(S)$ such that for every (a, t), $P_{a,t}$ is absolutely continuous with respect to Q. Specifically, we assume that there is a measurable $f : S \times A \times T \to [0, \infty)$ such that for every $(a, t) \in A \times T$ and for every measurable subset E of S,

$$P(E|a, t) = \int_E f(s|a, t) dQ(s). \tag{13.2}$$

This is the same assumption that was made in the single-agent model to arrive at the "standard case" there (see Section 6).

13.2 Incentive Compatibility

For any mechanism (α, κ), and for any agent i, by BS Proposition 7.27 we may decompose $\bar{H} \otimes \alpha \in \Delta(A \times T)$ as $\bar{H} \otimes \alpha = H_i \otimes \alpha_i \otimes \beta_i$ for some transition probabilities $\alpha_i : T_i \to \Delta(A_i)$ and $\beta_i : A_i \times T_i \to A_{-i} \times T_{-i}$. Moreover, the transition probability α_i is unique up to an H_i measure zero set of types, and the transition probability β_i is unique up to subset of $A_i \times T_i$ that has measure zero according to the marginal of $\bar{H} \otimes \alpha$ on $A_i \times T_i$, which, it should be noted, is equal to $H_i \otimes \alpha_i$. Because α_i is unique only H_i almost everywhere, it is most natural to define incentive compatibility so that only almost all (rather than all) untruthful reports are non-improving.[50]

The mechanism (α, κ) is *incentive compatible* if for every agent i, there is a measurable

[49] This important and useful condition was first introduced by Milgrom and Weber (1985).

[50] Analagous to footnote 15, one can ensure that, for each agent i, the IC inequality (13.3) below holds for all reports $t'_i \in T_i$ by first choosing any $t^0_i \in T^0_i$ and treating any report $t'_i \in T_i \backslash T^0_i$ made by agent i as if he had reported t^0_i. Adjusting the mechanism in this way for every agent leaves the mechanism unchanged on the probablity-one set $\times_{i=1}^n T^0_i$, and makes any report $t'_i \in T_i$ by any agent i payoff-equivalent (for all agents and for the principal) to some report in T^0_i.

$T_i^0 \subseteq T_i$ such that $H_i(T_i^0) = 1$, and, for all $t_i, t_i' \in T_i^0$,

$$\int_{R\times S\times A\times T_{-i}} u_i(r,s,a,t_i,t_{-i})h(t_i,t_{-i})f(s|a,t_i,t_{-i})d\kappa(r|s,a,t_i,t_{-i})dQ(s)d\beta_i(a_{-i},t_{-i}|a_i,t_i)d\alpha_i(da_i|t_i)$$

$$\geq \int_{A_i} \sup_{a_i\in A_i} \left\{ \int_{\substack{R\times S \\ \times A_{-i}\times T_{-i}}} u_i(r,s,a_i,a_{-i},t_i,t_{-i})h(t_i,t_{-i})f(s|a_i,a_{-i},t_i,t_{-i}) \right.$$

$$\left. d\kappa(r|s,a_i',a_{-i},t_i',t_{-i})dQ(s)d\beta_i(a_{-i},t_{-i}|a_i',t_i') \right\} d\alpha_i(a_i'|t_i'),^{51} \qquad (13.3)$$

Remark 13.1 *To obtain type t_i's conditional expected utility, divide both sides of (13.3) by $\int h(t_i,t_{-i})d\left(\times_{j\neq i}H_j(t_j)\right)$ when this quantity is positive.*

Remark 13.2 *Condition (13.1) has the important implication that certain conditional distributions depend on i's true type t_i only through the Radon-Nikodym derivative $h(t_i,t_{-i})$. These conditional distributions are, first, the conditional distribution over the recommended action a_i' given that i's type is t_i and that he reported t_i', and second, the conditional distribution over the others' actions and types (a_{-i},t_{-i}) given that i's type is t_i and that he was asked to take action a_i' after reporting type t_i'. This is why $\alpha_i(\cdot|t_i')$ and $\beta_i(\cdot|\cdot,t_i')$ in (13.3) depend only on the reported type t_i' and not also on the true type t_i.*

Remark 13.3 *Changing the transition probabilities α_i or β_i on any measure zero sets where they are not uniquely defined does not affect whether (α,κ) satisfies (13.3). Therefore, the incentive compatibility of any (α,κ) does not depend on which particular versions of the transition probabilities α_i and β_i are chosen.*

13.3 The Induced Single-Agent-i Model

In this section, we show how the incentive constraint (13.3) can be written in a useful form that resembles the single-agent incentive constraint from Definition 3.2.

Given the multi-agent model above, let us define, for any agent i, the *single-agent-i* model to be a principal-agent model with a single agent as in Section 3, but where the reward space is $R_i = R \times A_{-i} \times T_{-i}$, the signal space is S, the signal is always drawn according to Q, the action space is A_i, the type space is T_i with prior H_i and, for any

[51] As discussed previously in footnote 14, the outer integral on the right-hand side is with respect to the completion of the measure $\alpha_i(\cdot|t_i')$, although, at the optimum, this technical caveat is not required.

$(r_i, s, a_i, t_i) = ((r, a_{-i}, t_{-i}), s, a_i, t_i) \in R_i \times S \times A_i \times T_i$ the agent's utility is $\tilde{u}_i(r_i, s, a_i, t_i) = u_i(r, s, a, t)h(t)f(s|a, t)$ and the principal's loss is $\tilde{l}(r_i, s, a_i, t_i) = l(r, s, a, t)h(t)f(s|a, t)$.

Let (α, κ) be any mechanism for the multi-agent model, and choose transition probabilities $\alpha_i : T_i \to \Delta(A_i)$ and $\beta_i : A_i \times T_i \to A_{-i} \times T_{-i}$ so that $\bar{H} \otimes \alpha = H_i \otimes \alpha_i \otimes \beta_i$.

As previously noted, the marginal of $\bar{H} \otimes \alpha$ on $A_i \times T_i$ is $H_i \otimes \alpha_i$. Consequently, because $Q \in \Delta(S)$ is constant as a transition probability into $\Delta(S)$, the marginal of $\bar{H} \otimes \alpha \otimes Q$ on $S \times A_i \times T_i$ is $H_i \otimes \alpha_i \otimes Q$. But then $H_i \otimes \alpha_i \otimes Q$ is also the marginal of $\bar{H} \otimes \alpha \otimes Q \otimes \kappa$ on $S \times A_i \times T_i$. Hence, by BS Proposition 7.27, there is a transition probability $\kappa_i : S \times A_i \times T_i \to \Delta(R \times A_{-i} \times T_{-i})$, such that $\bar{H} \otimes \alpha \otimes Q \otimes \kappa = H_i \otimes \alpha_i \otimes Q \otimes \kappa_i \in \Delta(R \times S \times T \times A)$. Since $\bar{H} \otimes \alpha = H_i \otimes \alpha_i \otimes \beta_i$ we have, for every agent i,

$$H_i \otimes \alpha_i \otimes \beta_i \otimes Q \otimes \kappa = H_i \otimes \alpha_i \otimes Q \otimes \kappa_i \in \Delta(R \times S \times T \times A). \tag{13.4}$$

In particular, for every agent i and for H_i-a.e. $t_i \in T_i$,

$$\alpha_i(\cdot|t_i) \otimes \beta_i(\cdot|\cdot, t_i) \otimes Q \otimes \kappa(\cdot|\cdot, t_i) = \alpha_i(\cdot|t_i) \otimes Q \otimes \kappa_i(\cdot|\cdot, t_i), \tag{13.5}$$

and, for H_i-a.e. $t'_i \in T_i$ and for $\alpha_i(\cdot|t'_i)$ a.e. $a'_i \in A_i$.

$$\beta_i(\cdot|a'_i, t'_i) \otimes Q \otimes \kappa(\cdot|\cdot, a'_i, t'_i) = Q \otimes \kappa_i(\cdot|\cdot, a'_i, t'_i). \tag{13.6}$$

By (13.5), for every agent i, for H_i-a.e. $t_i \in T_i$,

$$\int_{R \times S \times A \times T_{-i}} u_i(r, s, a, t)h(t)f(s|a, t)d\kappa(r|s, a, t)dQ(s)d\beta_i(a_{-i}, t_{-i}|a_i, t_i)d\alpha_i(a_i|t_i)$$

$$= \int_{R_i \times S} \tilde{u}_i(r_i, s, a_i, t_i)d\kappa_i(r_i|s, a_i, t_i)dQ(s)d\alpha_i(a_i|t_i). \tag{13.7}$$

By (13.6), for every agent i, for every $t_i \in T_i$ and for H_i-a.e. $t'_i \in T_i$,

$$\int_{A_i} \left(\sup_{a_i \in A_i} \int_{\substack{R \times S \\ \times A_{-i} \times T_{-i}}} u_i(r, s, a_i, a_{-i}, t)h(t)f(s|a, t)d\kappa(r|s, a'_i, a_{-i}, t'_i, t_{-i})dQ(s)d\beta_i(a_{-i}, t_{-i}|a'_i, t'_i) \right) d\alpha_i(a'_i|t'_i)$$

$$= \int_{A_i} \left(\sup_{a_i \in A_i} \int_{R_i \times S} \tilde{u}_i(r_i, s, a_i, t_i)d\kappa_i(r_i|s, a'_i, t'_i)dQ(s) \right) d\alpha_i(a'_i|t'_i). \tag{13.8}$$

In view of (13.7) and (13.8), (α, κ) is incentive compatible iff for every i, for H_i-a.e.

$t_i \in T_i$, and for H_i-a.e. $t'_i \in T_i$,

$$\int_{R_i \times S} \tilde{u}_i(r_i, s, a_i, t_i) d\kappa_i(r_i|s, a'_i, t'_i) dQ(s) d\alpha_i(a_i|t_i)$$

$$\geq \int_{A_i} \left(\sup_{a_i \in A_i} \int_{R_i \times S} \tilde{u}_i(r_i, s, a_i, t_i) d\kappa_i(r_i|s, a'_i, t'_i) dQ(s) \right) d\alpha_i(a'_i|t'_i), \qquad (13.9)$$

where $\alpha_i : T_i \to \Delta(A_i)$ and $\kappa_i : S \times A_i \times T_i \to \Delta(A_{-i} \times T_{-i})$ are transition probabilities such that $\bar{H} \otimes \alpha \otimes Q \otimes \kappa = H_i \otimes \alpha_i \otimes Q \otimes \kappa_i$

We have therefore established the following.

Proposition 13.4 *A mechanism* (α, κ) *is incentive compatible for the multi-agent model iff for every agent* i*, there are transition probabilities* $\alpha_i : T_i \to \Delta(A_i)$ *and* $\kappa_i : S \times A_i \times T_i \to \Delta(A_{-i} \times T_{-i})$ *such that* $\bar{H} \otimes \alpha \otimes Q \otimes \kappa = H_i \otimes \alpha_i \otimes Q \otimes \kappa_i$ *and* (α_i, κ_i) *is incentive compatible for the single-agent-*i *model.*

The proof of the following result is in the appendix and uses Proposition 13.4 and Theorem 8.3.

Theorem 13.5 *Suppose that for each agent* $i \in \{1, ..., I\}$*, Assumptions 4.1-4.6 hold for the single-agent-*i *model. Then, provided that at least one multi-agent incentive-compatible mechanism exists, a mechanism that minimizes the principal's expected losses exists for the multi-agent model.*[52]

Conditions for the existence of at least one multi-agent incentive-compatible mechanism are provided by the following.

Proposition 13.6 *Suppose that Assumptions 4.1, 4.3, and 4.4 hold and that there are measurable functions* $\phi : S \to R$ *and* $\xi : T \to [0, \infty)$*, such that* $\phi(s) \in \Phi_s$ *for every* s*,* $\int \xi(t) dH(t) < \infty$*, and, letting* $v_i(a, t) = \int_S u_i(\phi(s), s, a, t) dP(s|a, t)$ *for each agent* $i \in \{1, ..., I\}$*,* $|v_i(a, t)| \leq \xi(t)$ *for every* a *and* t*, and* $v_i(a, t)$ *is continuous in* a *for each* t*. Then at least one multi-agent incentive-compatible mechanism exists.*

A Appendix

Lemma A.1 *Let* X *and* Y *be Polish spaces, let* $\{y_n\} \subseteq Y$ *converge to* $\hat{y}$ *and let* $\{\gamma_n\} \subseteq \Delta(X)$ *weak* converge to* $\hat{\gamma} \in \Delta(X)$*. If* $\zeta : X \times Y \to \mathbb{R}$ *is bounded below and is lower semicontinuous at* $(x, \hat{y})$ *for* $\hat{\gamma}$ *a.e.* $x \in X$*, then*

$$\underline{\lim}_n \int_X \zeta(x, y_n) d\gamma_n(x) \geq \int_X \zeta(x, \hat{y}) d\hat{\gamma}(x).$$

[52] Conditions on u_i, l, f, and h under which $\tilde{u}_i$ and $\tilde{l}$ satisfy the assumptions of this theorem are analogous to those presented in Remark 6.2.

Proof. Let $\underline{\zeta}(x,y) = \liminf_{(x',y')\to(x,y)} \zeta(x',y')$ for every $(x,y) \in X \times Y$. Then $\underline{\zeta}$ is l.s.c., $\underline{\zeta} \leq \zeta$, and $\underline{\zeta}(x,\hat{y}) = \zeta(x,\hat{y})$ for $\hat{\gamma}$ a.e. $x \in X$. Hence,

$$\begin{aligned} \underline{\lim}_n \int_X \zeta(x,y_n)d\gamma_n(x) &\geq \underline{\lim}_n \int_X \underline{\zeta}(x,y_n)d\gamma_n(x) \\ &\geq \int_X \underline{\zeta}(x,\hat{y})d\hat{\gamma}(x) \\ &= \int_X \zeta(x,\hat{y})d\hat{\gamma}(x), \end{aligned}$$

where the first inequality follows because $\underline{\zeta} \leq \zeta$, the second follows by BS Proposition 7.31 because $\underline{\zeta}$ is bounded below and lower semicontinuous, and the equality follows because $\underline{\zeta}(x,\hat{y}) = \zeta(x,\hat{y})$ for $\hat{\gamma}$ a.e. $x \in X$. Q.E.D.

Proof of Proposition 8.4. Since $\int_S u(\phi(s),s,a,t)dP(s|a,t)$ is continuous in a for each t, Wagner (1977, Theorem 9.1 part (ii)) implies that there is a measurable function $\hat{a}: T \to A$ such that, for each t, $\hat{a}(t)$ maximizes $\int_S u(\phi(s),s,a,t)dP(s|a,t)$ among all $a \in A$ (recall that A is compact, and hence a maximum exists for each t). Then, the mechanism that recommends $\hat{a}(t)$ when the report is t and assigns the reward $\phi(s)$ when the signal is s is incentive compatible. Q.E.D.

Recall that $W_t = \{(\nu,a) \in \Delta(R\times S)\times A : \nu(\Phi) = 1 \text{ and } \text{marg}_S\nu = P_{a,t}\}$ is closed. Define nonnegative functions $\mathcal{L}$ and $\mathcal{U}$, each mapping $\Delta(R\times S)\times A\times T$ into $[0,\infty)$, by $\mathcal{L}(\nu,a,t) = \int_{R\times S} l(r,s,a,t)d\nu(r,s)$, and $\mathcal{U}(\nu,a,t) = \int_{R\times S} u(r,s,a,t)d\nu(r,s)$.

Lemma A.2 *For any $t \in T$, the functions $\mathcal{L}(\nu,a,t)$ and $\mathcal{U}(\nu,a,t)$ are nonnegative and lower semicontinuous in (ν,a) on W_t, and, for any $\varepsilon > 0$ such that $\varepsilon l(r,s,a,t) - u(r,s,a,t)$ is bounded below in (r,s,a), the function $\varepsilon\mathcal{L}(\nu,a,t) - \mathcal{U}(\nu,a,t)$ is lower semicontinuous in (ν,a) on W_t.*

Proof. Fix $t \in T$. We give the proof only for $\mathcal{L}(\nu,a,t)$ since the others are similar. The nonnegativity of $\mathcal{L}(\nu,a,t)$ follows from the nonnegativity of $l(r,s,a,t)$ so we need only show lower semicontinuity. Suppose that $(\nu_n,a_n) \to (\tilde{\nu},\tilde{a}) \in W_t \subseteq \Delta(R\times S)\times A$. We wish to show that $\underline{\lim}_n\mathcal{L}(\nu_n,a_n,t) \geq \mathcal{L}(\tilde{\nu},\tilde{a},t)$.

Let $D = \{(r,s,a) \in R\times S\times A : l(\cdot,t) \text{ is not l.s.c. at } (r,s,a)\}$. Then D is measurable since it is the set on which $l(\cdot,t)$ is not equal to its (lower semicontinuous) lower envelope. Let $D_{\tilde{a}} = \{(r,s) : (r,s,\tilde{a}) \in D\}$ be the slice of D through $\tilde{a}$, and let $D_{(s,\tilde{a})} = \{r : (r,s,\tilde{a}) \in D\}$ be the slice of D through $(s,\tilde{a})$.

Since $(\tilde{\nu},\tilde{a}) \in W_t$, the marginal of $\tilde{\nu}$ on S is $P_{\tilde{a},t}$. By Assumption 4.4, there is a measurable subset of signals, $\tilde{S}$ say, such that $P(\tilde{S}|\tilde{a},t) = 1$ and such that $l(\cdot,t)$ is lower semicontinuous at $(r,s,\tilde{a})$ for every $(r,s) \in R\times\tilde{S}$. Then,

$$\begin{aligned} \tilde{\nu}(D_{\tilde{a}}) &= \int_{R\times S} \mathbf{1}_{D_{(s,\tilde{a})}}(r)d\tilde{\nu}(r,s) \\ &= \int_{R\times\tilde{S}} \mathbf{1}_{D_{(s,\tilde{a})}}(r)d\tilde{\nu}(r,s) \\ &= 0, \end{aligned}$$

where the second equality follows since $\tilde{v}(R \times \tilde{S}) = P(\tilde{S}|\tilde{a}, t) = 1$ and the third because $D_{(s,\tilde{a})} = \emptyset$ for every $s \in \tilde{S}$.

Therefore, $l(\cdot, t)$ is lower semicontinuous at $(r, s, \tilde{a})$ except for $(r, s) \in D_{\tilde{a}}$, a subset of $R \times S$ that has $\tilde{\nu}$ measure zero. Hence,

$$\begin{aligned} \underline{\lim}_n \mathcal{L}(\nu_n, a_n, t) &= \underline{\lim}_n \int_{R \times S} l(r, s, a_n, t) d\nu_n(r, s) \\ &\geq \int_{R \times S} l(r, s, \tilde{a}, t) d\tilde{\nu}(r, s) \\ &= \mathcal{L}(\tilde{\nu}, \tilde{a}, t), \end{aligned}$$

where the inequality follows from Lemma A.1. Q.E.D.

Lemma A.3 *For every $t \in T$ and for every $c \in \mathbb{R}$, the set $\{\mu \in \Delta(W_t) : \int \mathcal{L}(\nu, a, t) d\mu(\nu, a) \leq c\}$ is compact.*

Proof. Let us first establish,

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

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

Fix $t \in T$, fix $c \in R$, and let $C = \{(\nu, a) \in W_t : L(\nu, a, t) \leq c\}$. Since, by Lemma A.2, $\mathcal{L}(\nu, a, t)$ is lower semicontinous in (ν, a) on the closed set W_t, the set C is closed. Hence, it suffices, by ($*$), to show that C is compact.

Fix any $\varepsilon > 0$. 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. Hence, there is a compact subset Y of S such that $P(Y|a, t) > 1 - \varepsilon/2$ for every $a \in A$. Let $D = \{(r, s) \in \Phi : s \in Y$ and there exists $a \in A$ such that $l(r, s, a, t) \leq 2c/\varepsilon\}$. Then, the closure of D is compact by Assumptions 4.1 and 4.5.

Consider any sequence $(\nu_n, a_n) \in C$. We must show that (ν_n, a_n) has a subsequence that converges to a point in C. Since $l \geq 0$ and $Y \subseteq S$,we have

$$\begin{aligned} c &\geq \int_{R \times S} l(r, s, a_n, t) d\nu_n(r, s) \\ &\geq \int_{R \times Y} l(r, s, a_n, t) d\nu_n(r, s). \end{aligned} \tag{A.1}$$

Since $\nu_n(\Phi) = 1$, we have $l(r, s, a_n, t) > 2c/\varepsilon$ for ν_n a.e. $(r, s) \in (R \times Y) \backslash D$. Therefore, since $l \geq 0$, (A.1) implies that $\nu_n((R \times Y) \backslash D) \leq \varepsilon/2$. Since $D \subseteq R \times Y$, $\nu_n((R \times Y) \backslash D) = \nu_n(R \times Y) - \nu_n(D)$ and so

$$\begin{aligned} \nu_n(D) &\geq \nu_n(R \times Y) - \varepsilon/2 \\ &= P(Y|a_n, t) - \varepsilon/2 \\ &> 1 - \varepsilon, \end{aligned}$$

where the equality follows because $(\nu_n, a_n) \in C \subseteq W_t$ implies that the marginal of ν_n on S is $P(\cdot|a_n, t)$. *A fortiori*, each ν_n gives the closure of D, a compact set, probability at least $1-\varepsilon$.

Thus $\{\nu_n\}$ is a tight set of measures and hence, by Prohorov's theorem, the sequence ν_n has a convergent subsequence. The sequence a_n, being in the compact set A, also has a convergent subsequence. Thus, $(\nu_n, a_n) \in C$ has a convergent subsequence whose limit, because C is closed, is in C. Hence, C is compact. Q.E.D.

The next lemma is based heavily on Balder (1990). Indeed, conclusions (ii) and (iii) are a direct application of Balder's Theorem 2.1. But conclusion (i) is new, and is a consequence of our more specialized environment.

Lemma A.4 *Let $\{\mu_n\}$ be a sequence of transition probabilities from T to $\Delta(\Delta(R\times S)\times A)$ such that $\mu_n(\cdot|t) \in \Delta(W_t)$ for every $t \in T$, and $\sup_n \int_{\Delta(R\times S)\times A\times T} \mathcal{L}(\nu, a, t)d\mu_n(\nu, a|t)dH(t) < \infty$. Then there is a transition probability μ^* from T to $\Delta(\Delta(R\times S)\times A)$, there is a measurable function $\phi : T \to \mathbb{R}$, and there is a subsequence $\{n_j\}$ of $\{n\}$, such that, (i) $H \otimes \mu_{n_j} \to H \otimes \mu^*$, (ii) the Cesaro mean of $\{\mu_{n_j}(\cdot|t)\}$ converges to $\mu^*(\cdot|t)$, H a.e. $t \in T$, (iii) the Cesaro mean of $\{\int_{\Delta(R\times S)\times A} \mathcal{L}(\nu, a, t)d\mu_{n_j}(\nu, a|t)\}$ converges to $\phi(t)$, H a.e. $t \in T$, and (iv) $\int_{\Delta(R\times S)\times A} \mathcal{L}(\nu, a, t)d\mu^*(\nu, a|t) \le \phi(t)$, H a.e. $t \in T$.*[53]

Proof. Lemma A.3 implies that for every $c \in \mathbb{R}$ and for every $t \in T$ the set of measures $\Lambda_{c,t} = \{\mu \in \Delta(W_t) : \int_{\Delta(R\times S)\times A} \mathcal{L}(\nu, a, t)d\mu(\nu, a) \le c\}$ is relatively compact.

Let $X = \Delta(R \times S) \times A$. Since X is Polish, it contains countably many open subsets $U_1, U_2, \ldots$ that generate its topology. Let d be any metric on X and for each $i, k \in \mathbb{N}$, let $\xi_{i,k} : X \to \mathbb{R}$ be the continuous function defined by $\xi_{i,k}(x) = \frac{1}{1+kd(x,U_i)}$, and for each $\mu \in \Delta(X)$, define $\alpha_{i,k}(\mu) = \int_Y \xi_{i,k}(x)d\mu(x)$. Then $\mathcal{A} = \{\alpha_{i,k}\}$ is a countable set of affine continuous functions that countably separates $\Delta(X)$ as defined in Balder (1990). The desired results (ii), (iii), and (iv) can now be obtained by following the proof of Theorem 2.1 in Balder (1990), whose equation (2.6) shows that there is $c \in \mathbb{R}$, which may depend on t, such that the Cesaro mean of $\{\mu_n(\cdot|t)\}$ is contained in $\Lambda_{c,t}$.[54] This latter fact, since $\Lambda_{c,t}$ is relatively compact, can then be used in place of Balder's inf-compactness assumption on $h(t, \mu) = \int_{\Delta(R\times S)\times A} \mathcal{L}(\nu, a, t)d\mu(\nu, a)$.

Thus, we may conclude that there are mappings $\mu^* : T \to \Delta(X)$ and $\phi : T \to \mathbb{R}$, and that there is a subsequence $\{n'_j\}$ of $\{n\}$ such that the Cesaro mean of any subsequence of $\{\mu_{n'_j}(\cdot|t)\}$ converges to $\mu^*(\cdot|t)$ and the Cesaro mean of any subsequence of $\{\int_X \mathcal{L}(x, t)d\mu_{n'_j}(x|t)\}$ converges to $\phi(t)$, with both limits holding for H a.e. $t \in T$.[55]

Since ϕ is the pointwise H-a.e. limit of measurable functions, it is measurable. Each $\mu_n(\cdot|\cdot)$, being a transition probability, is a measurable function from T to $\Delta(X)$ by BS, Proposition 7.26. Therefore, since $\mu^*(\cdot|t)$ is the pointwise H-a.e. limit of the Cesaro mean of $\{\mu_{n'_j}(\cdot|t)\}$, $\mu^* : T \to \Delta(X)$ is measurable and hence a transition probability by the same proposition.

[53] In fact, the subsequence $\{n_j\}$ can be chosen so that, for all further subsequences, the two Cesaro mean convergence results, (ii) and (iii), hold also for the further subsequence. But we will not need this stronger result.

[54] For (iv), note that $L(\nu, a, t)$ is lower semicontinous in (ν, a) on W_t by Lemma A.2.

[55] The H measure zero set of types for which the limits fail to hold can depend on the subsequence of $\{n'_j\}$ that is chosen.

To complete the proof, it suffices to show that there is a subsequence $\{n_j\}$ of $\{n'_j\}$ such that (i) holds. Then, by what we have already proven, (ii) and (iii) will also hold. So, we turn to (i).

Since the Cesaro mean of any subsequence of $\{\mu_{n'_j}(\cdot|t)\}$ converges to $\mu^*(\cdot|t)$ for H a.e. $t \in T$, the Cesaro mean of any subsequence of $\{H \otimes \mu_{n'_j}\}$ converges to $H \otimes \mu^*$ (use the definition of weak* convergence and apply the dominated convergence theorem). To show that (i) holds, it suffices to show that there is a subsequence $\{n_j\}$ of $\{n'_j\}$ such that $\{H \otimes \mu_{n_j}\}$ converges since, if it converges, it must converge to $H \otimes \mu^*$ because, as we have just seen, its Cesaro mean converges to $H \otimes \mu^*$. By Prohorov's theorem, it therefore suffices to show that the sequence $\{H \otimes \mu_{n'_j}\}$ of probability measures in $\Delta(X \times T)$ has a tight subsequence.

For every j, let $v_j = H \otimes \mu_{n'_j}$, and for every m, let $\bar{\nu}_m = \frac{1}{m}\sum_{j=1}^{m} \nu_j$ be the m-th Cesaro mean of $\{H \otimes \mu_{n'_j}\}$. As argued in the previous paragraph, $\bar{\nu}_m \to H \otimes \mu^*$. Therefore, by Prohorov's theorem, $\{\bar{\nu}_m\}$ is a tight set in $\Delta(X \times T)$.

Let α_k and β_k be strictly increasing sequences of positive real numbers that converge to 1 and that satisfy $(1-\alpha_k)/(1-\beta_k) = (1/2)^{k+1}$.[56] The tightness of the set $\{\bar{\nu}_m\}$ implies that, for each positive integer k, there is a compact subset C_k of $X \times T$ such that $\bar{\nu}_m(C_k) > \alpha_k$ holds for every $m = 1, 2, ...$ Therefore,

$$\frac{\nu_1(C_k) + ... + \nu_m(C_k)}{m} > \alpha_k, \text{ for every } k, m \in \{1, 2, ...\}.$$

Letting $J_{k,m} = \{j \leq m : \nu_j(C_k) \leq \beta_k\}$, and $J^C_{k,m} = \{1, \ldots, m\} \setminus J_{k,m}$, we have

$$\begin{aligned} m\alpha_k &< \sum_{j \in J_{k,m}} \nu_j(C_k) + \sum_{j \in J^C_{k,m}} \nu_j(C_k) \\ &\leq \#J_{k,m}\beta_k + (m - \#J_{k,m})1, \end{aligned}$$

so that

$$\#J_{k,m} \leq \frac{1-\alpha_k}{1-\beta_k} m = \frac{1}{2^{k+1}} m, \text{ for every } k, m \in \{1, 2, ...\}.$$

Consequently,

$$\# \cup_{k=1}^{\infty} J_{k,m} \leq \sum_{k=1}^{\infty} \frac{1}{2^{k+1}} m = \frac{m}{2}, \text{ for every } m \in \{1, 2, ...\}.$$

But then, since $\# \{j \leq m\} = m$,

$$\# \cap_{k=1}^{\infty} J^C_{k,m} \geq m - \frac{m}{2} = \frac{m}{2}, \text{ for every } m \in \{1, 2, ...\}.$$

The last inequality, since it holds for every positive integer m, implies that there are infinitely many indices $j \in \{1, 2, ...\}$ such that,

$$\nu_j(C_k) > \beta_k \text{ for all } k \in \{1, 2, ...\}.$$

Since $\beta_k \to 1$, the collection of these indices j furnishes a tight subsequence of the sequence

[56] For example, $\alpha_k = 1 - \frac{1}{k2^{k+1}}$ and $\beta_k = 1 - \frac{1}{k}$.

$\nu_1, \nu_2, ...$, as desired. Q.E.D.

Lemma A.5 *Let $\{\mu_n\}$, μ^*, and ϕ be as in Lemma A.4 and suppose that the subsequence $\{n_j\}$ there is reindexed as $\{n\}$. Recall that $\mathcal{U}(\nu, a, t) = \int_{R\times S} u(r, s, a, t) d\nu(r, s)$. Then, for H a.e. $t \in T$,*

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

where $\bar{\mu}_n$ denotes the n^{th} Cesaro mean of $\{\mu_n\}$.

Proof. By Lemma A.4 (ii) and (iii), we may let $t \in T$ be any type from the H measure one set types such that $\bar{\mu}_n(\cdot|t) \to \mu^*(\cdot|t)$ and $\int_{\Delta(R\times S)\times A} \mathcal{L}(\nu, a, t) d\bar{\mu}_n(\nu, a|t) \to \phi(t)$. Fix this $t \in T$ for the remainder of the proof.

The function $\mathcal{U}(\cdot, t)$ is, by Lemma A.2, lower semicontinuous. Therefore, because $\bar{\mu}_n(\cdot|t) \to \mu^*(\cdot|t)$,

$$\liminf_n \int_X \mathcal{U}(x, t) d\bar{\mu}_n(x|t) \geq \int_X \mathcal{U}(x, t) d\mu^*(x|t), \ H \text{ a.e. } t \in T. \tag{A.2}$$

It remains to establish the reverse inequality for the limsup.

Let $Z = R \times S \times A$. By Assumption 4.6,

$$\limsup_{z\in Z: u(z,t)\to\infty} \frac{u(z, t)}{l(z, t)} = 0. \tag{A.3}$$

We claim that, for every $\varepsilon > 0$, there exists $c_\varepsilon > 0$ (c_ε may depend on t) such that

$$u(z, t) \leq \varepsilon l(z, t) + c_\varepsilon, \ \forall z \in Z. \tag{A.4}$$

To see this, suppose the contrary. Then, there exists $\varepsilon_0 > 0$, and, for every positive integer m, there exists $z_m \in Z$ such that,

$$u(z_m, t) > \varepsilon_0 l(z_m, t) + m. \tag{A.5}$$

Then, $u(z_m, t) \to_m \infty$. But (A.5) implies that $u(z_m, t) > \varepsilon_0 l(z_m, t)$, contradicting (A.3) and establishing the claim.

Inequality (A.4) is employed in Balder (1990), and we follow his usage here. By the definitions of $\mathcal{U}$ and $\mathcal{L}$, it follows that, for every $\varepsilon > 0$,

$$\mathcal{U}(x, t) \leq \varepsilon\mathcal{L}(x, t) + c_\varepsilon, \ \forall x \in X. \tag{A.6}$$

Consequently,

$$\begin{aligned} \int \mathcal{U}(x, t) d\bar{\mu}_n(x|t) &\leq \varepsilon \int \mathcal{L}(x, t) d\mu_n(x|t) + c_\varepsilon \\ &\to \varepsilon\phi(t) + c_\varepsilon < \infty. \end{aligned}$$

Therefore, the nonnegative sequence $\{\int \mathcal{U}(x, t) d\bar{\mu}_n(x|t)\}$ is bounded and so if we let $\sigma = \limsup_n \int \mathcal{U}(x, t) d\bar{\mu}_n(x|t)$, then $\sigma \in [0, \infty)$.

Note that for t fixed and for any $\varepsilon > 0$, $\varepsilon l(z, t) - u(z, t)$, as a function of z is bounded below (by (A.4)). It follows that $\varepsilon\mathcal{L}(x, t) - \mathcal{U}(x, t)$ is lower semicontinuous in x on W_t (by

Lemma A.2) and is bounded below (by (A.6)). Thus, for any $\varepsilon > 0$,

$$\begin{aligned}\varepsilon\phi(t) - \sigma &= \lim_n \int \varepsilon\mathcal{L}(x,t)d\mu_n(x|t) - \limsup_n \int \mathcal{U}(x,t)d\bar{\mu}_n(x|t) \\ &= \liminf_n \int (\varepsilon\mathcal{L}(x,t) - \mathcal{U}(x,t))\, d\bar{\mu}_n(x|t) \\ &\geq \int (\varepsilon\mathcal{L}(x,t) - \mathcal{U}(x,t))\, d\mu^*(x|t) \\ &= \varepsilon \int \mathcal{L}(x,t)d\mu^*(x|t) - \int \mathcal{U}(x,t)d\mu^*(x|t),\end{aligned}$$

where the inequality follows because $\varepsilon\mathcal{L}(x,t) - \mathcal{U}(x,t)$ is bounded below and lower semicontinuous on W_t. Taking the limit as $\varepsilon \to 0$ of the inequality $\varepsilon\phi(t) - \sigma \geq \varepsilon \int \mathcal{L}(x,t)d\mu^*(x|t) - \int \mathcal{U}(x,t)d\mu^*(x|t)$ implies, since $\int \mathcal{L}(x,t)d\mu^*(x|t)$ is finite by Lemma A.4 (iv), that $\sigma \leq \int \mathcal{U}(x,t)d\mu^*(x|t)$, as desired. Q.E.D.

Recall from Step 5 of the proof of Theorem 8.3 that

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

where $u(r,s,a,t) - u_*(s,a,t) \geq 0$ for all (r,s,a,t). Also, either $u_*(s,a,t) = u(r_*,s,a,t)$ for all (s,a,t) or $u_*(s,a,t) = 0$ for all (s,a,t), and so in either case (using Assumption 4.4 for the first case), for every $(a,t) \in A\times T$, $u_*(\cdot,t)$ is continuous at (s,a) for $P_{a,t}$ a.e. $s \in S$.

Lemma A.6 *Let $\{\bar{\mu}_n\}$ and μ^* be as in Lemma A.5. For H a.e. $t, t' \in T$,*

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

Proof. Because $\bar{\mu}_n(W_{t'}|t') = \mu^*(W_{t'}|t') = 1$, 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 $U_*(\nu_0,a_0,t,a'_0,t') + \varepsilon \geq \sup_{a\in A} U_*(\nu_0,a,t,a'_0,t')$. Choose a sequence a_n in A converging to a_0 according to Assumption 4.8. For any $(a,t),(a',t') \in A\times T$, recall that $g_{a,t/a',t'} : S \to [0,\infty]$ denotes the Radon-Nikodym derivative of $P_{a,t}(\cdot \cap S_{a',t'})$ with respect to $P_{a',t'}$.

For every $m \in \{1,2,...\}$ define $\xi_m : S \to [0,+\infty]$ and $\underline{\xi}_m : S \to [0,+\infty]$ by

$$\xi_m(s) = \inf\{g_{a_m,t/a'_m,t'}(s), g_{a_{m+1},t/a'_{m+1},t'}(s), ...\} \text{ and } \underline{\xi}_m(s) = \liminf_{s'\to s} \xi_m(s').$$

Then, $\underline{\xi}_m$ is lower semicontinuous (it is the lower envelope of the function ξ_m) and $\underline{\xi}_m \leq \xi_m \leq g_{a_n,t/a'_n,t'}$ for all $n \geq m$.

Since $\xi_m \leq \xi_{m+1}$ implies $\underline{\xi}_m \leq \underline{\xi}_{m+1}$, we may define $\underline{\xi} : S \to [0,+\infty]$ by

$$\underline{\xi}(s) = \lim_m \underline{\xi}_m(s).$$

By the definition of U_*, for each n,

$$U_*(\nu_n, a_n, t, a_n', t') = \int_S u_*(s, a_n, t)\, dP(s|a_n, t) + \int_{R\times S} [u(r, s, a_n, t) - u_*(s, a_n, t)] g_{a_n,t/a_n',t'}(s)\, d\nu_n(r, s). \tag{A.7}$$

Consider the first term on the right-hand side of (A.7) Because $u_*(\cdot, t)$ is nonnegative and continuous at (s, a_0) for $P_{a_0,t}$ a.e. $s \in S$, Lemma A.1 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 (A.7), letting $\hat{u}$ denote the nonnegative function $u - u_*$, we have for every m,

$$\begin{aligned}
\underline{\lim}_n \int_{R\times S} \hat{u}(r, s, a_n, t)\, g_{a_n,t/a_n',t'}(s)\, d\nu_n(r, s) &\geq \underline{\lim}_n \int_{R\times S} \hat{u}(r, s, a_n, t)\, \underline{\xi}_m(s) d\nu_n(r, s) \\
&\geq \int_{R\times S} \hat{u}(r, s, a_0, t)\, \underline{\xi}_m(s) d\nu_0(r, s) \\
&\to_m \int_{R\times S} \hat{u}(r, s, a_0, t)\, \underline{\xi}(s) d\nu_0(r, s) \\
&\geq \int_{R\times S} \hat{u}(r, s, a_0, t)\, g_{a_0,t/a_0',t'}(s) d\nu_0(r, s),
\end{aligned}$$

where the first inequality follows because $g_{a_n,t/a_n',t'}(s) \geq \underline{\xi}_m(s)$ for all $n \geq m$, the second inequality follows by Lemma A.1 because $\hat{u}(\cdot, t)\underline{\xi}_m(\cdot)$ is nonnegative and lower semicontinuous at (r, s, a_0) for ν_0 a.e. (r, s) (since the marginal of ν_0 on S is $P_{a_0,t}$ and by Assumption 4.4; see the proof of Lemma A.1 for a complete and similar argument), the limit follows by the monotone convergence theorem since $\underline{\xi}_m(s) \uparrow \underline{\xi}(s)$, and the final inequality follows because $\hat{u}$ is nonnegative and Assumption 4.8 implies that $\underline{\xi}(s) \geq g_{a_0,t/a_0',t'}(s)$ for all $s \in S$. Hence,

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

Since $\varepsilon > 0$ is arbitrary, we are done. Q.E.D.

Lemma A.7 *Recall that our space of types T is Polish with prior $H \in \Delta(T)$. Let X be any Polish space, and let $\beta_n \in \Delta(X \times T)$ be any sequence that weak* converges to $\beta \in \Delta(X \times T)$ such that the marginal of each β_n on T is H. Then,*
(i) $\lim_n \int g(x, t) d\beta_n = \int g(x, t) d\beta$ for every measurable $g : X \times T \to \mathbb{R}$ such that $g(x, t)$ is continuous in x for each t and such that there exists a measurable $\xi : T \to \mathbb{R}$ such that $\int |\xi(t)|\, dH(t) < \infty$ and $|g(x, t)| \leq |\xi(t)|$ for every (x, t), and,
(ii) $\underline{\lim}_n \int g(x, t) d\beta_n \geq \int g(x, t) d\beta$ for every measurable $g : X \times T \to \mathbb{R}$ such that $g(x, t)$ is lower semicontinuous in x for each t and such that there exists a measurable $\xi : T \to \mathbb{R}$ such that $\int |\xi(t)|\, dH(t) < \infty$ and $g(x, t) \geq \xi(t)$ for every (x, t).

Proof. Since $\beta_n \to \beta$ and the marginal of each β_n on T is H, the marginal of β on T is also H. So, by Corollary 7.27.2 of Bertsekas and Shreve and by parts (a) and (b) of Theorem 2.2 in Balder (1988), (i) and (ii) are equivalent. Hence it suffices to prove (i). Suppose then that $g : X \times T \to \mathbb{R}$ and $\xi : T \to \mathbb{R}$ are measurable, that $g(x,t)$ is continuous in x for each t, that $|g(x,t)| \leq |\xi(t)|$ for every (x,t), and that $\int |\xi|\, dH < \infty$.

For any positive integer m define

$$g_m(x,t) = \begin{cases} m, & \text{if } g(x,t) > m \\ -m, & \text{if } g(x,t) < -m \\ g(x,t), & \text{otherwise,} \end{cases}$$

and let $T_m = \{t \in T : |\xi(t)| \leq m\}$.

Fix any $\varepsilon > 0$. Since $\int |\xi(t)|\, dH(t) < \infty$, we may choose m so that $\int_{T \setminus T_m} |\xi(t)|\, dH(t) < \varepsilon/2$. Since $g - g_m = 0$ on $X \times T_m$ and $|g - g_m| \leq |g| \leq |\xi|$, and since the marginal of β on T is H, we have

$$\begin{aligned} \left| \int g d\beta_n - \int g d\beta \right| &\leq \left| \int (g - g_m) d\beta_n \right| + \left| \int g_m d\beta_n - \int g_m d\beta \right| + \left| \int (g_m - g) d\beta \right| \\ &\leq \int_{X \times (T \setminus T_m)} |g|\, d\beta_n + \left| \int g_m d\beta_n - \int g_m d\beta \right| + \int_{X \times (T \setminus T_m)} |g|\, d\beta \\ &\leq \left| \int g_m d\beta_n - \int g_m d\beta \right| + 2 \int_{T \setminus T_m} |\xi|\, dH \\ &\leq \left| \int g_m d\beta_n - \int g_m d\beta \right| + \varepsilon. \end{aligned} \tag{A.8}$$

So, if for every positive integer m,

$$\int g_m d\beta_n \to_n \int g_m d\beta, \tag{A.9}$$

then, by (A.8), $\limsup_n \left| \int g d\beta_n - \int g d\beta \right| \leq \varepsilon$. Since $\varepsilon > 0$ is arbitrary, (i) would hold and the proof would be complete. Hence, it suffices to prove (A.9).

Fix any positive integer m and any $\varepsilon > 0$. Since β_n weak* converges to β, the set of probability measures $\{\beta, \beta_1, \beta_2, ...\}$ is compact and hence tight by Prohorov's theorem. Therefore, there is a compact $C \subseteq X \times T$ such that $\beta(C) > 1 - \varepsilon/8m$ and $\beta_n(C) > 1 - \varepsilon/8m$ for every n. Let X_1 be the projection of C on X and let T_1 be the projection of C on T. Then $X_1 \times T_1$ is compact and contains C, and so $\beta(X_1 \times T_1) > 1 - \varepsilon/8m$ and $\beta_n(X_1 \times T_1) > 1 - \varepsilon/8m$ for every n.

Let γ be the restriction of g_m to $X_1 \times T_1$ and let H_1 be the restriction of H to the Borel subsets of T_1. By Theorem 2.1 in Jacobs (1967) (and because a continuous function on a compact set is uniformly continuous), there is an open subset U_1 of T_1 (i.e., $U_1 = U \cap T_1$ for some open subset U of T) such that $H_1(U_1) < \varepsilon/8m$ and such that γ_1, the restriction of γ to $D := X_1 \times (T_1 \setminus U_1)$, is continuous. By the Tietze extension theorem, we may extend γ_1 to a continuous function $\tilde{\gamma} : X \times T \to \mathbb{R}$ such that $|\tilde{\gamma}| \leq m$ (since $|\gamma_1| \leq m$ on $X_1 \times (T_1 \setminus U_1)$).

Note that $g_m = \tilde{\gamma}$ on D and that $|g_m - \tilde{\gamma}| \leq |g_m| + |\tilde{\gamma}| \leq 2m$. Hence, for every n,

$$\begin{aligned}\left|\int g_m d\beta_n - \int g_m d\beta\right| &\leq \left|\int (g_m - \tilde{\gamma}) d\beta_n\right| + \left|\int \tilde{\gamma} d\beta_n - \int \tilde{\gamma} d\beta\right| + \left|\int (\tilde{\gamma} - g_m) d\beta\right| \\ &= \left|\int_{(X\times T)\backslash D} (g_m - \tilde{\gamma}) d\beta_n\right| + \left|\int \tilde{\gamma} d\beta_n - \int \tilde{\gamma} d\beta\right| + \left|\int_{(X\times T)\backslash D} (\tilde{\gamma} - g_m) d\beta\right| \\ &\leq 2m\beta_n((X\times T)\backslash D) + \left|\int \tilde{\gamma} d\beta_n - \int \tilde{\gamma} d\beta\right| + 2m\beta((X\times T)\backslash D), \quad \text{(A.10)}\end{aligned}$$

and, because the marginal of β_n on T is H,

$$\begin{aligned}\beta_n((X\times T)\backslash D) &= \beta_n(X\times T) - \beta_n(D) \\ &= 1 - \beta_n(X_1 \times (T_1\backslash U_1)) \\ &= 1 - \beta_n(X_1\times T_1) + \beta_n(X_1 \times U_1) \\ &< 1 - (1-\varepsilon/8m) + \beta_n(X\times U_1) \\ &= \varepsilon/8m + H_1(U_1) \\ &< \varepsilon/4m, \quad \text{(A.11)}\end{aligned}$$

and similarly,

$$\beta((X\times T)\backslash D) < \varepsilon/4m. \quad \text{(A.12)}$$

Substituting (A.11) and (A.12) into (A.10) gives, for every n,

$$\left|\int g_m d\beta_n - \int g_m d\beta\right| < \left|\int \tilde{\gamma} d\beta_n - \int \tilde{\gamma} d\beta\right| + \varepsilon.$$

Since $\tilde{\gamma}$ is continuous and bounded on $X \times T$, $\lim_n \left|\int \tilde{\gamma} d\beta_n - \int \tilde{\gamma} d\beta\right| = 0$ and so $\limsup_n \left|\int g_m d\beta_n - \int g_m d\beta\right| \leq \varepsilon$. Since $\varepsilon > 0$ was arbitrary, this proves (A.9) and completes the proof. Q.E.D.

Proof for Example 11.6. To see that M' is d_M-closed, suppose that the sequence $(\alpha_n, \kappa_n) \in M'$ d_M-converges to $(\hat{\alpha}, \hat{\kappa})$. 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

$$H \otimes \hat{\alpha} \otimes P \otimes \hat{\kappa} = H \otimes \alpha_1 \otimes \gamma \otimes \kappa_1, \quad \text{(A.13)}$$

where $\alpha_1 : T \to \Delta(R_2 \times A)$, $\gamma : R_2 \times A \times T \to \Delta(S)$, and $\kappa_1 : R_2 \times S \times A \times T \to \Delta(R_1)$ are transition probabilities. Since $(\alpha_n, \kappa_n) \in M'$ d_M-converges to $(\hat{\alpha}, \hat{\kappa})$, the sequence $H \otimes \alpha_{1n} \otimes P \otimes \kappa_{1n}$ weak* converges to $H \otimes \alpha_1 \otimes \gamma \otimes \kappa_1$ and so all of the various marginals also weak* converge. In particular, $H \otimes \alpha_{1n} \otimes P$ weak* converges to $H \otimes \alpha_1 \otimes \gamma$ and $H \otimes \alpha_{1n}$ weak* converges to $H \otimes \alpha_1$. Consequently, for any continuous and bounded $g : R_2 \times S \times A \times T \to \mathbb{R}$, $\int g d\gamma d\alpha_1 dH = \lim_n \int g dP d\alpha_{1n} dH = \int g dP d\alpha_1 dH$, where the second equality follows from Lemma A.7 part (i), since the continuity of $P_{a,t}$ in a for each t implies that $\int_S g(r_2, s, a, t) dP(s|a, t)$, a measurable function of (r_2, a, t), is continuous in (r_2, a) for each t. Hence, $\int g d\gamma d\alpha_1 dH = \int g dP d\alpha_1 dH$ for any bounded and continuous g, and so $H \otimes \alpha_1 \otimes \gamma = H \otimes \alpha_1 \otimes P$ which, by (A.13) implies that $H \otimes \hat{\alpha} \otimes P \otimes \hat{\kappa} = H \otimes \alpha_1 \otimes P \otimes \kappa_1$. We conclude that $(\hat{\alpha}, \hat{\kappa}) \in M'$, as desired. Q.E.D.

Proof for Example 11.7. We will show that $M' \cap M'' \cap M_c$ is d_M-closed for every $c \in \mathbb{R}$. Since M' is d_M-closed by Example 11.6, 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, suppose that (α_n, κ_n) is any sequence in $M'' \cap M_c$. The proof of Theorem 8.3 establishes that there is a subsequence $(\alpha_{n_j}, \kappa_{n_j})$ that d_M-converges to some $(\alpha^*, \kappa^*) \in M_c$ such that (see Lemma A.5) the Cesaro mean of $\int u\,(r,s,a,t)\,d\kappa_{n_j}\,(r|s,a,t)\,dP(s|a,t)d\alpha_{n_j}\,(a|t)$ converges to $\int u\,(r,s,a,t)\,d\kappa^*\,(r|s,a,t)\,dP(s|a,t)d\alpha^*\,(a|t)$ for H a.e. t. Since, for each j, $\int u\,(r,s,a,t)\,d\kappa_{n_j}\,(r|s,a,t)\,dP(s|a,t)d\alpha_n\,(a|t) \geq u_O\,(t)$ for H a.e. t, we may conclude that $\int u\,(r,s,a,t)\,d\kappa^*\,(r|s,a,t)\,dP(s|a,t)d\alpha^*\,(a|t) \geq u_O(t)$ for H a.e. t, i.e., that $(\alpha^*, \kappa^*) \in M''$. It follows that $M'' \cap M_c$ is d_M-compact. Q.E.D.

Proof for Example 11.8. We will show that $M' \cap M_c$ is d_M-closed for every $c \in \mathbb{R}$. But in fact, $M' \cap M_c$ is d_M-compact. Indeed, suppose that (α_n, κ_n) is any sequence in $M' \cap M_c$. The proof of Theorem 8.3 establishes that there is a subsequence $(\alpha_{n_j}, \kappa_{n_j})$ that d_M-converges to some $(\alpha^*, \kappa^*) \in M_c$ such that (see Lemma A.5) the Cesaro mean of $\int u\,(r,s,a,t)\,d\kappa_{n_j}\,(r|s,a,t)\,dP(s|a,t)d\alpha_{n_j}\,(a|t)$ converges to $\int u\,(r,s,a,t)\,d\kappa^*\,(r|s,a,t)\,dP(s|a,t)d\alpha^*\,(a|t)$ for H a.e. t. By the dominated convergence theorem (recall that u is nonnegative), $\int ud[H \otimes \alpha_{n_j} \otimes P \otimes \kappa_{n_j}]$ converges to $\int ud[H \otimes \alpha^* \otimes P \otimes \kappa^*]$. Since, for each j, $\int ud[H \otimes \alpha_{n_j} \otimes P \otimes \kappa_{n_j}] \geq u_O$, we may conclude that $\int ud[H \otimes \alpha^* \otimes P \otimes \kappa^*] \geq u_O$, i.e., that $(\alpha^*, \kappa^*) \in M'$. It follows that $M' \cap M_c$ is d_M-compact. Q.E.D.

Proof of Theorem 13.5. Let (α^n, κ^n) be a sequence of incentive compatible mechanisms for the multi-agent model such that the principal's expected loss along the sequence converges to the infimum of her expected losses, l_*, among all incentive compatible mechanisms. That is,

$$\lim_n \int l(r,s,a,t)d\kappa^n(dr|s,a,t)dP(s|a,t)d\alpha^n(a|t)dH(t) = l_*.$$

As in Section 13.3, for every agent i and for every n, we may define transition probabilities $\alpha_i^n : T_i \to \Delta(A_i)$ and $\kappa_i^n : S \times A_i \times T_i \to \Delta(A_{-i} \times T_{-i})$ so that $\bar{H} \otimes \alpha^n \otimes Q \otimes \kappa^n = H_i \otimes \alpha_i^n \otimes Q \otimes \kappa_i^n \in \Delta(R \times S \times T \times A)$. Then, by Proposition 13.4, (α_i^n, κ_i^n) is incentive compatible in the single-agent-i model.

By hypothesis, for every i, the single-agent-i model satisfies Assumptions 4.1-4.6. Moreover, since in the single-agent-i model the signal is drawn according to $Q \in \Delta(S)$ for every a, t, Assumptions 4.7, 4.8, and 4.9 hold trivially. Consequently, the hypotheses of Theorem 8.3 hold in the in the single-agent-i model for every i.

By the distributional mechanism compactness result established in Theorem 8.3 we may choose a single subsequence $\{n_j\}$ of $\{n\}$ such that, for each i, there is a mechanism (α_i^*, κ_i^*) that is incentive compatible for the single-agent-i model and

$$\lim_j H_i \otimes \alpha_i^{n_j} \otimes Q \otimes \kappa_i^{n_j} = H_i \otimes \alpha_i^* \otimes Q \otimes \kappa_i^*, \text{ for every } i,$$

where, for each i, convergence is with respect to the weak* topology on $\Delta(R_i \times S \times A_i \times T_i) = \Delta(R \times S \times A \times T)$.

Since $\bar{H} \otimes \alpha^{n_j} \otimes Q \otimes \kappa^{n_j} = H_i \otimes \alpha_i^{n_j} \otimes Q \otimes \kappa_i^{n_j}$ for every i and j,

$$\lim_j \bar{H} \otimes \alpha^{n_j} \otimes Q \otimes \kappa^{n_j} = \lim_j H_i \otimes \alpha_i^{n_j} \otimes Q \otimes \kappa_i^{n_j} = H_i \otimes \alpha_i^* \otimes Q \otimes \kappa_i^*, \text{ for every } i.$$

Because for every j the marginal of $\bar{H} \otimes \alpha^{n_j} \otimes Q \otimes \kappa^{n_j}$ on T is $\bar{H}$, the marginal of $\lim_j \bar{H} \otimes \alpha^{n_j} \otimes Q \otimes \kappa^{n_j}$ on T is also $\bar{H}$. By BS Proposition 7.27, we may therefore decompose

$\lim_j \bar{H}\otimes\alpha^{n_j}\otimes Q\otimes\kappa^{n_j}$ (an element of $\Delta(R\times S\times A\times T)$) as $\bar{H}\otimes\hat{\alpha}\otimes\tilde{P}\otimes\hat{\kappa}$ for some transition probabilities $\hat{\alpha}: T\to\Delta(A)$, $\tilde{P}: A\times T\to S$, and $\hat{\kappa}: S\times A\times T\to\Delta(R)$.

Since $\bar{H}\otimes\alpha^{n_j}\otimes Q\otimes\kappa^{n_j}\to^* \bar{H}\otimes\hat{\alpha}\otimes\tilde{P}\otimes\hat{\kappa}$ implies that $\bar{H}\otimes\alpha^{n_j}\otimes Q\to^* \bar{H}\otimes\hat{\alpha}\otimes\tilde{P}$ and that $\bar{H}\otimes\alpha^{n_j}\to^* \bar{H}\otimes\hat{\alpha}$, we claim that $\bar{H}\otimes\hat{\alpha}\otimes\tilde{P}\otimes\hat{\kappa}=\bar{H}\otimes\hat{\alpha}\otimes Q\otimes\hat{\kappa}$. To see this, let $g: S\times A\times T\to\mathbb{R}$ be continuous and bounded. Then

$$\begin{aligned}\lim_j\int g(s,a,t)dQ(s)d\alpha^{n_j}(a|t)d\bar{H}(t) &= \lim_j\int_{A\times T}\left(\int_S g(s,a,t)dQ(s)\right)d\alpha^{n_j}(a|t)d\bar{H}(t)\\ &= \int_{A\times T}\left(\int_S g(s,a,t)dQ(s)\right)d\hat{\alpha}(a|t)d\bar{H}(t)\\ &= \int g(s,a,t)dQ(s)d\hat{\alpha}(a|t)d\bar{H}(t),\end{aligned}$$

where the second equality follows by (A.7) part (i) because $\bar{H}\otimes\alpha^{n_j}\to^* \bar{H}\otimes\hat{\alpha}$ and the function of (a,t) in parentheses is continuous in a for each t. Since g is an arbitrary continuous and bounded function and the left-hand side limit is equal to the last expression on the right-hand side, we may conclude that $\bar{H}\otimes\alpha^{n_j}\otimes Q\to^* \bar{H}\otimes\hat{\alpha}\otimes Q$. But since $\bar{H}\otimes\alpha^{n_j}\otimes Q\to^* \bar{H}\otimes\hat{\alpha}\otimes\tilde{P}$ we must have $\bar{H}\otimes\hat{\alpha}\otimes\tilde{P}=\bar{H}\otimes\hat{\alpha}\otimes Q$ and so $\bar{H}\otimes\hat{\alpha}\otimes\tilde{P}\otimes\hat{\kappa}=\bar{H}\otimes\hat{\alpha}\otimes Q\otimes\hat{\kappa}$ as claimed.

So, we have shown that

$$\begin{aligned}\bar{H}\otimes\hat{\alpha}\otimes Q\otimes\hat{\kappa} &= \lim_j \bar{H}\otimes\alpha^{n_j}\otimes Q\otimes\kappa^{n_j}\\ &= \lim_j H_i\otimes\alpha_i^{n_j}\otimes Q\otimes\kappa_i^{n_j}\\ &= H_i\otimes\alpha_i^*\otimes Q\otimes\kappa_i^*,\ \text{for every } i. \qquad (A.14)\end{aligned}$$

Since each (α_i^*,κ_i^*) is incentive compatible for the single-agent-i model, Proposition 13.4 implies that $(\hat{\alpha},\hat{\kappa})$ is incentive compatible for the multi-agent model.

Finally, for any agent i,

$$\begin{aligned}l_* &= \lim_j\int l(r,s,a,t)d\kappa^{n_j}(r|s,a,t)dP(s|a,t)d\alpha^{n_j}(a|t)dH(t)\\ &= \lim_j\int l(r,s,a,t)h(t)f(s|a,t)d\kappa^{n_j}(r|s,a,t)dQ(s)d\alpha^{n_j}(a|t)d\bar{H}(t)\\ &= \lim_j\int \tilde{l}(r_i,s,a_i,t_i)d\kappa_i^{n_j}(r_i|s,a_i,t_i)dQ(s)d\alpha_i^{n_j}(a_i|t_i)dH_i(t_i)\\ &\geq \int \tilde{l}(r_i,s,a_i,t_i)d\kappa_i^*(r_i|s,a_i,t_i)dQ(s)d\alpha_i^*(a_i|t_i)dH_i(t_i)\\ &= \int l(r,s,a,t)h(t)f(s|a,t)\hat{\kappa}(dr|s,a,t)dQ(s)d\hat{\alpha}(a|t)d\bar{H}(t)\\ &= \int l(r,s,a,t)\hat{\kappa}(dr|s,a,t)dP(s|a,t)d\hat{\alpha}(a|t)dH(t),\end{aligned}$$

where the third equality follows from the definition of (α_i^n,κ_i^n), the inequality follows from the lower semicontinuity (established in Theorem 8.3) of the principal's loss function in the single-agent-i model, and the second-last equality follows from (A.14). Hence the incentive compatible mechanism $(\hat{\alpha},\hat{\kappa})$ yields losses no greater than l_* for the principal, which, by the definition of l_* implies that $(\hat{\alpha},\hat{\kappa})$ is loss minimizing among all incentive compatible mechanisms. Q.E.D.

Proof of Proposition 13.6 Consider the I-player Bayesian game between the agents in which each agent i's utility as a function of the profiles of actions and types is $v_i(a,t)$ and in which types are drawn according to H. By hypothesis, $v_i(a,t)$ is jointly measurable, continuous in a for each t, and H-integrably bounded. By Balder (1988, Proposition 3.1) a Nash equilibrium exists for this Bayesian game. Let $(\alpha_i)_{i\in I}$ be this equilibrium, where each $\alpha_i : T_i \to \Delta(A_i)$ is a transition probability. Then, the multi-agent mechanism (α, κ) defined by $\alpha(\cdot|t) = \times_{i=1}^{I}\alpha_i(\cdot|t_i)$ and $\kappa(\cdot|s,a,t) = \delta_{\phi(s)}$ for every (s,a,t), is incentive compatible. Q.E.D.

References

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

Balder, E. J., (1988): "Generalized Equilibrium Results for Games with Incomplete Information," *Mathematics of Operations Research*, 13, 265-276.

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

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

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.

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

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

Chaigneau, P., A. Edmans, and D. Gottlieb (2014): "The Generalized Informativeness Principle," working paper HEC Montreal, London Business School, and Wharton Business School.

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

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

Fellingham, J. C., Young K. K., and D. P. Newman (1984): "Ex ante randomization in agency models," *The RAND Journal of Economics,* 15, 290-301.

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

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

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

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

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

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

Ke, R., and C. T. Ryan (2015): "A General Solution Method for Moral Hazard Problems," working paper, University of Chicago.

Kierkegaard, R. (2014): "A Complete Solution to the Moral Hazard Problem under the Spanning Condition," working paper, University of Guelph.

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

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

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

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

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

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

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

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

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

Page, F. H. (1994): "Dominant Strategy Mechanisms for Contract Auctions with Risk Aversion and Moral Hazard, *International Journal of Game Theory*, 23, 25-42.

Renner, P. and K. Schmedders (2015): "A Polynomial Optimization Approach to Principal-Agent Problems," *Econometrica*, 83, 729-769.

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

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

Strausz, R. (2006): "Deterministic Versus Stochastic Mechanisms in Principal–Agent Models," *Journal of Economic Theory,* 128, 306-314.

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