

University of Toronto

From the Selected Works of Michael D Ryall

Summer July 8, 2016

Contract structure for joint production: risk and ambiguity under compensatory damages

Michael D Ryall, *University of Toronto*

Rachelle C Sampson, *University of Maryland*



Available at: https://works.bepress.com/michael_ryall/13/

Contract Structure for Joint Production: Risk and Ambiguity Under Compensatory Damages

Michael D. Ryall

Rotman School of Management, University of Toronto, Toronto, Ontario M5S 3E6, Canada,
m.ryall@mikeryall.com

Rachelle C. Sampson

Robert H. Smith School of Business, University of Maryland, College Park, Maryland 20742,
rsampson@rhsmith.umd.edu

We develop a model in which the parties to a joint production project have a choice of specifying contractual performance in terms of actions or deliverables. Penalties for noncompliance are not specified; rather, they are left to the courts under the legal doctrine of compensatory damages. We analyze three scenarios of increasing uncertainty: *full information*, where implications of partner actions are known; *risk*, where implications can be probabilistically quantified; and *ambiguity*, where implications cannot be so quantified. Under full information, action requirements dominate: they always induce the maximum economic value. This dominance vanishes in the risk scenario. Under ambiguity, deliverables specifications can interact with compensatory damages to create a form of “ambiguity insurance,” where ambiguity aversion is assuaged in a way that increases the aggregate, perceived value of the project. This effect does not arise under contracts specifying action requirements. Thus, deliverables contracts may facilitate highly novel joint projects that would otherwise be foregone as a result of excessive uncertainty. Suggested empirical implications include the choice of contract clause type depending on the level of uncertainty in a joint development project, one application being the level of partner experience with interfirm collaborations.

Keywords: organizational studies; strategy; games–group decisions; technology; ambiguity; contracts

History: Received July 14, 2011; accepted October 26, 2015, by Bruno Cassiman, business strategy. Published online in *Articles in Advance* July 8, 2016.

1. Introduction

Firms often use formal contracts in the governance of their collaborative projects. In both management and economics, significant bodies of scholarly work have arisen to explain the shapes taken and purposes served by such documents. As discussed below, empirical work in the management literature documents actual variation in the structure of written contracts, including features such as performance clauses, monitoring conventions, and financial contingencies under different transactional settings. By contrast, the energy in economics has largely focused on developing theory to provide explanations of the kinds of economic problems contracts can be used to solve. Between the two disciplines lies a gap: the bundle of contractual features typically documented in management’s empirical work is difficult to map to the abstract representations of contracts commonly used in economic theory. The purpose of this paper is to forge a closer link between theory and empirics via development of a formal model that maps onto observed contractual data.

There is over a decade of management research on the link between relational capital and the overall

degree of contract detail.¹ While some authors find prior relationships lead to more detailed contracts (e.g., Poppo and Zenger 2002, Mayer and Argyres 2004, Argyres et al. 2007), others find the opposite (e.g., Gulati 1995, Gulati and Singh 1998, Uzzi 1997, Corts and Singh 2004). Investigating these different conclusions, Ryall and Sampson (2009) and Vanneste and Puranam (2010) demonstrate that measures of overall contract detail blunt the ability to observe dynamics of individual contract terms. Both of these later studies find that the relationship between experience and the inclusion of a clause varies by clause type. For example, studying technology development alliances, Ryall and Sampson (2009) find action requirements (e.g., contribution of specific, preexisting intellectual property) are more likely with increasing partner experience, yet the opposite holds for deliverable requirements (e.g., completion dates for tasks and/or specified end-of-project achievements for each firm). It is noteworthy that only

¹ “Relational capital” is usually measured as prior experience between the contracting parties.

11% of contracts in this study contain any mention of financial contingencies whatsoever.²

Learning-by-doing reasoning has been used to explain the use of contract terms in many of the earlier management studies: experience with a partner leads to learning, which lowers the marginal cost of adding clauses to contracts involving that partner, thereby inducing a positive correlation between experience and contract detail (e.g., Argyres et al. 2007).³ Alternatively, sociologists posit that trust develops between parties through experience, which reduces the need to explicate performance requirements, thereby inducing a negative correlation (e.g., Granovetter 1985). The work on reputational capital in economics is closely related. This literature finds that, similar to trust, reputation may also act to reduce the need for legal enforcement mechanisms. Radner (1981) provides one of the earliest examples in this large stream of literature, which includes refinements by Bull (1987), MacLeod and Malcolmson (1993), Baker et al. (1994), Bernheim and Whinston (1998b), Levin (2003), and Kvaløy and Olsen (2009). It is worth noting that the economics findings are based on game-theoretic models where the “contracts” are typically represented as schedules of state-contingent transfer payments.⁴

Salient empirical features of contracts shown in recent management studies include the specification of required actions without additional contingencies, the use of deliverable requirements, variation in the use of action versus deliverable requirements, correlations in this variation with partner experience, and the general rarity of specified financial transfers. The connections between these findings and the closest theoretical streams across the disciplines are loose and incomplete. The theories imply different relationships between experience and the overall number of clauses used in a contract, but they are silent on variation in the relative use of clauses across contracts. Moreover, the formal representations of contracts in the economics literature on reputation (as state-contingent transfer payments) are qualitatively different than those observed in practice, since many real-world contracts

lack financial contingencies altogether. This weak connection and our earlier empirical work motivate the present paper, which analyzes variation in the use of action versus deliverables performance terms in contracts. The variation here arises as a result of qualitatively different degrees of uncertainty regarding the joint project. Although we do not model experience per se, our uncertainty scenarios can be interpreted as arising as a result of it.

Our focus on uncertainty touches on several streams within the broader economics literature on contracts, which are too substantial to review in detail here. However, there are two general ideas from this literature that are particularly relevant to this study: moral hazard and adverse selection. In the context of contracts, moral hazard arises when actions by contracting parties are not observable or verifiable. An example of this would be when the quantity and/or quality of an input into a joint project cannot be verified by anyone other than the contributing party. Lack of observability implies that terms surrounding such actions are not externally enforceable and eliminates any direct reason to specify them in contracts. At first glance this offers a possible explanation for the use of action versus deliverable requirements in contracts: firms include action requirements when these actions are observable, but they exclude action requirements when not observable and, thus, not enforceable. However, a moral hazard approach is unable to explain observed contracts for joint technology development that sometimes include action clauses and in other cases deliverables only (Ryall and Sampson 2009). In these cases, many actions are observable, hence their inclusion in some contracts. Yet in other contracts for similar projects—i.e., in which these actions should be also observable—such clauses are not included. We attempt to explain such cases by constructing a model in which actions are always observable (i.e., our results with respect to clause selection do not depend on whether actions are hidden or not).

The literature on adverse selection, where actions and deliverables may be observed but some inherent feature of a partner is not, is closer to our approach. For example, contracts in Khalil and Lawarrée (1995) are represented as schedules of effort- or output-contingent payments, depending on which is selected for monitoring. The hidden feature in Khalil and Lawarrée (1995) is a productivity endowment. In our setup, the potential unknown is how partner actions are linked to different deliverable outcomes—essentially, managers face uncertainty about their partner’s “production technology” or type. A firm’s type determines how likely particular outcomes are for the various possible actions that the firm may take.

The essential features of our approach are as follows. We build a model that depicts two firms faced with

² Typically, these financial contingencies take the form of performance penalties (in legal terms, *specified damages*) or bonuses.

³ These ideas are in line with those of Williamson (1975), who argues that real-world complexity implies incomplete contracts, which, combined with specific relationships and opportunism, create favorable conditions for the existence of firms. Grossman and Hart (1986) begin the formalization of these ideas. A more recent example is the work of Tadelis (2002), who presents a model in which an exogenous parameter captures the likelihood that a contract sufficiently covers ex post contingencies and proceeds to analyze the contract-versus-integrate implications.

⁴ One line of work in this vein concerns the extent to which it is possible to articulate the state-contingent transfers in simple form, such as a linear price schedule (e.g., Levin 2003, Bernheim and Whinston 1998a).

a joint economic opportunity.⁵ They may agree to a formal contract if they wish. If so, they then write a contract that specifies either action or deliverable performance requirements. *Actions* are defined as managerial decision or choice variables; *deliverables* are the stochastic, payoff-relevant consequences of those decisions. Both actions and deliverables are verifiable by an outside party and thus are practical to include in a contract. To allow for contracts that do not elaborate explicit schedules of state-contingent financial transfers, we introduce the legal institution of *compensatory damages*, under which contractual violations are assessed. This doctrine is prevalent in many real-world jurisdictions and is a common remedy for contract breach. Finally, we create a setting in which the parties face different degrees of uncertainty regarding the consequences of each other's actions. Our earlier empirical investigations suggest that uncertainty of this kind is an important obstacle to the organization of many forms of economic activity. In the presence of uncertainty, firms have difficulty selecting a mutually agreeable set of action requirements, even when actions are verifiable.

Thus, our firms choose actions that, in turn, affect the probabilities on a set of possible deliverables or outcomes. We refer to the mapping from a firm's actions to probability distributions on deliverables as its *type*. Implicitly, a firm's type is determined by its assets (both real and intangible), prior experience, and investments made—that is, the capabilities of a firm that determine its ability to produce deliverables. In our model, whether a contract stipulates action or deliverable performance is driven by the extent of the uncertainty the parties face regarding each others' type. Here, when a firm knows its partner's type, it also knows how partner actions affect the likelihood of desired outcomes. We consider three fundamental cases, representing increasing levels of uncertainty. The first is *full information*, in which each firm has exact knowledge of its partner's type. (We assume that firms always know their own type.) The second is *risk*, in which firms do not know their partner's type yet are able to assign a probability to each of their partner's possible types. The last is *ambiguity*, in which the partner's type is unknown and, moreover, cannot be assessed with probabilistic precision. The sources of these differences in uncertainty are not modeled explicitly. We simply suggest that they are sufficiently broad to be relevant in a number of empirically interesting settings. For example, uncertainty could be driven by a firm's level of relational or collaborative experience, a premise

⁵ The opportunity could be anything from a buy–sell agreement to a complex technology development alliance. By “firm,” we typically mean the decision makers responsible for implementing the agreement.

that links our results to our earlier empirical work motivating this paper. Alternatively, such uncertainty might also be project related, being higher in the earlier phases of a project or with less mature technology areas.

Theoretical work on ambiguity is an active area of inquiry in economic theory.⁶ As a technical matter, we adopt the maxmin expected utility (MMEU) approach of Gilboa and Schmeidler (1989). This approach has gained traction in economics, particularly in principal–agent settings. However, despite the fact that ambiguity about partner capabilities is a common issue in joint development projects, there are presently very few papers that include ambiguity as a category of uncertainty in contracting. A notable exception is Mukerji (1998), who analyzes buyer–seller transactions with ambiguous investment holdup problems. He demonstrates (i) that ambiguity can result in inefficient and incomplete arms-length contracts and (ii) that vertical integration can result in strict efficiency improvements. Our analysis departs from Mukerji (1998) in that we examine joint production with adverse selection instead of moral hazard. More specifically, we assume that both actions and outcomes are observable and, thus, contractible, which is not the case in the moral hazard setting used in Mukerji (1998). Here, partner type is uncertain (except under full information), which, in turn, leads to uncertainty about what the best partner actions are. We also assume, in contrast to Mukerji (1998), that firms assess contract violations under the specific legal institution of compensatory damages. These appear to be reasonable assumptions given that clauses specifying both actions and deliverables are prevalent in actual contracts and that contracting parties typically do have recourse to courts for compensatory damages.

As mentioned above, one of the important novelties of our model is that it does not require contracts to include schedules of state-contingent transfer payments (in contrast to much of the extant contract theory based on asymmetric information). Rather, the penalties and

⁶ Readers interested in digging deeper may find the following papers useful: Kelsey and Quiggin (1992), Mukerji (1997), Casadesu-Masanell et al. (2000), Ghirardato and Marinacci (2002), Aragones et al. (2005), Maccheroni et al. (2006), and Gilboa et al. (2008). Klibanoff et al. (2005) provide a general framework that allows distinct preferences for risk and uncertainty. Chamberlain (2000) develops an econometric algorithm for estimating MMEU choices and applies it in an autoregressive, random-effects model for panel data. Actual applications of the theory are much sparser. Epstein and Zin (1989) provide a well-known examination of asset returns. Lopomo et al. (2014) study mechanism design under ambiguity and show conditions under which a mechanism designer can extract full information rents from an agent. More recently, Kellner (2015) examined a principal–agent problem with multiple agents under ambiguity and showed that ambiguity aversion favors tournament-style compensation over wage schemes that only depend on output levels. De Castro and Yannelis (2011) show, in a general setting, any efficient allocation is incentive compatible under MMEU preferences.

awards associated with noncompliance are left to the courts according to the *doctrine of compensatory damages*, as they often are in the real world. Under this doctrine, a party to a contract that suffers an economic loss as a result of noncompliance by its partner must receive compensation. The compensation is paid by the breaching party to the performing party in an amount that keeps the latter whole relative to its expected economic outcome under full compliance. It is important to note that this does not imply that courts always ensure that actual payoffs (of the performing party) equal expected payoffs. This is because the doctrine contains an asymmetry: compensating transfers are required when one party fails to perform *and* the compliant party receives a worse-than-expected outcome (i.e., is “damaged”). If the one party breaches but the compliant party receives a better-than-expected payoff (which is possible in our model because of the stochastic nature of deliverables), no penalty is imposed on the former. This relaxes the usual assumption that firms automatically obey their contracts due to absolute enforcement by the courts.⁷ We do assume that, at the time of signing, firms believe any court-based damages assessments will be accurate. Finally, although limited liability issues are potentially important, we do not consider them here. Limitations of liability largely constrain punitive damages, not compensatory damages and, as such, are not as critical to our analysis as under different assumptions about the operative legal institutions.

Our results make several theoretical contributions to the theory of contracts, the most novel of which are summarized below by case. (Empirical implications are suggested throughout this paper but summarized in our conclusion, Section 7.)

- *Full information*: With full knowledge, firms know their partner’s type and, thus, know the precise stochastic implications of a partner’s actions its deliverables. In this case, we show that first-best actions (i.e., those that maximize the expected value of the project) are always induced by a contract that requires them. Thus, detailed state-contingent transfers need not be specified in such settings—the legal institution of compensatory damages suffices. By contrast, deliverable-based contracts do not always guarantee the implementation of value maximizing actions. Intuitively, a firm may choose an opportunistic action that benefits its own expected payoff at the expense of the aggregate, provided that the required deliverable occur with sufficient likelihood under that action.

⁷ Hence, our *full information* case is not trivial. This modeling choice was originally intended as a realistic way to determine the financial implications of noncompliance in contracts without specified transfer payments. However, the asymmetry inherent in compensatory damages makes its overall effect on firm behavior an interesting, open theoretical question in its own right.

- *Risk*: Action requirements lose their dominance under uncertainty with respect to partner type—deliverable requirements can induce efficient behavior in situations where action requirements fail to do so. This is because firms do not know which actions are required of their partner when they are uncertain about partner type. The alternative to contracting on actions is to specify a desired outcome and leave it to each partner to figure out how best to deliver. We describe sufficient conditions under which the deliverables approach is effective.

- *Ambiguity*: Under ambiguity, the incentive effects of action versus deliverable requirements are similar to those in the case of risk. However, in contrast to the case of risk, we show that deliverables contracts may increase the ex ante perceived value of an ambiguous project beyond the value attainable under an action contract. Essentially, deliverable requirements under compensatory damages create a form of “ambiguity insurance” that is mutually value increasing for the parties involved in the deal. Not only does contracting on deliverables reduce the payoff from inefficient actions under ambiguity but it may also increase the value of efficient actions. Both effects work to increase the expected value of the project. To the best of our knowledge, we are the first to discover this connection.⁸ The implication of this insurance effect is that contracting on deliverables under compensatory damages may encourage firms to undertake valuable development projects that would have otherwise been foregone as a result of ambiguity’s negative effect on perceived value.

Returning to the empirical work that motivates our analysis, these three cases can be related to collaborative experience. *Full knowledge* can be interpreted as arising because both firms have significant experience with each other on similar projects. Each knows the implications of the other’s actions. *Risk* is consistent with an empirical setting in which firms have much experience with similar projects but not with each other. Actions have different implications depending on partner type, and experience across a large number of partners results in accurate assessments regarding the probability of one type versus another. *Ambiguity* represents situations in which the firms are very

⁸ Note that this ambiguity insurance effect is similar to one described by Bose et al. (2006). In that model, ambiguity insurance is provided to bidders as an explicit part of an auction mechanism to increase the seller’s expected payoff. The source of this effect is the same ours. What we show is that the doctrine of compensating damages can create this effect in a natural way, thereby providing a rough form of ambiguity insurance with positive welfare implications. Kotowski (2012) goes in the in the opposite direction from us by showing how terms that increase ambiguity in contracts between principals and agents relax incentive constraints and, as a result, implement actions at lower cost than otherwise possible.

inexperienced in projects of the kind under consideration. Here, firms lack the ability to even formulate the likelihood of their partner’s type. This interpretation suggests an empirical mapping from our findings to correlations between experience and contract structure (e.g., Ryall and Sampson 2009).

Finally, our paper is organized in a somewhat non-traditional way. We begin with an extended numerical example that illustrates all of the main findings of the general theoretical model that follows (i.e., instead of having a running example interdispersed throughout the formal theory to illustrate results as they are demonstrated). This permits us to sketch out the main findings before getting into the more technical details of the model. Readers should keep in mind that whereas the extended example adopts specific parameters to illustrate our points, the theoretical model is much more general. Our choice of numbers was designed to highlight all of the major findings of the model across the three cases in a single example. Readers interested in the general formulation should skip ahead to the actual model, starting in Section 3.

2. Overview and Extended Example

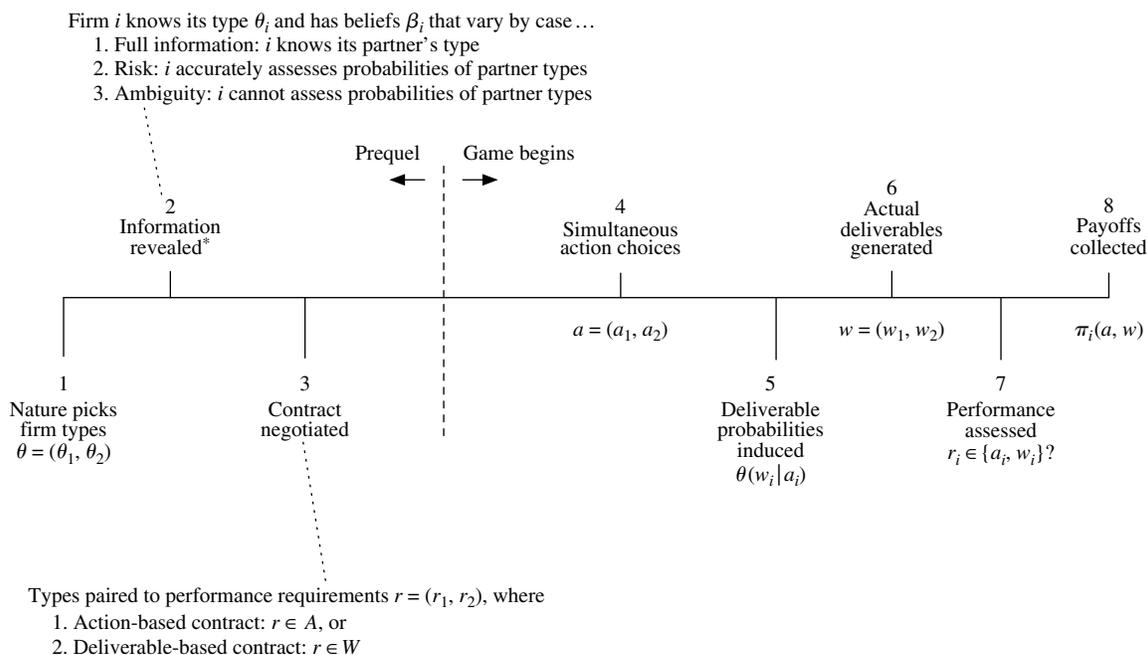
Our model, set out in detail in the next section, proceeds as follows. Two firms consider undertaking a joint production project. If they proceed, each takes actions that induce stochastic outcomes according to its type. These outcomes, which we refer to as deliverables, jointly determine the economic value of the project. Firms know their own type and begin by negotiating a contract via some process that is not modeled.

We consider two types of contracts: those specifying required actions and those specifying deliverables. With the contract set, the firms simultaneously choose their actions. Each firm’s action generates a deliverable according to its type. Once the actions are chosen and deliverables generated, firm performance is assessed via the contract. Finally, payoffs are distributed. Each firm receives its share of the realized joint economic value, less any individual cost, plus or minus any court-ordered adjustments due to performance. The timing of events is summarized in Figure 1. (The notation illustrated is elaborated in the next section.)

As mentioned earlier, both actions and deliverables are verifiable. There is no assumption of forced compliance by the courts. Rather, penalties are assessed according to compensatory damages: when one party meets its contractual obligations and the other does not, the compliant party is eligible for a damages payment from the noncompliant partner. If the actual outcome is such that the compliant party receives less than its expected payoff (under full compliance), then the breaching party must make up the difference to the complying party via a transfer payment.

Our objective is to connect contract structure—the use of deliverables versus action requirements—with uncertainty regarding partner type (i.e., the stochastic consequences of their partner’s actions). The three cases we examine are *full information*, where firms know their partner’s type with certainty; *risk*, where each firm is able to assign probabilities to candidate partner types; and *ambiguity*, where firms are unable even to formulate the likelihoods of partner types.

Figure 1 Timing of Events in the Model



Downloaded from informs.org by [129.2.19.102] on 08 July 2016, at 11:52. For personal use only, all rights reserved.

To give a preliminary sense of our findings and to provide a map of the more general model, we examine these issues in the context of a familiar example—the prisoner’s dilemma. Aside from its simplicity and familiarity, the prisoner’s dilemma is a canonical example of exactly the kind of problem contracts are expected to solve—namely, the situation is such that the best expected outcomes require firms to adopt “cooperative” strategies. However, in the absence of effective contracts, payoffs are such that the firms have individually rational incentives to “shirk” (e.g., provide low-quality inputs). Thus, in the following example, we assume that firms know the expected value of the project when the jointly efficient actions are implemented. What differs among the three scenarios we now explore is the extent of uncertainty surrounding partner type (i.e., the stochastic consequences of partner actions with respect to deliverables). We begin with the case of full information, in which the probability distribution on deliverables induced by each partner action is precisely known.

2.1. Full Information

There are two symmetric firms, firms 1 and 2, which simultaneously choose between actions **a** and **b**. These actions generate one of two possible deliverables, **x** and **y**. Actual payoffs depend on the firms’ joint deliverables as summarized in Table 1.⁹ Here, for example, “**x, y**” means firm 1 delivers **x** and firm 2 delivers **y**. These payoffs have a prisoner’s dilemma flavor: each firm prefers to generate **y** when its partner generates **x**, thereby attaining the highest individual payoff of \$37.5; however, both firms delivering **y** results in the worst payoffs (i.e., \$15, \$15), both individually and in aggregate. Keep in mind that the firms’ action choices will ultimately depend on the *expected* payoffs associated with the probabilities with which these deliverables are generated.

The mapping from actions to deliverable probabilities is summarized by firm type. In this illustration, both firms are endowed with the same type, θ . Assume the probability of deliverable **x** given action **a** is 90%: $\theta(x|a) = 90\%$ and $\theta(y|a) = 10\%$. Alternatively, assume action **b** maximizes the likelihood of **y**: $\theta(x|b) = 5\%$ and $\theta(y|b) = 95\%$. These effects are summarized in Table 2, which, in conjunction with Table 1, can be used to compute the expected payoffs for firm 1 (F_1) and firm 2 (F_2), as shown in the last two columns.

⁹The example is abstract to help readers understand the technical flow and grasp the main ideas. Payoffs such as those in Table 1 could arise on a project in which parties agree to split the aggregate value, which is high for (x, x), medium for (x, y) or (y, x), and low for (y, y), and in which deliverables costs depend on quality (low for y, high for x). The highest individual payoff, 37.5, arises under medium value and low individual cost.

Table 1 Deliverable-Contingent Payoffs

	Deliverable outcomes			
	x, x	x, y	y, x	y, y
Firm 1	25.0	14.0	37.5	15.0
Firm 2	25.0	37.5	14.0	15.0

Table 3 rearranges the expected payoffs into a conventional two-by-two format (row player payoffs are listed on the left, column player on the right). This game has the familiar prisoner’s dilemma structure—playing the inefficient action is the dominant strategy for both players.

The effectiveness of introducing a contract in such situations depends on a number of considerations, not least being the effectiveness of the available adjudication process. Typically, theorists assume that some unmodeled legal institution has the power to force parties to take the actions or produce the deliverables specified in their formal agreements. Here, we take a different approach by including an explicit legal institution—compensatory damages—as part of our model. Under compensatory damages, mitigating transfer payments are imposed from a defaulting party to a performing party when the failure of the former damages the latter. In our setting, a party is “damaged” when it receives an actual payoff below that which was expected when both parties performed. Thus, two conditions must be met for a transfer payment to be imposed: (i) one party breaches, and (ii) the other party is damaged. Whether such an arrangement provides sufficient incentive for firms to comply with a contract is a question for analysis that we examine here.

The firms would like a contract that provides mutual inducement to choose **a**, thereby maximizing the aggregate expected payoff. Therefore, suppose they simply

Table 2 Action-Induced Probabilities on Deliverables and Expected Payoffs

Actions	Deliverables (%)				E[Payoff]	
	x, x	x, y	y, x	y, y	F_1	F_2
a, a	81.0	9.0	9.0	1.0	25.0	25.0
a, b	4.5	85.5	0.5	9.5	14.7	34.7
b, a	4.5	0.5	85.5	9.5	34.7	14.7
b, b	0.3	4.8	4.8	90.3	16.0	16.0

Table 3 Expected Payoffs with No Contract

Firm 1	Firm 2	
	a	b
a	25.0, 25.0	14.7, 34.7
b	34.7, 14.7	16.0, 16.0

Notes. Equilibrium payoffs are shown in bold. See Definition 1 for a formal description of our equilibrium concept, ambiguity-averse Bayesian equilibrium.

Downloaded from informs.org by [129.2.19.102] on 08 July 2016, at 11:52. For personal use only, all rights reserved.

Table 4 Damages-Adjusted Payoffs: Action-Based Contract Specifies a, a, But F_2 Plays a, b

	Deliverables			
	x, x	x, y	y, x	y, y
Firm 1	25.0	25.0	37.5	25.0
Firm 2	25.0	26.5	14.0	5.0

Table 5 Expected Payoffs Under Action-Based Contract a, a

Firm 1	Firm 2	
	a	b
a	25.0, 25.0	25.1, 24.3
b	24.3, 25.1	16.0, 16.0

Note. Equilibrium payoffs are shown in bold.

Table 6 Damages-Adjusted Payoffs: Contract Specifies Deliverables x, x

	Deliverables			
	x, x	x, y	y, x	y, y
Firm 1	25.0	25.0	26.5	5.0
Firm 2	25.0	26.5	25.0	5.0

Table 7 Expected Payoffs Under Deliverable-Based Contract x, x

Firm 1	Firm 2	
	a	b
a	25.0, 25.0	24.1, 25.3
b	25.3, 24.1	16.0, 16.0

Note. Equilibrium payoffs are shown in bold.

sign an agreement requiring each to implement a. When both firms perform under this agreement, the expected payoff is \$25 each. Thus, under compensating damages, if F_1 performs by choosing a and F_2 breaches by choosing b, then F_2 must make F_1 whole under any combination of deliverables in which F_1 receives less than \$25. The payoffs, adjusted for damages transfers, are summarized in Table 4. For example, the courts award F_1 \$11 if the deliverables are x, y and \$10 if they are y, y, each of which brings F_1 's payoff up to \$25. Note that failure to perform is not, by itself, sufficient to trigger a penalty. Firm 2 is not required to make any payment under deliverables x, x or y, x because firm 1 is not damaged in those cases; its payoffs are equal to or greater than \$25 even though firm 2 breached. Working through all the action combinations, the expected payoffs under a contract requiring a, a are shown in Table 5. Note that we assume breach by both parties results in no damages awards.¹⁰ The important conclusion is that the contract requiring a, a (under compensatory damages) is, indeed, sufficient to induce the efficient actions.¹¹

Now, let us consider the alternative of a deliverable-based contract. Here, the intuitive choice is a contract requiring each firm to deliver x. Under such a contract, the courts ignore the actions taken by the firms and focus only on actual deliverables: a firm breaches when it delivers y. The damages-adjusted payoffs associated with deliverable outcomes are shown in Table 6. Combining these with the action-induced probabilities from Table 2, we can compute the expected payoffs

associated with each action (Table 7). Unfortunately, this contract does not induce the desired actions. Moreover, it can be shown that none of the other deliverable specifications work either.

These results continue to hold in the more general case developed below. When the stochastic implications of partner actions are well understood, action-based contracts under compensating damages can always be used to induce efficient behavior. Deliverable-based contracts are not similarly reliable. This leads, for example, to the empirical conjecture that, other things equal, extensive partner/project experience is correlated with action (versus deliverable) requirements.

2.2. Risk

The risk case represents situations in which managers are uncertain about the stochastic effects of partner actions. To see the difference with the full information case, imagine a firm considering a joint product development project with a partner. Suppose the partner is responsible for an innovation deliverable, to be generated using one of two approaches: exploration or exploitation. Full information is when the firm knows the likelihoods of partner success associated with the adoption of either approach. Risk allows for situations in which some partners have superior exploration skills, whereas others are better with exploitation. The assumption in this case is that managers, perhaps from prior experience in such settings, are able to formulate probabilities on partner types (e.g., "there is a 90% chance *this* partner has superior exploration capabilities").¹²

Return to the previous example with deliverable-contingent payoffs as elaborated in Table 1 (assume

¹⁰ In real life, mutual breach requires courts to assess "contributory negligence"—i.e., the share of one's damages resulting from one's own failure to perform. We sidestep this complication by implicitly assuming that mutual breaches cancel each other out.

¹¹ It is sufficient in the sense that the actions corresponding to the bold cells in Table 5 constitute a Nash equilibrium of the adjusted game.

¹² These are called "second-order" probabilities because they quantify uncertainty with respect to types, which, in turn, specify the deliverable probabilities induced by actions (the latter being "first-order" probabilities).

Table 8 Deliverable Probabilities Based on Firm Types and Actions Taken

Action	Firm = θ_1		Firm = θ_2	
	Deliverable		Deliverable	
	x	y	x	y
a	0.90	0.10	a	0.05
b	0.05	0.95	b	0.50

Table 9 Expected Payoffs Under the Four Firm-Type Combinations

	Partner = θ_1		Partner = θ_2	
	a	b	a	b
Firm $i = \theta_1$				
a	25.0, 25.0	14.7, 34.7	14.7, 34.7	20.2, 29.6
b	34.7, 14.7	16.0, 16.0	16.0, 16.0	25.9, 15.3
Firm $i = \theta_2$				
a	34.7, 14.7	16.0, 16.0	16.0, 16.0	25.9, 15.3
b	29.6, 20.2	15.3, 25.9	15.3, 25.9	22.9, 22.9

there is no contract at this point). Now, add a second possible type, θ_2 . The stochastic implications of the two types are presented Table 8: θ_1 is as before, and θ_2 is its mirror image. Thus, if a firm is type θ_1 , action **a** maximizes the chances of **x**. If the firm is type θ_2 , however, taking action **a** actually makes **y** more likely. Using the probabilities in Table 8 and deliverable-contingent payoffs from Table 1, we compute the expected payoffs to each firm for each of the various firm-type combinations (Table 9). Note that each of the four firm-type pairs generates its own simultaneous-move subgame. For example, the four cells in the upper left of Table 9 correspond to the game shown previously in Table 3 (in which both firms were what we are now calling type θ_1).

This scenario presents an additional complication: firms must take into account the fact that partner strategies may well vary by partner type. Other things equal, firms with different experience/capabilities, etc., have different preferences with respect to their actions. Thus, it is not simply that the consequences of partner actions vary by type, but, precisely because they do, so too may different types choose different actions under any given contract.¹³ Such heterogeneity is a realistic possibility when partners vary in their capabilities.

Our example brings this issue to the foreground in a direct way: inspecting Table 9, we see that **b** is dominant for θ_1 firms (i.e., strictly preferred regardless of partner type or choice of action), and **a** is dominant

¹³ The *risk* case is a Bayesian game in which nature assigns types to firms according to some distribution. Firms observe their own type and then select their actions. Thus, a “strategy” maps types to actions. Since firms know the distribution of partner types, they can compute expected payoffs in the context of a given partner strategy.

Table 10 Expected Payoffs Under Action-Based Contract a, a

	Partner = θ_1		Partner = θ_2	
	a	b	a	b
Firm $i = \theta_1$				
a	25.0, 25.0	25.1, 34.7	14.7, 34.7	20.5, 29.3
b	24.3, 25.1	16.0, 16.0	-2.7, 34.8	25.9, 15.3
Firm $i = \theta_2$				
a	34.7, 14.7	34.8, -2.7	16.0, 16.0	26.5, 14.8
b	29.3, 20.5	15.3, 25.9	14.8, 26.5	22.9, 22.9

for θ_2 types. Unfortunately, this is the opposite of the efficient strategy—to maximize expected aggregate payoff, type θ_1 firms should play **a** and θ_2 types **b**. Essentially, each of the four subgames contained in Table 9 is a variant on the prisoner’s dilemma. Moreover, since dominant actions do not depend on partner type or action, the overall game inherits this problem: its equilibrium is for every pair of types to play their inefficient actions.

Is it possible to remedy this problem with an action-based contract as was the case with full information? Presumably, the answer is no because type is not observed and, hence, is not a contractible contingency. Thus, specifying a particular action, say, **a**, induces the desired behavior only when the firm’s type is θ_1 . Nevertheless, let us consider a contract requiring both firms to choose **a**. The expected payoffs, contingent on each type pair, are summarized in Table 10.¹⁴

Finally, let us illustrate the calculation of the expected payoffs given an unknown partner type. If firm i is type θ_1 and anticipates partner compliance (i.e., partner plays **a** regardless of type), the expected payoffs to **a** and **b** are, respectively,

$$\text{firm } i \text{ chooses a: } 90\% \times 25.0 + 10\% \times 14.7 = \$24.3 \quad \text{and}$$

$$\text{firm } i \text{ chooses b: } 90\% \times 24.3 + 10\% \times (-2.7) = \$21.6.$$

If firm i is type θ_2 , the calculations are $90\% \times 34.7 + 10\% \times 16.0 = \32.8 and $90\% \times 29.3 + 10\% \times 14.8 = \27.8 , respectively. As expected, both firms choose **a**, regardless of type, in equilibrium under this contract (note that damages penalties are never imposed). Running through the various possibilities, it can be shown that requiring both firms to do **a** is the action-based contract that results in the highest aggregate expected payoff on an a priori basis.

¹⁴ Note that the adjusted entries in Table 10 are based on damages transfers computed using the expected payoff under **a**, **a** for the given firm-type pair and imposed at the *deliverables* level of detail. As a result, the cross-diagonal figures may not conform to immediate intuitions.

Table 11 Expected Payoffs Under Deliverable Requirements x , x ; All Firm-Type Pairs

	Partner = θ_1		Partner = θ_2	
	a	b	a	b
Firm $i = \theta_1$				
a	25.0, 25.0	24.1, 25.3	24.1, 25.3	24.6, 25.2
b	25.3, 24.1	16.0, 16.0	16.0, 16.0	21.0, 20.3
Firm $i = \theta_2$				
a	25.3, 24.1	16.0, 16.0	16.0, 16.0	21.0, 20.3
b	25.2, 24.6	20.3, 21.0	20.3, 21.0	22.9, 22.9

There is another notable point. Consider the case in which both firms are type θ_1 . Under both the risk and full information scenarios, the contract requiring the efficient action, **a**, does succeed in getting both firms to choose **a**. Under full information, the expected payoff to each firm is \$25.0, whereas under risk it is \$24.3. The difference between the two is driven by the fact that under the risk scenario, the action-based contract requiring **a** actually dictates that type θ_2 firms play their individually rational (inefficient) action. The 10% probability of this occurrence reduces the expected payoffs to both firms.¹⁵

Interestingly, a deliverable-based contract does succeed in inducing efficient behavior. Consider the contract requiring both firms to deliver x . Table 11 adjusts the payoffs in Table 9 for the deliverable-based damages transfers (as shown in Table 6). Suppose the strategy of firm i 's partner is to play **a** when its type is θ_1 and **b** when it is θ_2 . This is the efficient strategy. To evaluate the best reply for a type θ_1 firm i , we calculate the expected payoffs as follows:

$$\text{Firm } i \text{ chooses } \mathbf{a}: 90\% \times 25.0 + 10\% \times 24.6 = \$25.0,$$

$$\text{Firm } i \text{ chooses } \mathbf{b}: 90\% \times 25.3 + 10\% \times 21.0 = \$24.9.$$

Similarly, firm i 's payoffs to actions **a** and **b** when its type is θ_2 are 24.9 and 25.0, respectively. Therefore, the best reply strategy for firm i is also to take action **a** when its type is θ_1 and **b** when it is θ_2 . Thus, these strategies constitute a Bayesian Nash equilibrium in which firms of either type play efficiently—that is, they achieve the greatest possible expected payoff, which strictly exceeds the payoff under the best action-based contract in the risk scenario.

Summing up, risk represents situations where agents face well-calibrated uncertainty regarding the stochastic consequences of each others' actions. A contract

¹⁵ Note the potential for adverse selection problems to arise as a result. Managers entering into contract negotiations are well aware that their partner knows its own type. Thus, they may be rightly suspicious of, for example, a partner *claiming* to be θ_1 and arguing in favor of writing **a** into the contract as the appropriate requirement for themselves.

specifying actions under compensating damages is efficacious in the sense that it induces compliance. This carries through to the more general setting set out in our formal model below. The problem is that managers are uncertain of their partner's type, and, if efficient actions vary by type, then the possibility of specifying the wrong actions devalues the a priori expected value of any action-based agreement. Thus, deliverable requirements—which set outcome goals without taking a stand on how they must be achieved—may very well succeed where action requirements fall short. This is a reversal of the full information result.

2.3. Ambiguity

At least since Knight (1921), a distinction has been made between two types of uncertain situations—those in which agents are able to quantify their uncertainty in terms of probabilities and those in which they are unable to do so. For example, most people know that the probability of rolling a five on a balanced, six-sided die is one in six. However, when managers are asked to assess the probability that a particular research and development project will result in a patent, a common response is to identify a range of probabilities (e.g., “between 5% and 15%”). The inability to make a precise assessment is likely due to some combination of limited individual skill, imperfections in human cognitive capacity, and a lack of information and experience. Our ambiguity case is meant to capture situations in which this kind of ambiguity arises with respect to partner type.

The imprecise probability assessments under ambiguity are assumed to be irreducible: managers are unable to quantify their uncertainty in the form of a single, second-order probability distribution on partner types (as was the case in the risk scenario). Rather, they are only able to narrow their uncertainty down to a set of possible such distributions. This raises a technical problem: a firm's preferred action now depends on which distribution, from the set of possibilities, is used to resolve the ambiguity. The solution is provided by the fact that ambiguity aversion appears to be a salient psychological trait shared by many, if not most, decision makers. Apparently, humans tend to penalize choices with ambiguous consequences—the greater the ambiguity (i.e., the larger the set of second-order possibilities), the greater the penalty.¹⁶

This can be operationalized in our setting by assuming that firms evaluate each of their own actions using a distribution on partner types (from the set thought possible) that delivers the minimum expected payoff for that action. Then, the action that yields the greatest expected payoff (so evaluated) is chosen. In

¹⁶ See Ellsberg (1961). This finding has been replicated in a variety of settings.

essence, ambiguity-averse managers are subjective expected payoff maximizers who happen to penalize ambiguous choices by associating them with the most pessimistic evaluations possible. This approach, which we adopt below, is called the *maxmin expected utility* formalization (Gilboa and Schmeidler 1989). Note the implication that the probability distribution being used to evaluate an action’s expected payoff is only correct when the partner’s true type happens to be the one that generates the worst expected payoff (for the focal action). Accordingly, we refer to these evaluations as “subjective” expected payoffs below because they are, typically, not equal to the true (or objective) expected payoffs.

To proceed, assume that firms face extreme ambiguity in the sense that they simply have no idea what their partner type is—*any* distribution on partner types is considered possible. Under this assumption, the assessment of subjective expected payoff consequences is fairly straightforward: for each of the firm’s action choices, it assumes that its partner is the type that yields the firm the lowest expected payoff. The firm then selects its action based on these assessments. Continue to assume the firm types are those given in Table 8, which result in the type-contingent expected payoffs as shown in Table 9. The problem facing managers is that they only know that the true (or, objective) probability that their partner’s type is θ_1 is something between zero and one. As in the previous cases, the equilibria we consider are Nash-like in the sense that firms best reply to each others’ actions. The difference here is that the “best reply” is calculated in terms of subjective expected payoffs, which reflect the ambiguity aversion discount.¹⁷

For example, suppose F_1 is type θ_1 and we want to identify its ambiguity-averse best reply to an objectively efficient partner strategy (again, this is the strategy that θ_1 and θ_2 types choose actions **a** and **b**, respectively). Using the payoffs in Table 9, F_1 ’s ambiguity-averse expected payoff to choosing **a** is $\min\{25.0, 20.2\} = 20.2$; that is, if its partner is type θ_1 and chooses **a**, then F_1 also choosing **a** results in expected payoff of 25.0. However, if its partner is θ_2 and chooses **b**, then F_1 ’s selection of **a** results in an

¹⁷ We elaborate the epistemological consistency conditions required for equilibria in our model below. Firms face ambiguity regarding partner type and know their partners do as well but understand that both parties are ambiguity-averse profit maximizers. In equilibrium, this knowledge leads to strategies that are ambiguity-averse best replies. Even if managers cannot accurately estimate the probability of partner type and, thus, objectively optimal actions, they know their counterparts are as rational as they are. Thus, the firms may grasp, for example, that a contract specifying **a**, **a** causes their partner to choose **a** regardless of firm type and, given this understanding, resolve the consequences of their own choice of action assuming the partner type that minimizes their payoff in this context. Our equilibrium notion is “boundedly rational” in this sense.

Table 12 Ambiguity-Averse Expected Payoffs by Type in the No-Contract Equilibrium

	$F_2 = \theta_1$		$F_2 = \theta_2$	
	<i>a</i>	<i>b</i>	<i>a</i>	<i>b</i>
$F_1 = \theta_1$				
<i>a</i>	—, —	14.7, —	14.7, —	—, —
b	—, 14.7	16.0, 16.0	16.0, 16.0	—, 15.3
$F_1 = \theta_2$				
a	—, 14.7	16.0, 16.0	16.0, 16.0	—, 15.3
<i>b</i>	—, —	15.3, —	15.3, —	—, —

Note. Equilibrium payoffs are shown in bold.

expected payoff of 20.2. Given the extreme ambiguity facing F_1 , it computes the subjective expected payoff to choosing **a** simply by associating it with the worst partner type (i.e., F_2 is θ_2)—i.e., a value of 20.2.¹⁸ Using similar reasoning, F_1 ’s ambiguity-averse payoff to **b** is $\min\{34.7, 25.9\} = 25.9$. Therefore, if F_1 is θ_1 , it chooses **b**. This is the objectively inefficient action but, nevertheless, the one that maximizes its subjective expected payoff.

It is, consequently, not possible for firms to implement efficient strategies in equilibrium in the no-contract scenario. Here, the only equilibrium is the objectively inefficient one: firms with θ_1 choose **b**, and those with θ_2 choose **a**. The ambiguity-averse equilibrium payoffs for all type combinations are shown in Table 12. To see what is going on, note that all the relevant information available to F_1 at the time it must take an action is summarized in the third column of the table (under “ $F_2 = \theta_1$ ” and “**b**”). If F_1 is a type θ_1 , it computes its subjective expected payoff to **a** as 14.7 and to **b** as 16.0. If F_1 is a θ_2 , it calculates the payoffs to be 16.0 and 15.3, respectively. Thus, if F_1 is θ_1 , it does **b**, and if it is θ_2 , it does **a**.¹⁹ The payoffs in Table 12 are those in Table 9, adjusted for ambiguity aversion (i.e., minimized over partner types). Since the firms are symmetric, play of the inefficient strategy is an equilibrium.²⁰ The important highlight here is that the aggregate subjective expected value of the partnership is now only 32.0.²¹

¹⁸ The “worst” partner type is the one taken in the context of F_1 ’s own type, a focal choice of action by F_1 , and a focal strategy by F_2 (a mapping from types to actions; in this case, the objectively efficient one).

¹⁹ These payoffs are simply repeated in the fifth column of Table 12 (under “ $F_2 = \theta_2$ ” and “**a**”) because F_1 ’s calculations cannot incorporate F_2 ’s unobserved type.

²⁰ Again, in equilibrium, each firm’s subjective expectations must be consistent with the equilibrium strategy of its partner ($\theta_1 \rightarrow \mathbf{b}$, $\theta_2 \rightarrow \mathbf{a}$). This is why partner play of ($\theta_1 \rightarrow \mathbf{b}$, $\theta_2 \rightarrow \mathbf{a}$) is presupposed in the preceding calculations of each firm’s ambiguity-averse payoffs.

²¹ As an interesting aside, this equilibrium is “subjectively efficient” in the sense that 32.0 is the maximum aggregate subjective expected

Table 13 Ambiguity Payoffs Under Action-Based Contract Requiring a, a

	$F_2 = \theta_1$		$F_2 = \theta_2$	
	a	b	a	b
$F_1 = \theta_1$				
a	14.7, 14.7	—, -2.7	14.7, 16.0	—, 14.8
b	-2.7, —	—, —	-2.7, —	—, —
$F_1 = \theta_2$				
a	16.0, 14.7	—, -2.7	16.0, 16.0	—, 14.8
b	14.8, —	—, —	14.8, —	—, —

Note. Equilibrium payoffs are shown in bold.

Considering whether an action-based contract can correct this problem, we evaluate a contract requiring both firms to play a.²² The objective expected payoffs under this contract (that is, payoffs absent ambiguity) are those shown in Table 10. From these payoffs, we construct the ambiguity-averse expected payoffs in Table 13, evaluated under the contracted actions a, a and in the presence of compensating damages. Table 13 demonstrates that, as has been the case throughout, compensating damages are effective at inducing compliance with an action-based contract. Comparing outcomes under this contract to those without any contract (Table 12), a striking result emerges: θ_1 -type firms actually prefer no contract (equilibrium payoffs of 16.0) to the action-based contract (equilibrium payoffs of 14.7); θ_2 -type firms are indifferent. Ambiguity aversion amplifies the problem, observed earlier in the risk case, that there is some chance the partner is the type for which the contractually required action is the most damaging to the firm. When contract compliance yields these undesirable effects, action-based contracts may actually be spurned by the collaborating firms.

Turning to a deliverables contract, consider the one requiring outcome x from both firms.²³ We construct the ambiguity-averse expected payoffs in Table 14 via the objective expected payoffs under the deliverables contract shown in Table 11. Once again, consider the objectively efficient strategy: θ_1 types choose a and θ_2 types choose b. We use the same subjective evaluation procedure as before; i.e., given the strategy of its partner ($\theta_1 \rightarrow a, \theta_2 \rightarrow b$), F_1 assesses the payoff to a as $\min\{25.0, 24.6\} = 24.6$. Under this contract, the efficient strategies do constitute an equilibrium. This is the same behavior as in the risk case and works under a similar

payoff possible without a contract. This is consistent with the findings of De Castro and Yannelis (2011), which show that ambiguity-averse mechanism design problems often have this feature.

²² We illustrate this version of the action-based contract because, as was the case in the risk scenario, this happens to be the most efficient contract (now, in terms of subjective expected payoffs).

²³ This is the best-case deliverables contract (which is by design to simplify cross-case comparisons).

Table 14 Ambiguity Payoffs Under Deliverable-Based Contract Requiring x, x: Both Firms Choose Efficient Strategy

	$F_2 = \theta_1$		$F_2 = \theta_2$	
	a	b	a	b
$F_1 = \theta_1$				
a	24.6, 24.6	—, 21.0	—, 21.0	24.6, 22.9
b	21.0, —	—, —	—, —	21.0, —
$F_1 = \theta_2$				
a	21.0, —	—, —	—, —	21.0, —
b	22.9, 24.6	—, 21.0	—, 21.0	22.9, 22.9

Note. Equilibrium payoffs are shown in bold.

logic, in which compensating damages sufficiently curb the individual incentive to take the inefficient actions.

Comparing Tables 12–14 shows that the expected payoffs for both firms are substantially higher under the deliverable-based contract than either no contract or the best action-based contract. This is due to two effects that can be isolated in our model: (i) gains in subjective expected payoff resulting from the shift to efficient strategies as a result of incentive alignment and (ii) additional gains resulting from a form of mutual “ambiguity insurance” induced by the pattern of transfer payments arising under compensating damages. The latter, ambiguity insurance effect, a discovery that is novel to this paper and elaborated on in greater detail below stems from the ability of compensating damages to mitigate the worst-case deliverable outcomes when deliverables are the items being contracted upon. Even if partner type is the one that maximizes the probability of the “bad” deliverable, the harm that occurs when that deliverable arises is lessened by a transfer from the partner to the firm and, thus, the appearance of the ambiguity insurance effect. As we prove later, action-based contracts cannot achieve this effect. The implications are significant—not only may deliverable-based contracts provide the best incentives as they do under risk, but, by raising the absolute level of subjective expected payoff a priori, these contracts may also lead firms to undertake objectively productive projects that would otherwise be abandoned as a result of ambiguity aversion.

To better illustrate why deliverable requirements lead to the most efficient outcome under ambiguity aversion, consider the situation from the perspective of F_1 . Without a contract, F_1 anticipates its partner playing the inefficient, dominant strategy ($\theta_1 \rightarrow b, \theta_2 \rightarrow a$). Under these assumptions, F_1 's ambiguity-averse payoff for taking action a is 14.7, assuming F_1 is type θ_1 . F_1 's payoff to each action-type combination is listed in the first column of Table 15. The second column of Table 15 displays the payoffs if, instead, F_2 adopts the objectively efficient strategy (i.e., $\theta_1 \rightarrow a, \theta_2 \rightarrow b$). These payoffs are taken from those in Table 9 by finding the ambiguity-averse payoff to F_1 's actions

Table 15 Ambiguity Payoffs to F_1 with No Contract; F_2 Plays Inefficient vs. Efficient Strategies

	F_2 Strategy		Change (%)
	Inefficient: $\theta_1 \rightarrow \mathbf{b}, \theta_2 \rightarrow \mathbf{a}$	Efficient: $\theta_1 \rightarrow \mathbf{a}, \theta_2 \rightarrow \mathbf{b}$	
$F_1 = \theta_1$			
a	14.7	20.2	43.3
b	16.0	25.9	61.9
$F_1 = \theta_2$			
a	16.0	25.9	61.9
b	15.3	22.9	49.7

Table 16 Ambiguity Payoffs to F_1 Assuming F_2 Plays Efficient Strategies

	F_2 plays efficient strategy ($\theta_1 \rightarrow \mathbf{a}, \theta_2 \rightarrow \mathbf{b}$)		Change (%)
	No contract	Require \mathbf{x}, \mathbf{x}	
$F_1 = \theta_1$			
a	20.2	24.6	21.8
b	25.9	21.0	-18.9
$F_1 = \theta_2$			
a	25.9	21.0	-18.9
b	22.9	22.9	0.0

given F_2 's efficient play. For example, if F_1 is type θ_1 and chooses action **a**, its payoff is $\min\{25.0, 20.2\} = 20.2$.²⁴ Thus, getting its partner to play the efficient strategy increases a type θ_1 F_1 's payoff from playing **a** from 14.7 to 20.2—an increase of 43.3%. This improvement is solely attributable to getting the partner to play efficiently.

It is possible, however, for F_1 to do better by adopting a deliverables contract, as demonstrated earlier. Recall that the penalties associated with breaching this contract are sufficient to ensure that both firms adopt the efficient strategy. The pure effects of incentive alignment were just shown. Comparing these with the subjective expected payoffs associated with the deliverables contract (Table 14) yields the residual, ambiguity insurance effect. This effect is summarized in Table 16. Thus, we see that compensating damages can have positive effects in addition to those accruing to incentive alignment. For example, F_1 's payoff for taking action **a** as a θ_1 type increases by 21.8% above and beyond the incentive alignment effect. In this sense, contracts with compensating damages under ambiguity aversion work not only to incentivize efficient behavior but also to increase subjective expected payoffs by reducing the ambiguity aversion discount.

The ambiguity insurance effect seems similar to insurance effects resulting from risk aversion; that

²⁴ F_1 's payoff to **a** is 25.0 when F_2 is type θ_1 and does **a**; F_1 's payoff to **a** is 20.2 when F_2 is type θ_2 and does **b**. Ambiguity aversion leads F_1 to assume that F_2 is type θ_2 .

is, by smoothing payoffs over deliverable outcomes, compensating damages would also raise the expected value of a project to risk averse firms. This similarity between a more traditional insurance effect and our ambiguity insurance effect based on MMEU preferences is only superficial, however. First, our model setup does not require firms to be risk averse; hence, the outcomes discussed above are not due to payoff smoothing per se. Rather, the ambiguity insurance improvement depends on the probabilities on deliverable outcomes relative to partner type. If types are equivalent, such that the deliverable probabilities are identical, then a deliverables contract with nontrivial smoothing as a result of compensating damages transfers would induce no insurance effect on subjective expected payoffs. However, the same would *not* be true under risk aversion: payoff smoothing would increase a project's expected value even under type equivalence. Thus, risk and ambiguity aversion can lead to very different payoffs, expected values, and, hence, firm behavior.

3. The General Model

3.1. Setup

Our notational conventions are as follows. Sets are represented by capital letters; elements of sets and functions are represented by lowercase letters. Unless indicated otherwise, all sets are finite. When firm i is understood from the context, a “ $-i$ ” subscript refers to its partner. For $i \in \{1, 2\}$, firm i chooses from a set of $m \geq 2$ feasible actions, A_i , with typical element a_i .²⁵ Each a_i should be interpreted as a collection of activities, processes, resource allocations, etc., to be conducted by i on behalf of the project. An *action profile* is a list of actions $a \equiv (a_1, a_2)$, with $A \equiv A_1 \times A_2$ denoting the set of all such profiles. Similarly, W_i is firm i 's set of $n \geq m$ deliverables. A *deliverables profile* is a pair $w = (w_1, w_2) \in W$, where $W \equiv W_1 \times W_2$.

The joint project is described by a *value function* v , where $v(w)$ is the joint economic value generated when the deliverables profile is w , and a *cost profile* $c \equiv (c_1, c_2)$, where $c_i(w)$ is the cost to firm i when the deliverable profile is w .²⁶ In the absence of a contract, firms split the joint value evenly: the *payoff to i given w* is $u_i(w) \equiv \frac{1}{2}v(w) - c_i(w)$. To avoid trivial cases, we require that neither the u_i 's nor the sum of the u_i 's be constant.

The set of types associated with firms capable of taking on the project role of firm i is denoted by Θ_i and

²⁵ Our results do not depend on the assumption that both firms have the same number of actions.

²⁶ Cost $c_i(w)$ may or may not depend only on w_i . Since the w_i 's can be multidimensional themselves, they may include components to influence v and c_i so that actions have distinct effects on payoffs and costs, respectively.

contains at least two elements. The set of *type profiles* is $\Theta \equiv \Theta_1 \times \Theta_2$ with typical element $\theta \equiv (\theta_1, \theta_2)$. The *distribution of types within* Θ_i , denoted by ρ_i , is strictly positive. Profile θ is selected at the start by *nature's probability*, $\rho(\theta) \equiv \rho_1(\theta_1)\rho_2(\theta_2)$.

Types map actions to probability distributions over deliverables. Specifically, if firm i chooses a_i , then deliverable w_i occurs with probability $\theta_i(w_i | a_i)$. The *joint probability* of $w = (w_1, w_2)$ given $a = (a_1, a_2)$ and $\theta = (\theta_1, \theta_2)$ is $p(w | a, \theta) \equiv \theta_1(w_1 | a_1)\theta_2(w_2 | a_2)$.²⁷ The expected value of u_i given a and θ is thus $\mathbb{E}(u_i | a, \theta) = \sum_w u_i(w)p(w | a, \theta)$. Assume partner types matter in that, for any action profile, there is at least one pair of partner types that result in different expected payoffs.²⁸

Contracts are defined by a *performance requirements profile* $r \equiv (r_1, r_2)$, where $r \in A \cup W$. If $r \in A$, the contract is classified as *action based*; if $r \in W$, it is *deliverable based*. If a is the action profile played and w is the resulting deliverable outcome, then firm i *performs* under r if and only if $r_i \in \{a_i, w_i\}$. Otherwise, i *breaches* under r .

Suppose $r \in W$ and the deliverable outcome, w , is such that i performs and its partner breaches. Then, i receives a *damages* payment of $d_i \equiv \max\{u_i(r) - u_i(w), 0\} \equiv -d_{-i}$.²⁹ Alternatively, suppose $r \in A$. If the chosen action profile a is such that i performs when its partner breaches and the deliverable outcome is w , then $d_i \equiv \max\{\mathbb{E}(u_i | r, \theta) - u_i(w), 0\} \equiv -d_{-i}$; that is, the damages benchmark is the expected payoff given full performance of required actions. If both parties perform or both breach, $d_i \equiv -d_{-i} \equiv 0$. Under the *null contract*, denoted by r_\emptyset , damages are always zero.

From the preceding, d_i is a well-defined function of the actions taken, a ; the resulting deliverables, w ; the types, θ ; and the contract, r . Since r is always clear from the context, we write $d_i(w | a, \theta)$. Note that damages are balanced ($d_1 = -d_2$). This allows us to define firm i 's *net payoff* as $\pi_i(w | a, \theta) \equiv u_i(w) + d_i(w | a, \theta)$.

At the point of choosing actions, firms know their own type and the terms of their contract. They may or may not know their partner's type. In cases of uncertainty, let $b_i(\theta_{-i})$ denote i 's *belief* that its partner's type is θ_{-i} . Firm i is initially endowed with a *set* of prior beliefs, B_i . We consider three cases. With *full information*, types are known: given θ , $B_i = \{\theta_{-i}\}$, where $b_i(\theta_{-i}) = 1$. In the *risk* scenario, firms know the true distribution of types: $B_i = \{\rho_{-i}\}$, where $b_i = \rho_{-i}$. Under *ambiguity*,

firms do not know the distribution of types, and, consequently, no beliefs can be ruled out: $B_i = \Delta(\Theta_{-i})$.³⁰

Firms develop their beliefs prior to negotiating their contracts. Thus, contracts may vary with θ , but not beyond the limitations implied by the information available in the case at hand.³¹ We abuse notation and write $r(\theta)$ to indicate the requirements agreed to by firms whose type profile is θ . We skip the complexities of modeling the negotiation process explicitly and simply assume r (the function) is set as an exogenous feature of the game.

Thus, a firm's chosen action may depend on the contract, its type, and what it believes about its partner's type. Since the contract and beliefs are fixed for each case, let $s_i(\theta) = a_i$ indicate that firm i plays a_i when the type profile is θ . Note that, under risk and ambiguity, s_i must be constant on the θ_{-i} s (because i does not observe them). In those cases, we write $s_i(\theta_i)$ to emphasize this fact. The set of all strategies for firm i is denoted by S_i .³² The set of all *strategy profiles* is $S \equiv S_1 \times S_2$ with typical element $s \equiv (s_1, s_2)$.

3.2. Equilibrium

Our equilibrium concept is Nash adjusted to introduce ambiguity aversion to the model, using maximin expected utilities (Gilboa and Schmeidler 1989). Faced with multiple prior beliefs, $B_i = \Delta(\Theta_{-i})$, firm i applies an "ambiguity discount" by assessing expected payoffs using the minimizing beliefs in B_i . Specifically, given a type profile θ and with r understood from the context, define firm i 's *ambiguity-averse expected payoff* under s as

$$\mathbb{E}_i(s | \theta) \equiv \min_{b_i \in B_i} \sum_{\theta' \in \{\theta_i\} \times \Theta_{-i}} \sum_{w \in W} \pi_i(w | s, \theta') p(w | s, \theta') b_i(\theta'_{-i}), \quad (1)$$

where we simplify notation by writing, e.g., $\pi_i(w | s, \theta')$ rather than $\pi_i(w | s(\theta'), \theta')$. Under the first summation in (1), the θ' 's all take the form $\theta' = (\theta_i, \theta'_{-i})$, varying only in the second component.³³

DEFINITION 1 (AMBIGUITY-AVERSE BAYESIAN EQUILIBRIUM). Given r and B_i , $i \in \{1, 2\}$, a strategy profile s is an *ambiguity-averse Bayesian equilibrium* (AABE) if, for $i \in \{1, 2\}$, all $\theta \in \Theta$ and all $s'_i \in S_i$, $\mathbb{E}_i(s_i, s_{-i} | \theta) \geq \mathbb{E}_i(s'_i, s_{-i} | \theta)$. If s is an AABE under requirements mapping r , we say that r *sustains* s .

³⁰ Here, $\Delta(\cdot)$ indicates the set of all probability distributions on a set. This is an extreme form of ambiguity.

³¹ That is, r cannot vary in a way that permits a firm to refine its beliefs. This is a consistency issue: if the partner type is not known at the time of negotiation, it is difficult to imagine how such contractual variation would arise in the first place.

³² Our attention is largely limited to pure strategies. We note when mixed strategies are required for a result.

³³ The i subscript in (1) indicates ambiguity-averse expectations from the perspective of firm i . This is a generalization: ambiguity aversion arises only if B_i has multiple priors; otherwise, \mathbb{E}_i is equal to the expected net payoff, computed in the usual way.

²⁷ Although the stochastic effects of actions are independent, deliverables interactions can arise in v and the c_i 's.

²⁸ That is, for $i \in \{1, 2\}$ and all $a \in A$, $\theta \in \Theta$, there exist θ'_{-i} such that $\mathbb{E}(u_i | a, \theta) \neq \mathbb{E}(u_i | a, (\theta_i, \theta'_{-i}))$.

²⁹ When $u_i(r) - u_i(w) \leq 0$, firm i is not "damaged," and no payment is awarded even if $-i$ breached. Note also that this specification implicitly ignores limited liability issues. This is not without substance. For example, under limited liability, first-best results such as Proposition 1 would require additional conditions to be true.

3.3. Efficiency

Firms prefer contracts that induce larger over smaller aggregate expected payoffs since the former allow both firms to achieve strictly greater expected payoffs (if necessary, via lump-sum transfers at contract signing). The finiteness of A implies that, for every θ , there exists an *efficient action profile*

$$a_\theta^* \in \arg \max_{a \in A} \mathbb{E}(v - c_1 - c_2 | a, \theta). \quad (2)$$

Here, a_θ^* indicates an arbitrary efficient action profile when there is more than one. A *first-best strategy profile* is denoted s^* , where $s^*(\theta) = a_\theta^*$ for all θ . Assume, for both firms and all θ , $\mathbb{E}(u_i | a_\theta^*, \theta) > 0$; that is, efficiently implemented, the project is economically sensible for both parties (it satisfies their individual participation constraints).

4. Results for Full Information

When actions are verifiable, we know that the efficient action profile can be sustained by an action-based contract that imposes sufficient penalties on breachers. However, our first proposition shows that contractually specified penalties are not necessary in the presence of compensatory damages, as long as firms know each other's type. For this section, assume full information: given θ , $B_i = \{b_i\}$, where $b_i(\theta_{-i}) = 1$ for $i \in \{1, 2\}$.

PROPOSITION 1. *The requirements mapping $r(\theta) = a_\theta^*$ for all $\theta \in \Theta$ sustains the first-best strategy s^* .*

Thus, action-based contracts under full information and compensatory damages are sufficient to sustain efficient behavior. Damages guarantee an expected payoff to the compliant firm of at least its expected share of the efficient payoff. Consequently, any expected gains to the deviating firm are erased through court-imposed damages.

This explains a frequent observation in the management literature: firms with substantial experience with one another on similar projects often write contracts that elaborate detailed action plans but do not specify penalties for unfaithful implementation. Research suggests that negotiating such contracts serves an important planning role (e.g., Argyres and Mayer 2007). Proposition 1 demonstrates that it is precisely because these contractualized action plans have legal ramifications (often embodied as a “statement of work”) that they may also induce efficient outcomes.

When firms understand the implications of partner type in the context of the project, action-based contracts are sufficient to induce efficient actions under compensatory damages. The following proposition says that compensatory damages are also necessary, in the sense that any guarantee below that ensured by compensatory damages may cause Proposition 1

to fail. To see this, given some $\varepsilon > 0$, suppose the courts award firm i the following payment when it performs under an action-based contract and its partner breaches: $d_i^\varepsilon \equiv \max\{\mathbb{E}(u_i | r, \cdot) - u_i(w) - \varepsilon, 0\} \equiv -d_{-i}^\varepsilon$, i.e., a guarantee of ε less than compensatory damages. For all other cases, keep $d_i^\varepsilon \equiv d_i$. Let $\pi_i^\varepsilon \equiv u_i + d_i^\varepsilon$.

PROPOSITION 2. *For all $\varepsilon > 0$, there exist type profiles θ and projects (v, c) such that, if π_i is replaced by π_i^ε in (1), then the requirements mapping $r(\theta) = a_\theta^*$ for all $\theta \in \Theta$ do not sustain the first-best strategy $s^*(\theta) = a_\theta^*$ for all $\theta \in \Theta$.*

Proposition 1 implies that deliverable-based contracts can never *outperform* action contracts under full information. Are they always at least as good? The answer is no. We demonstrate this by, first, demonstrating how our adverse selection model under full information maps to a well-known partnership model with moral hazard found in the economics literature. Then, by citing an extent result on moral hazard, we establish conditions under which deliverable-based contracts in our model cannot induce efficient behavior.

The following partnership model is from Legros and Matsushima (1991; hereafter LM91). A finite set of partners is indexed by $\bar{N} \equiv \{1, \dots, \bar{n}\}$.³⁴ Agent i 's feasible actions are given by a finite set \bar{A}_i . Each joint action profile $\bar{a} \in \bar{A} \equiv \times_i \bar{A}_i$ induces a probability distribution $\bar{p}(\cdot | \bar{a})$ on a finite set of outcomes \bar{W} . The payoff to i is $\bar{\pi}_i(\bar{w}, \bar{a}) \equiv \bar{t}_i(\bar{w}) + \bar{u}_i(\bar{w}, \bar{a})$, where \bar{t}_i is the i th component of a balanced transfer rule \bar{t} , and \bar{u}_i is i 's utility function. A *mixed strategy* for partner i is a distribution $\bar{\sigma}_i \in \bar{\Sigma}_i \equiv \Delta(\bar{A}_i)$; the set of all such profiles is $\bar{\Sigma} \equiv \times_i \bar{\Sigma}_i$. In this game, the partners simultaneously choose their actions; these induce a distribution over outcomes; the realized outcome determines final payoffs. The “contract” is the transfer rule \bar{t} . Unverifiable actions are implicit in the restriction of transfers to outcome contingencies. The following theorem applies to games of this kind.

THEOREM 1 (LM91). *There exists a balanced transfer rule \bar{t} such that an efficient pure strategy profile \bar{a}^* is a Nash equilibrium of the resulting game if, and only if, $\gamma^* < +\infty$, where γ^* is an index implied by $\bar{\Sigma}$ (see LM91 for a precise definition).³⁵*

Consider a specific instance of our game in the full information case under a deliverable-based contract r . Begin with an arbitrary θ selected by nature. This subgame is a special instance of LM91. To see this,

³⁴ We use the bar to decorate the objects in the LM91 partnership game and thereby distinguish them from ours.

³⁵ Loosely put, Theorem 1 says that when the strategies available to partners are “too close” in the sense of the probability distributions induced by deviations from efficiency, then partners can mimic each others' deviations in ways that make it impossible to solve the team incentive problem via a balanced transfer scheme based on outcomes.

set $\bar{n} = 2$, $\bar{A} = A$, and $\bar{W} = W$. For firm i , $\bar{A}_i = A_i$, $\bar{u}_i(\cdot) = u_i(\cdot)$, and $\bar{t}_i(\cdot) = d_i(\cdot | \cdot, \theta)$. Let $\bar{p}(\cdot | \cdot) = p(\cdot | \cdot, \theta)$. Last, for any pure strategy profile s in our game, set $\bar{a}_i = s_i(\theta)$. Thus, there is a one-to-one correspondence between pure strategy Nash equilibria of this instance of the LM91 game and any θ subgame in our model.

Following this logic, if the necessary and sufficient condition in Theorem 1 is not met for the LM91 game implied by one of the θ subgames in a particular instance of our model, then none of the potential deliverable requirements for the firm-type pair θ is capable of inducing efficient action choices. Since efficiency fails for at least one pair, it fails for the game overall (i.e., there is no requirements profile r that sustains an efficient strategy). This demonstrates that in situations where firms understand the consequences of each other's actions, deliverable-based contracts are weakly dominated by action-based contracts. This is consistent with our intuition from the moral hazard literature: if firms can base penalty schemes on actions, first-best actions can be induced; if not, then outcome-based transfers are sometimes (but not always) efficacious. We extend this intuition to the institution of compensatory damages.³⁶

5. Results for Risk

In the case of risk, partner type is unknown, and beliefs are consistent with the true type-generating distributions: $B_i = \{\rho_{-i}\}$. Knowledge of ρ_{-i} can be interpreted as firm i being well calibrated as a result of substantial experience with projects of the present kind across a large population of partners (but not firm $-i$ specifically).

We begin by illustrating a substantive obstacle to achieving efficiency under risk. In the present case, each firm's belief is equal to the true distribution of partner types. Because this distribution is strictly positive, no partner types can be ruled out. This imposes a restriction: $s_i(\theta_i, \theta_{-i})$ must be constant on the θ_{-i} 's.³⁷ For this case to make sense, we must also assume that i 's lack of knowledge about its partner's type is not contradicted by the contract. Therefore, we assume that $r(\cdot)$ is also constant (on the θ_{-i} 's). Otherwise, i could refine its knowledge of its partner's type based on the terms of the contract—a situation that would make little sense. Since this is true for both firms, $r(\cdot)$

is constant; the contract terms do not vary by type profile.

What this means is that, under uncertainty, the only time first-best efficiency has any hope of attainment is when a firm's first-best actions depend only on its own type and not partner type. Formally, a project is said to exhibit *first-best action independence* if for $i \in \{1, 2\}$ and all pairs of firm-type profiles $\theta = (\theta_i, \theta_{-i})$ and $\theta' = (\theta'_i, \theta'_{-i})$, $\theta_i = \theta'_i$ implies that there exist $a_{\theta}^* = (a_i, a_{-i})$ and $a_{\theta'}^* = (a'_i, a'_{-i})$ such that $a_i = a'_i$. This leads to the following proposition.

PROPOSITION 3. *Assume first-best action independence does not hold. Let $\bar{r} \in A \cup W$ be an arbitrary requirements profile and assume $r(\theta) = \bar{r}$ for all $\theta \in \Theta$. If s is sustained by r , then $s \neq s^*$.*

As is true in so many settings, Proposition 3 shows that the presence of uncertainty may well bound firms away from first-best strategies. It is worth emphasizing that the issue is that limited information forces a firm's action choice to vary only with its *own* type. Failing action independence, the best that can be hoped for in the present context is a strategy profile that solves

$$\max_{s \in S} \sum_{\theta \in \Theta} \mathbb{E}(v - c_1 - c_2 | s, \theta) \rho(\theta). \quad (3)$$

With unknown partner type, firms can condition on their own type but not on that of their partner. This is implicit in (3), which maximizes over strategies that conform to the game's information structure. In some cases, (3) is equal to the expected amount under the first-best strategy s^* , but not in general; e.g., Proposition (3) demonstrates that first-best action independence is a necessary condition.

We return our attention to contracts, starting with action requirements. Unfortunately, the problem of insufficiently responsive action choices just discussed is compounded by action-based contracts. Why? As we have already explained, contract requirements under risk cannot vary with the firm-type profile. Therefore, *compliance* with an action-based contract implies that a firm always plays the same action—regardless of even its own type. Thus, assuming an action-based contract is successful in inducing compliance, the best it can do is to sustain a strategy profile $s(\theta) = a_{\rho}$ for all θ , where

$$a_{\rho} \in \arg \max_{a \in A} \sum_{\theta \in \Theta} \mathbb{E}(v - c_1 - c_2 | a, \theta) \rho(\theta). \quad (4)$$

When can an action-based contract achieve the level of performance set out in (4)? To provide an answer, consider a type profile θ under an arbitrary action-based contract $r(\theta) = a = (a_i, a_{-i})$ for all θ . Suppose the type profile is θ and firm i considers a deviation to a'_i under which $\mathbb{E}(u_i | a', \theta) > \mathbb{E}(u_i | a, \theta)$, where $a' = (a'_i, a_{-i})$. Then, i 's partner must be kept whole with respect to its expected payoff under a . Thus, for the deviation to be attractive, the difference in i 's expected

³⁶ And, as an aside, we demonstrate the relationship between our model under full information and the more familiar moral hazard models in the economics literature.

³⁷ In the extensive form representation, firm i 's information sets imply a partition of Θ made up of sets of the form $\{\theta_i\} \times \Theta_{-i}$. Actions must be constant on the elements of these sets. Put differently, i cannot vary its strategy in response to information it does not have about its partner's type.

unadjusted payoff must be greater than the difference in its partner's expected loss:

$$\mathbb{E}(u_i | a', \theta) - \mathbb{E}(u_i | a, \theta) \geq \mathbb{E}(u_{-i} | a, \theta) - \mathbb{E}(u_{-i} | a', \theta),$$

otherwise, damages payments from i to its partner more than wipe out any expected gain in unadjusted payoffs. Rearranging terms, this condition can be stated in terms of the available surplus in both cases; that is, the deviation is not profitable if

$$\mathbb{E}(v - c_1 - c_2 | a, \theta) - \mathbb{E}(v - c_1 - c_2 | a', \theta) \geq 0. \quad (5)$$

When this is true, there is not enough surplus available under the deviation both to increase i 's expected payoff and to keep its partner whole. The following proposition says if condition (6) holds on average (across partner types) with respect to a_p , then the maximum in (4) can be sustained.

PROPOSITION 4. *If, for all $i \in \{1, 2\}$, $\theta_i \in \Theta_i$, and $a'_i \in A_i$,*

$$\sum_{\theta'_{-i} \in \Theta_{-i}} [\mathbb{E}(v - c_1 - c_2 | a_p, (\theta_i, \theta'_{-i})) - \mathbb{E}(v - c_1 - c_2 | (a'_i, a_{p,-i}), (\theta_i, \theta'_{-i}))] p_{-i}(\theta'_{-i}) \geq 0, \quad (6)$$

then $r(\theta) = a_p$ for all θ sustains $s(\theta) = a_p$ for all θ .

Proposition 4 implies an essential moral hazard problem facing firms that wish to use an action-based contract under uncertainty with respect to the consequences of partner actions. Suppose θ is the type profile and that firm i truthfully represents its type. It is willing to sign a contract specifying its contribution in a_θ^* . Unfortunately, its partner may have an incentive to misrepresent its type in an effort to insert a lower cost action requirement for itself that is an inefficient best reply to its contribution in a_θ^* . Once the contract goes into force, i has no legal recourse—indeed, the contract requires the partner's deviation from first-best actions. Thus, we should expect to observe firms using action requirements that are robust to the spectrum of possible partner types, i.e., that sustain solutions to (4).

Shifting attention to deliverable-based contracts, we now show that they may now do strictly better than action-based contracts. The reason is intuitive. As with action-based requirements, information consistency requires us to assume that r imposes a single deliverable requirement on all firm-type pairs. The difference is that contractual specification of outcomes rather than actions removes the possibility of penalties for action choices per se. Firms are free to choose the most efficacious actions according to their own type to achieve outcomes that avoid deliverable penalties. Thus, instead of biasing firms toward (4), as do action-based damages, deliverable-based penalties do not present any inherent hindrance to the implementation of strategies that solve (3).

One may suspect that the present case in which partner actions are observed but have uncertain consequences is similar to one with unobserved actions but known consequences (i.e., the moral hazard model described in the previous section). This suspicion is correct: when the structure of our model meets a simple set of conditions, it is equivalent (in the sense of Kuhn 1953) to an instance of the moral hazard model elaborated earlier. Once this equivalence is established, Theorem 1 again characterizes the conditions under which deliverable-based damages are efficacious in inducing efficient actions from a given firm-type pair.

The preceding assertion suggests that the full information and risk cases under deliverable requirements are equivalent, since both are equivalent to the moral hazard model. However, the incorporation of uncertain type implies fundamental differences between the two cases, for two reasons. First, the information consistency conditions under risk are in significant contrast to those under full information. Unlike under full information, contractual performance requirements under risk cannot vary by firm-type pair and firm strategies cannot vary by partner type. Taken together, these imply (at least weakly) that the set of games for which efficient behavior can be induced by deliverable requirements is smaller under risk than full information. Second, although the equivalence conditions that we present are simple, they are not without substance. Although the games that meet these conditions strike us as an economically interesting class, they are special, in a way that happens to overcome the constraints mentioned in the first point.

With this in mind, let us now state the conditions, to which we refer collectively as *symmetric consequences*. They are as follows, for $i \in \{1, 2\}$: (1) each type is a bijection from the m actions in A_i to an arbitrary subset P_i of m probability distributions on W_i , and (2) Θ_i contains all the permutations of mappings from A_i to P_i .³⁸ *Symmetric consequences* describes situations in which firms know the possible distributions over partner deliverables (perhaps due to experience or research) but not which partner actions induce them. The partner's type is a "black box"—output potential is known, but the actions required to achieve that potential are not.

For example, imagine a construction market in which a general contractor wishes to subcontract to a painting firm that employs two painters. The contractor knows that the economics of the painting business typically leads subcontractors to employ one person who does high-quality work 90% of the time and another who does so only 20% of the time. This situation exhibits symmetric consequences: one type maps the first and

³⁸ Note that Θ_i is such that $\theta_i \in \Theta_i$ if and only if θ_i is a bijection from A_i to P_i . Thus, there are $m!$ elements in Θ_i .

second painter to probabilities of high quality of 0.9 and 0.2, respectively; the other reverses these, mapping the first and second painter to 0.2 and 0.9, respectively. Even though a contract could specify a specific painter be assigned to the job, the contractor does not know which one is better. Hence, a contract specifying delivery of “high-quality work” may be preferred.

Let us define an LM91 moral hazard game (where objects are indicated with overbars) using an arbitrary instance of our game under risk with constant deliverables performance requirements ($r(\theta) = w$ for all $\theta \in \Theta$). Let Γ denote this instance of our game. (Note that, unlike with full information, there are no subgames in Γ .) Set $\bar{n} = 2$, $\bar{A} = A$, and $\bar{W} = W$. For firm i , $\bar{A}_i = A_i$ and $\bar{u}_i(\cdot) = u_i(\cdot)$. Pick an arbitrary θ and define $\bar{t}_i(\cdot) = d_i(\cdot | \cdot, \theta)$ as well as $\bar{p}(\cdot | \cdot) = p(\cdot | \cdot, \theta)$.

PROPOSITION 5. *If Γ satisfies symmetric consequences, then, for every $\bar{a} \in \bar{A}$, there exist $s \in S$ such that, for all $\theta \in \Theta$, (1) $p(\cdot | s(\theta), \theta) = \bar{p}(\cdot | \bar{a})$, and (2) for $i \in \{1, 2\}$, $\mathbb{E}_i(s | \theta) = \mathbb{E}(\bar{\pi}_i | \bar{a})$. Moreover, if \bar{s} is a Nash equilibrium in the LM91 game, then s is an ambiguity-averse Bayesian equilibrium in Γ .*

With Proposition 5 in hand, Theorem 1 can be invoked to characterize the conditions under which a deliverables contract is capable of delivering first-best outcomes. (Note that symmetric consequences imply first-best action independence, up to a permutation of action labels). In the case of risk, action-based contracts are not optimal as they are under full information. As a firm’s uncertainty about the consequences of a partner’s actions increases, contracting on action requirements becomes more difficult. In this case, efficiency may be achieved by establishing outcome requirements and then leaving each partner to figure out how best to achieve them.

6. Results for Ambiguity

Under ambiguity, the incentive alignment effect of action- versus deliverable-based contracts operates in much the same way as in the case of risk; that is, deliverable requirements may induce efficient actions by all firm-type pairs in settings where action requirements fail to do so—and for the same reasons, as illustrated in the example below. The substantive difference between risk and ambiguity is that the presence of ambiguity causes firms to discount their subjective expectations, resulting in assessments of partnering benefits that are lower than the true expected value. Thus, some valuable projects that would have been adopted under risk may be forgone under ambiguity.

An interesting finding, illustrated earlier, is that deliverable-based contracts under compensatory damages may have the added benefit of creating a natural kind of “ambiguity insurance” and, in so doing, increase the aggregate ex ante subjective value of the alliance. As a result, projects that are objectively

worthwhile but rejected as a result of ambiguity aversion may become viable under compensatory damages. As we show, such a result cannot be achieved under risk or under ambiguity with action-based contracts.

Our first goal is to establish conditions under which a constant deliverables contract $r(\theta) \equiv w$ for all θ results in strictly greater aggregate, ambiguity-averse expected payoffs than no contract at all. Note that there are two ways a change in contract regime can lead to different expected payoffs. The first is in the sense of mechanism design: different contracts encourage different strategy choices and, hence, alter expected payoffs in equilibrium. The second is in a more direct sense: different contracts may alter deliverable-contingent transfer payments, which, for a given strategy profile, affect aggregate expected payoffs according to those payments and their respective probabilities. We focus our analysis on this latter case: examining the value-increasing effect of deliverable-based contracts under ambiguity with transfer payments, behavior held constant.

We restrict our attention to partnerships that satisfy first-best action independence and fix the strategy profile to s^* . Ideally, $r(\cdot) = w$ sustains s^* , but our primary interest is in understanding the effect of deliverable-based damages on aggregate expected payoffs. Since we must now compare expectations across contracts, we augment our notation so that $\mathbb{E}_i(s | \theta, r)$ is the ambiguity-averse expected payoff to firm i when the firm types are θ , the strategy profile is s , and damages payments are based on the contractual requirements mapping $r(\cdot)$. We wish to describe conditions under which, given θ and $r(\cdot) = w$,

$$\begin{aligned} \mathbb{E}_1(s^* | \theta, r) + \mathbb{E}_2(s^* | \theta, r) \\ > \mathbb{E}_1(s^* | \theta, r_\emptyset) + \mathbb{E}_2(s^* | \theta, r_\emptyset). \end{aligned} \quad (7)$$

Next, we impose an order on W according to increasing values of d_1 under $r = w$, i.e., such that, for $w^g, w^h \in W$, $h > g$, $d_1(w^h, w) \geq d_1(w^g, w)$.³⁹ Then, given the order imposed by r , a strategy profile s , and two arbitrary type profiles θ and θ' understood from the context, we say p satisfies the *strict single crossing property (SSCP)* if $p(w | s^*, \theta) - p(w | s^*, \theta')$ crosses zero once at some $w \in W$ for which $d_1(w) = 0$.

Given a firm-type profile θ , let $\theta^{\min i}$ replace the type of i ’s partner with one that minimizes i ’s expected net payoff—that is, $\theta^{\min i} \equiv (\theta_i, \hat{\theta}_{-i})$, where $\hat{\theta}_{-i}$ is a solution

$$\hat{\theta}_{-i} \in \arg \min_{\theta'_{-i} \in \Theta_{-i}} \mathbb{E}(\pi_i | s^*, (\theta_i, \theta'_{-i})). \quad (8)$$

To avoid trivial cases, assume that every type profile does not result in the same payoff under r and s^* ; i.e., there exist θ such that $\theta \neq \theta^{\min i}$ for $i \in \{1, 2\}$.

³⁹ Recall that, under a deliverables contract, damages payments depend only on the deliverable outcomes.

Under ambiguity, each firm assesses its ex ante subjective expected value for a particular partner type that places greatest weight on deliverables yielding relatively low payoffs. In other words, each firm assumes the worst possible outcomes from a set defined by partner type. Suppose the low-payoff deliverables for one firm tend to be the high-payoff outcomes for its partner, a situation ripe for moral hazard. Then, the distributions by which the firms compute their expected payoffs will place asymmetric weights on the deliverables. In such cases, deliverable requirements may induce the SSCP. The implied schedule of transfer payments can mitigate low-payoff outcomes for both firms directly. Given that partners assign probabilities to different outcomes asymmetrically under ambiguity, transfer payments can also increase the aggregate expected value, ex ante. This is because these payments transfer value from one firm to the other for deliverables where the receiving firm places greater weight than the paying firm.

PROPOSITION 6. Assume p satisfies the SSCP for all $\theta^{\min 1}$ and $\theta^{\min 2}$. If θ is a type profile such that $\theta \neq \theta^{\min i}$ for $i \in \{1, 2\}$, then condition (7) is true.

Proposition 6 demonstrates the sense in which deliverables contracts have the capacity to create “ambiguity insurance.” The intuition follows from the preceding discussion: when the SSCP holds for minimizing firm-type profiles, damages payments are made from firm i to $-i$ following poor outcomes for $-i$. Given ambiguity aversion, firm $-i$ cannot estimate the probabilities of the different outcomes accurately (in contrast to the risk case) and places a higher probability on these poor outcomes than does firm i . Thus, every dollar transferred by firm i to its partner in such a state increases the aggregate ex ante expected value. Note that this effect is not limitless; with sufficiently large transfers, the minimizing firm-type profile is the same for both firms (i.e., the true profile), which causes the SSCP requirement to fail.

It is also worth mentioning that, in some games, $\mathbb{E}_i(s^* | \theta, r) > \mathbb{E}_i(s^* | \theta, r_\emptyset)$ for both firms under every type profile other than those that minimize payoffs under s^* (the $\theta^{\min i}$'s). This is interesting given the balanced nature of damages payments. Even if the increase is only in aggregate, the ability to make up-front transfer payments in such situations suggests that firms facing ambiguous partner type may well have a preference for these kinds of contracts.

COROLLARY 1. Assume p satisfies the SSCP for all $\theta^{\min 1}$ and $\theta^{\min 2}$. Then

$$\sum_{\theta} (\mathbb{E}_1(s | \theta, r) + \mathbb{E}_2(s | \theta, r)) \rho(\theta) > \sum_{\theta} (\mathbb{E}_1(s | \theta, r_\emptyset) + \mathbb{E}_2(s | \theta, r_\emptyset)) \rho(\theta). \quad (9)$$

Corollary 1 is interesting for the following reason: in the case of risk, Proposition 6 may be true for some θ but never on average, as in (9). This distinction has loose empirical implications. As mentioned above, the SSCP is intuitive in moral hazard settings—unilateral deviations from efficient actions tend to increase the likelihood of high outcomes for the deviator and low ones for its partner. There are many situations that may meet these conditions, especially those where a contract is used to provide incentives to adopt efficient actions. Corollary 1 suggests the use of deliverable-based contracts is “more likely” by inexperienced partners (i.e., those facing ambiguity) in the sense that they induce such partners to engage in projects that they would otherwise pass up as a result of ambiguity discounts large enough to swamp their objective benefits.

PROPOSITION 7. Corollary 1 fails if either of the following conditions hold: (1) $s = a$ is sustained by $r = a$ for some $a \in A$, or (2) for $i \in \{1, 2\}$, $B_i = \{\rho_{-i}\}$ (risk).

The first condition is important because there are cases in which an action-based contract under ambiguity can sustain the implementation of a “second-best” action profile (e.g., a_ρ in the previous section). Nonetheless, when an action profile is sustained by a contract requiring it, both parties are in compliance, regardless of the ultimate deliverable profile generated. Thus, $d_1 = d_2 = 0$ for all w , which means that effective action-based contracts under ambiguous partner types can never augment value as in Proposition 6. The second condition says that the average payoff-enhancing effect of deliverable requirements is never possible under risk. Note that this effect is relative to ambiguity-averse expected payoffs; that is, an ambiguity discount is necessary for the increase in the subjected expected value, ex ante. Thus, other things equal, firms prefer risk to ambiguity. Still, as discussed above, the potential payoff enhancement under ambiguity does suggest that those faced with an unfamiliar project type and unfamiliar partner may find deliverable-based contracts to be especially attractive.

7. Conclusions

Our analysis examines the effect of different types of joint production contract (action versus deliverables) on equilibrium behavior given varying degrees of uncertainty about the stochastic consequences of partner actions under the legal institution of compensatory damages. We begin by showing that when partners know each other's type or the probability that certain actions lead to particular deliverables (i.e., the full information case), a contract specifying partner actions is sufficient to induce the partners to take those contracted actions. This is an ideal outcome, given that in the full information case firms know what the efficient

actions are for themselves and their partners. Next, we demonstrate that this result can be reversed—in the sense that deliverable requirements become the uniquely efficient option—in situations where firms are uncertain about partner type, but this uncertainty is well calibrated (i.e., *risk*). In this situation, firms know the probability that their partner is of a particular type and thus are able to assess the expected value of a project with some accuracy. However, because the partner type (and thus the ideal actions for each partner given a desired outcome) is unknown, assignment of required actions may reduce the expected value of a project when compared with assigning deliverable outcomes, where each firm will choose an action to best achieve the desired deliverable, given its type. In developing this result for risk, we also make a technical connection from uncertainty about partner type (adverse selection) to unverifiable partner actions (moral hazard).

Under our third case, *ambiguity*, when firms are unable to make precise assessments of the probabilities of partner type, we prove that deliverable requirements under compensatory damages not only increase the likelihood that firms and their partners will take efficient actions but also may have the added benefit of generating a rough form of ambiguity insurance. This ambiguity insurance effect increases the aggregate, perceived value of the project, *ex ante*. This increased value occurs because the deliverables contract with compensating damages not only reduces the payoffs from inefficient actions (i.e., an incentive alignment effect) but may also increase the expected payoffs from efficient actions. This ambiguity insurance effect differs from a straightforward insurance effect, since our result does not require firms to be risk averse, but it does require that firm actions cannot lead to all possible deliverables with equal probability. An important implication of the ambiguity insurance effect is that the resulting increase in value may induce firms to undertake objectively valuable projects that would not otherwise occur, thereby increasing both social welfare and the welfare of the participating parties.

Our formalism is sufficiently general to admit a number of potentially interesting empirical settings. Assuming that our three categories of uncertainty could be the result of relational experience, with more experience implying less uncertainty, then our results predict the stylized empirical correlations summarized in Table 17.

All of our results are developed in the context of compensatory damages. This feature was included to provide a realistic context within which to judge noncompliance penalties in contracts lacking specified financial contingencies. We made this assumption of compensatory damages following a previously discovered empirical fact that many complex, economically

Table 17 Summary: Partner Experience and Contract Structure

Level of experience	Stylized facts suggested by our results
Extensive with project class and partner	Action-based contracts dominant
Experience with project class, new partner	Greater use of deliverable-based contracts
New to project class and partner	Deliverable-based contracts especially favored for marginal projects

significant contracts do not include any such contingencies. Our findings shed new light on compensatory damages, especially the discovery of their potential to create the ambiguity insurance effect. However, deep insight into compensatory damages requires analysis comparing them to important alternatives—e.g., no damages, punitive damages, and specified damages. We leave this enquiry to future work.

Our findings are intended to bridge our existing knowledge of contract design and formal economic theory with the empirical realities we have observed in our own and other related work. With the development of a broader classification of contracts that speaks to the fundamental structure of contracts, rather than specific mechanisms, we hope to provide insight on how contracts are designed that better matches empirical observation. We also demonstrate how contracts can create value by changing the perceived value of projects *ex ante* (to more closely align with the true value) that allows managers to move beyond the common behavioral trait of ambiguity aversion in evaluating partnerships. We expect that contracts are thus used as a means to encourage firms to undertake value creating, joint work that would otherwise be unexplored because of a firm's extreme uncertainty about the true value of a project.

Many relevant details are still unexplored that would allow even closer mapping of theory onto empirical regularities. For example, we have left unexplored the implications of asymmetry between partners in their level of uncertainty for contract design. It would be useful to know whether significant differences between partners in terms of prior experience affect the ability to use, for example, contracts specifying required actions versus deliverables. It appears that action-based clauses are only possible when both partners know what actions are required of each other, which necessitates that action-based contracts be reserved for partnerships with little asymmetry in terms of prior experience. However, the structure of contracts in circumstances of significant asymmetry (i.e., when one partner is very experienced and the other has little experience) is still an open question that merits further investigation, both theoretically and empirically. Further refinements to exemplify variance in partner experience and which types of experience (e.g., with projects and/or partners)

are more salient for contract design would enhance our understanding of how organization, via contracts, influences firm actions and, thus, performance of such deals.

Acknowledgments

The authors gratefully thank R. Casadesus-Masanell; B. Cassiman; A. Friedenber; A. Galasso; J. Gans; I. Horstman; P. Klibanoff; F. Lafontaine; S. Masten; J. Oxley; F. Ruiz-Aliseda; participants in École Polytechnique, European School of Management and Technology (ESMT), Instituto de Estudios Superiores de la Empresa (IESE), Barcelona, London Business School, University of Michigan, University of Toronto, and University of California, Los Angeles workshops; and colleagues at the Olin Foundations of Strategy conference and Arizona State University Economics Department for their excellent suggestions and insights at various stages of this project. The authors also thank the department and associate editors as well as two anonymous reviewers for their patience, several careful readings, and a wealth of thoughtful feedback.

Appendix. Proofs

In what follows, $|\cdot|$ indicates the cardinality of a set. Several of the proofs adopt vector/matrix notation to facilitate the analyses. Such notation is introduced within the proof when needed.

PROOF OF PROPOSITION 1. Begin with an arbitrary firm-type profile $\theta \in \Theta$. By the premise, $r(\theta) = a_\theta^* = (a_{\theta,1}^*, a_{\theta,2}^*)$. Under full information, $B_i = \{b_i\}$, $b_i(\theta_{-i}) = 1$ for $i \in \{1, 2\}$. Under full information, the balancedness of damages implies that, for any strategy profile s ,

$$\mathbb{E}_1(s|\theta) + \mathbb{E}_2(s|\theta) = \mathbb{E}(u_1 + u_2 | s(\theta), \theta). \quad (10)$$

By the efficiency of a_θ^* ,

$$\mathbb{E}(u_1 + u_2 | a_\theta^*, \theta) \geq \mathbb{E}(u_1 + u_2 | a, \theta), \quad \text{for all } a \neq a_\theta^*. \quad (11)$$

Without loss of generality, suppose there exists some profitable deviation for firm 1; i.e., there exists a strategy profile $s = (s_1, s_2^*)$ such that

$$\mathbb{E}_1(s|\theta) > \mathbb{E}_1^*(s^*|\theta). \quad (12)$$

Under s , firm 1 breaches and firm 2 performs. Therefore, $d_2 = -d_1$ guarantees that firm 2 achieves a payoff of at least $\mathbb{E}(u_2 | a_\theta^*, \theta)$ for all $w \in W$. Under full information, this implies

$$\mathbb{E}_2(s|\theta) \geq \mathbb{E}_2(s^*|\theta). \quad (13)$$

Adding (12) and (13), we get $\mathbb{E}_1(s|\theta) + \mathbb{E}_2(s|\theta) > \mathbb{E}_1(s^*|\theta) + \mathbb{E}_2(s^*|\theta)$. Invoking balancedness condition (10), we have

$$\mathbb{E}(u_1 + u_2 | s(\theta), \theta) > \mathbb{E}(u_1 + u_2 | a_\theta^*, \theta), \quad (14)$$

a contradiction of efficiency condition (11). \square

PROOF OF PROPOSITION 2. The proposition is true if, given any $\varepsilon > 0$ no matter how small, an example can be generated in which the conclusion holds by picking (v, c) and distributions θ for arbitrarily sized A and W . The smallest cardinality of W permitted by the setup is four. We require

only two deliverable profiles to make the example work, so the following parameters can be used in any situation allowed by the setup requirements. Select distinct deliverable profiles $w, w' \in W$. Set $v(w) = v(w') = V$ and set all other deliverable profiles to zero. Let $c_1(w) = x - \frac{1}{4}\varepsilon > 0$ and $c_2(w') = x > 0$, and let both equal zero for all other deliverables.

Pick an action profile to be the efficient one, a_θ^* . Set $p(w | a_\theta^*, \theta) = 1$ and zero for all other profiles in W . Without loss of generality, pick a deviation profile for firm 1, $a = (a_1, a_{\theta,2}^*)$. Set $p(w' | a_\theta^*, \theta) = 1$ and zero for all other profiles in W . For all other $a' \in A \setminus \{a_\theta^*, a\}$, ensure $p(w | a', \theta) = p(w' | a', \theta) = 0$ (there are at least two other deliverable profiles on which to put weights summing to 1). Then, under the efficient action profile, $\mathbb{E}(u_1 + u_2 | a_\theta^*, \theta) = V - x + \frac{1}{4}\varepsilon$. Under the deviation by firm 1, $\mathbb{E}(u_1 + u_2 | a, \theta) = V - x$. Under any other choices, the expected value is zero. Therefore, a_θ^* is, indeed, efficient.

Under a_θ^* , firm 1's expected payoff is $\frac{1}{2}V - (x - \frac{1}{4}\varepsilon)$. Under a , firm 1 must reimburse firm 2 its cost under w' , x , less the "slippage" of ε permitted by the courts according to the premise. Therefore, its payoff is $\frac{1}{2}V - (x - \varepsilon)$. Thus, firm 1's expected payoff under the deviation exceeds that under the efficient profile by $\frac{3}{4}\varepsilon$. Of course, more complicated examples can be constructed, including for strictly positive firm types. The key is managing the tension between shrinking ε and creating a profitable deviation for one of the firms given the tightening damages payment it must relinquish while, at the same time, keeping the efficient action strictly better than any other. \square

PROOF OF PROPOSITION 3. The conclusion follows immediately from the definition of first-best action independence and the requirement that strategies be constant on partner type. \square

PROOF OF PROPOSITION 4. Assume $r = a$ for some $a \in A$ and consider the constant strategy profile $s = a$. Then, s is an AABE if, for $i \in \{1, 2\}$, all $\theta \in \Theta$, and all $s'_i \in S_i$,

$$\mathbb{E}_i(s|\theta) \geq \mathbb{E}_i(s'|\theta), \quad \text{where } s' = (s'_i, s_{-i}). \quad (15)$$

When θ_i and θ'_{-i} are clear from the context, let $\theta' = (\theta_i, \theta'_{-i})$. Since $r = s = a$ and $B_i = \{\rho_{-i}\}$ for $i \in \{1, 2\}$, the left-hand side of (15) can be expanded as follows:

$$\begin{aligned} \mathbb{E}_i(s|\theta) &= \sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{w \in W} \pi_i(w|a, \theta') p(w|a, \theta') \rho_{-i}(\theta'_{-i}) \\ &= \sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{w \in W} u_i(w) p(w|a, \theta') \rho_{-i}(\theta'_{-i}), \end{aligned} \quad (16)$$

where the last step is because $d_i(w|a, \theta) = 0$ for all $w \in W$. Let $a' = (a'_i, a_{-i})$, where $s'_i(\theta) = a'_i \neq a_i$ is a deviation by i when the type profile is θ . Then, expanding the right-hand side of (15),

$$\begin{aligned} \mathbb{E}_i(s'|\theta) &= \sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{w \in W} \pi_i(w|a', \theta') p(w|a', \theta') \rho_{-i}(\theta'_{-i}) \\ &= \sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{w \in W} [u_i(w) + d_i(w|a', \theta')] p(w|a', \theta') \rho_{-i}(\theta'_{-i}) \\ &= \sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{w \in W} (u_i(w) - \max\{\mathbb{E}(u_{-i}|a, \theta') - u_{-i}(w), 0\}) \\ &\quad \cdot p(w|a', \theta') \rho_{-i}(\theta'_{-i}). \end{aligned} \quad (17)$$

To conserve on notation, let $\Delta_i(s' | \theta) \equiv \mathbb{E}_i(s | \theta) - \mathbb{E}_i(s' | \theta)$. Then, condition (15) is satisfied if, for all $s'_i \in S_i$, $\Delta_i(s' | \theta) \geq 0$. By (16) and (17),

$$\begin{aligned} \Delta_i(s' | \theta) &= \sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{w \in W} (u_i(w)p(w|a, \theta') \\ &\quad - (u_i(w) \max\{\mathbb{E}(u_{-i}|a, \theta') - u_{-i}(w), 0\})p(w|a', \theta')) \\ &\quad \cdot \rho_{-i}(\theta'_{-i}) \\ &\geq \sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{w \in W} (u_i(w)p(w|a, \theta') - (u_i(w) \\ &\quad - \mathbb{E}(u_{-i}|a, \theta') + u_{-i}(w))p(w|a', \theta'))\rho_{-i}(\theta'_{-i}) \quad (18) \\ &= \sum_{\theta'_{-i} \in \Theta_{-i}} (\mathbb{E}(u_i|a, \theta') + \mathbb{E}(u_{-i}|a, \theta')) \\ &\quad - \sum_{w \in W} (u_i(w) + u_{-i}(w))p(w|a', \theta'))\rho_{-i}(\theta'_{-i}) \\ &= \sum_{\theta'_{-i} \in \Theta_{-i}} (\mathbb{E}(v - c_1 - c_2|a, \theta') - \mathbb{E}(v - c_1 - c_2|a', \theta')) \\ &\quad \cdot \rho_{-i}(\theta'_{-i}). \quad (19) \end{aligned}$$

Step (18) is because negative values of d_{-i} are set to zero given that i is the breaching party under a deviation from a . Thus, $r = a$ supports $s = a$ if $\Delta_i(s' | \theta)$ is greater than or equal to (19) for $i \in \{1, 2\}$, all $\theta \in \Theta$, and all $s'_i(\theta) = a'_i \in A_i$. \square

PROOF OF PROPOSITION 5. Let θ be the firm-type pair that determines $\bar{p}(\cdot | \cdot) = p(\cdot | \cdot, \theta)$. For each firm i , impose an arbitrary order on the m distributions in P_i ; i.e., $P_i = \{p_{i,1}, \dots, p_{i,m}\}$. Let $M \equiv \{1, \dots, m\}$. Next, order the elements in A_i so that, for $j \in M$, $\theta_i(a_{i,j}) = p_{i,j}$. By the premise of symmetric consequences, Θ_i contains all the permutations of maps from A_i to P_i . Therefore, define $f_i: M \times \Theta \rightarrow M$ such that $f_i(j, \theta) = k$ where, for $\theta' = (\theta'_i, \theta'_{-i})$, $\theta'_i(a_{i,k}) = p_{i,j}$. To establish the first part of the proposition, let $\bar{a} \in \bar{A}$ be an arbitrary action in the LM91 game. Since $\bar{A} = A$ by construction, assume $\bar{a} = a = (a_{1,j}, a_{2,j})$ (without loss of generality). Then, construct s in Γ such that, for $i \in \{1, 2\}$ and all $\theta' \in \Theta$, $s_i(\theta'_i) = a_{i, f_i(j, \theta')}$. Then, for all $\theta' \in \Theta$, $p(\cdot | s(\theta'), \theta') = \bar{p}(\cdot | \bar{a})$.

For part (2), by the definition of expected value,

$$\mathbb{E}(\bar{\pi}_i | \bar{a}) = \sum_{\bar{w} \in \bar{W}} \bar{\pi}_i(\bar{w})\bar{p}(\bar{w} | \bar{a}). \quad (20)$$

Since $\bar{W} = W$, (20) can be written

$$\mathbb{E}(\bar{\pi}_i | \bar{a}) = \sum_{w \in W} \bar{\pi}_i(w)\bar{p}(w | \bar{a}).$$

By construction, $\bar{u}_i = u_i$ and $\bar{t}_i = d_i$. Note that the latter equality is feasible because deliverable-based damages depend only on deliverables (and not actions or firm types). Thus, for all w, a and θ' , $\bar{\pi}_i(w) = \pi_i(w | a, \theta')$. Therefore,

$$\mathbb{E}(\bar{\pi}_i | \bar{a}) = \sum_{w \in W} \pi_i(w | a, \theta')\bar{p}(w | \bar{a}), \quad \text{for all } a \in A, \theta' \in \Theta. \quad (21)$$

Let s be the strategy in Γ constructed for \bar{a} in the first part. Then,

$$\mathbb{E}(\bar{\pi}_i | \bar{a}) = \sum_{w \in W} \pi_i(w | s(\theta'), \theta')p(w | s(\theta'), \theta'), \quad \text{for all } \theta' \in \Theta. \quad (22)$$

Recalling (1) in the case of risk, for all $\theta' \in \Theta$,

$$\begin{aligned} \mathbb{E}_i(s | \theta') &= \sum_{\theta'' \in \Theta'_i} \sum_{w \in W} \pi_i(w | s(\theta''), \theta'')p(w | s(\theta''), \theta'')\rho_{-i}(\theta''_{-i}). \quad (23) \end{aligned}$$

Substituting by (22),

$$\begin{aligned} \mathbb{E}_i(s | \theta') &= \sum_{\theta'' \in \Theta'_i} \mathbb{E}(\bar{\pi}_i | \bar{a})\rho_{-i}(\theta''_{-i}) \\ &= \mathbb{E}(\bar{\pi}_i | \bar{a}). \end{aligned}$$

This completes the proof for part (2).

For the last part, assume $\bar{a} = a = (a_{1,j}, a_{2,j})$ is a Nash equilibrium in the LM91 game. By the definition of a Nash equilibrium, for $i \in \{1, 2\}$ and all $a_{i,k} \in A_i$, $\mathbb{E}(\bar{\pi}_i | a_{1,j}, a_{2,j}) \geq \mathbb{E}(\bar{\pi}_i | a_{i,k}, a_{-i,j})$. Now, suppose s , as constructed above with respect to \bar{a} , is not an AABE. Then, for at least one firm i , there exists a type profile θ' and a strategy $s' = (s'_i, s_{-i})$ such that $\mathbb{E}_i(s' | \theta') > \mathbb{E}_i(s | \theta')$. From part (2), this implies $\mathbb{E}_i(s' | \theta') > \mathbb{E}(\bar{\pi}_i | \bar{a})$. By (23), this in turn implies there is at least one firm-type profile $\theta'' = (\theta'_i, \theta''_{-i})$ for which

$$\sum_{w \in W} \pi_i(w | s'(\theta''), \theta'')p(w | s'(\theta''), \theta'') > \mathbb{E}(\bar{\pi}_i | \bar{a}). \quad (24)$$

As discussed above, (24) can be rewritten

$$\sum_{w \in W} \bar{\pi}_i(w)p(w | s'(\theta''), \theta'') > \mathbb{E}(\bar{\pi}_i | \bar{a}).$$

Or, writing out the strategies and firm types,

$$\sum_{w \in W} \bar{\pi}_i(w)p(w | (s'_i(\theta'_i), s_{-i}(\theta''_{-i})), (\theta'_i, \theta''_{-i})) > \mathbb{E}(\bar{\pi}_i | \bar{a}).$$

By the definition of p , the assumption of symmetric consequences, and the construction of s ,

$$\sum_{w \in W} \bar{\pi}_i(w)p_{i,k}p_{-i,j} > \mathbb{E}(\bar{\pi}_i | \bar{a}), \quad \text{for some } k \neq j.$$

Thus, $\mathbb{E}(\bar{\pi}_i | a_{i,k}, a_{-i,j}) > \mathbb{E}(\bar{\pi}_i | \bar{a})$, which contradicts the premise that \bar{a} is an equilibrium. \square

PROOF OF PROPOSITION 6. The ambiguity case assumes $B = \Delta(\Theta)$, all the distributions including those placing unit mass on the individual elements of Θ . Therefore, ambiguity-averse expected value (1) in the context of $r = w$ given a firm-type profile θ and a strategy profile s is given by $\mathbb{E}_i(s | \theta, r) = \mathbb{E}(\pi_i | s, \theta^{\min i})$, where $\theta^{\min i}$ is any firm-type profile satisfying (8) for $i \in \{1, 2\}$. Thus,

$$\begin{aligned} \sum_{i \in \{1,2\}} \mathbb{E}_i(s | \theta, r) &= \sum_{i \in \{1,2\}} \mathbb{E}(\pi_i | s, \theta^{\min i}) \\ &= \sum_w (u_1(w) + d_1(w))p(w | s, \theta^{\min 1}) \\ &\quad + (u_2(w) - d_1(w))p(w | s, \theta^{\min 2}) \\ &= \sum_{i \in \{1,2\}} \mathbb{E}(u_i | s, \theta^{\min i}) \\ &\quad + \sum_w d_1(w)[p(w | s, \theta^{\min 1}) - p(w | s, \theta^{\min 2})]. \end{aligned}$$

By the assumption that p satisfies the SSCP, $\text{sgn}[d_1] = \text{sgn}[p(w | s, \theta^{\min 1}) - p(w | s, \theta^{\min 2})]$. Therefore,

$$d_1(w)[p(w | s, \theta^{\min 1}) - p(w | s, \theta^{\min 2})] \geq 0,$$

for all w and strictly so for some w . Hence,

$$\sum_{i \in \{1,2\}} \mathbb{E}_i(s | \theta, r) > \sum_{i \in \{1,2\}} \mathbb{E}(u_i | s, \theta^{\min i}). \quad (25)$$

Under the null contract r_\emptyset ,

$$\mathbb{E}_i(s | \theta, r_\emptyset) = \min_{\theta' \in \{\theta_1\} \times \Theta_{-i}} \mathbb{E}(u_i | s, \theta').$$

If the $\theta^{\min i}$'s also happen to solve $\mathbb{E}_i(s | \theta, r_\emptyset) = \mathbb{E}(u_i | s, \theta^{\min i})$, then Condition (7) follows immediately from (25). Suppose instead that θ^1 and θ^2 do not satisfy (8) but do minimize $\mathbb{E}(u_1 | s, \cdot)$ and $\mathbb{E}(u_2 | s, \cdot)$, respectively. Then,

$$\sum_{i \in \{1,2\}} \mathbb{E}(u_i | s, \theta^{\min i}) > \sum_{i \in \{1,2\}} \mathbb{E}(u_i | s, \theta^i). \quad (26)$$

Inequalities (25) and (26) imply Condition (7) by transitivity. \square

PROOF OF PROPOSITION 7. Item (1) follows directly from the fact that under $r = a$, $s = a$ implies $d_1 = d_2 = 0$ for all w . This, in turn, implies $\mathbb{E}_i(s | \theta, r) = \mathbb{E}_i(s | \theta, r_\emptyset)$ for $i \in \{1, 2\}$. (It also implies p cannot satisfy the SSCP.) Item (2) assumes $B_i = \{\rho_{-i}\}$ for $i \in \{1, 2\}$. For this part of the proof, define the average probability of w_i for firm i under s and ρ as $\bar{\theta}_i(w_i | s_i) = \sum_{\theta_i \in \Theta_i} \theta_i(w_i | s_i(\cdot, \theta_i)) \rho(\theta_i)$. Then, for all $w \in W$,

$$\begin{aligned} & \sum_{\theta \in \Theta} \sum_{i \in \{1,2\}} \mathbb{E}_i(s | \theta, r) \rho(\theta) \\ &= \sum_{\theta \in \Theta} \rho(\theta) \sum_{i \in \{1,2\}} \sum_{\theta' \in \{\theta_i\} \times \Theta_{-i}} \sum_{w \in W} (u_i(w) + d_i(w)) p(w | s, \theta') \rho_{-i}(\theta'_{-i}) \\ &= \sum_{\theta \in \Theta} \rho(\theta) \left[\sum_{\theta'_2 \in \Theta_2} \sum_{w \in W} (u_1(w) + d_1(w)) \right. \\ & \quad \cdot \theta_1(w_1 | s_1(\theta_1)) \theta'_2(w_2 | s_2(\theta'_2)) \rho_2(\theta'_2) \\ & \quad \left. + \sum_{\theta'_1 \in \Theta_1} \sum_{w \in W} (u_2(w) - d_1(w)) \theta_2(w_2 | s_2(\theta_2)) \theta'_1(w_1 | s_1(\theta'_1)) \rho_1(\theta'_1) \right] \\ &= \sum_{\theta \in \Theta} \rho(\theta) \left[\sum_{w \in W} (u_1(w) + d_1(w)) \theta_1(w_1 | s_1(\theta_1)) \right. \\ & \quad \cdot \sum_{\theta'_2 \in \Theta_2} \theta'_2(w_2 | s_2(\theta'_2)) \rho_2(\theta'_2) \\ & \quad \left. + (u_2(w) - d_1(w)) \theta_2(w_2 | s_2(\theta_2)) \sum_{\theta'_1 \in \Theta_1} \theta'_1(w_1 | s_1(\theta'_1)) \rho_1(\theta'_1) \right] \\ &= \sum_{\theta \in \Theta} \rho(\theta) \sum_{w \in W} (u_1(w) + d_1(w)) \theta_1(w_1 | s_1) \bar{\theta}_2(w_2 | s_2) \\ & \quad + (u_2(w) - d_1(w)) \theta_2(w_2 | s_2) \bar{\theta}_1(w_1 | s_1) \\ &= \sum_{\theta \in \Theta} \rho(\theta) \sum_{w \in W} u_1(w) \theta_1(w_1 | s_1) \bar{\theta}_2(w_2 | s_2) + u_2(w) \bar{\theta}_1(w_1 | s_1) \\ & \quad \cdot \theta_2(w_2 | s_2) + d_1(w) (\theta_1(w_1 | s_1) \bar{\theta}_2(w_2 | s_2) - \theta_2(w_2 | s_2) \bar{\theta}_1(w_1 | s_1)) \\ &= \sum_{\theta \in \Theta} [\mathbb{E}_1(s | \theta, r_\emptyset) + \mathbb{E}_2(s | \theta, r_\emptyset)] \rho(\theta) \\ & \quad + \sum_{w \in W} d_1(w) \sum_{\theta_1 \in \Theta_1} \sum_{\theta_2 \in \Theta_2} (\theta_1(w_1 | s_1) \bar{\theta}_2(w_2 | s_2) \rho_1(\theta_1) \rho_2(\theta_2) \\ & \quad - \theta_2(w_2 | s_2) \bar{\theta}_1(w_1 | s_1) \rho_1(\theta_1) \rho_2(\theta_2)) \\ &= \sum_{\theta \in \Theta} [\mathbb{E}_1(s | \theta, r_\emptyset) + \mathbb{E}_2(s | \theta, r_\emptyset)] \rho(\theta) \end{aligned}$$

$$\begin{aligned} & + \sum_{w \in W} d_1(w) (\bar{\theta}_1(w_1 | s_1) \bar{\theta}_2(w_2 | s_2) - \bar{\theta}_1(w_1 | s_1) \bar{\theta}_2(w_2 | s_2)) \\ &= \sum_{\theta \in \Theta} [\mathbb{E}_1(s | \theta, r_\emptyset) + \mathbb{E}_2(s | \theta, r_\emptyset)] \rho(\theta). \end{aligned}$$

Therefore,

$$\sum_{\theta \in \Theta} \sum_{i \in \{1,2\}} \mathbb{E}_i(s | \theta, r) \rho(\theta) = \sum_{\theta \in \Theta} \sum_{i \in \{1,2\}} \mathbb{E}_i(s | \theta, r_\emptyset) \rho(\theta),$$

which completes the proof. \square

References

- Aragones E, Gilboa I, Postlewaite A, Schmeidler D (2005) Fact-free learning. *Amer. Econom. Rev.* 95(5):1355–1368.
- Argyres N, Mayer KJ (2007) Contract design as a firm capability: An integration of learning and transaction cost perspectives. *Acad. Management Rev.* 32(4):1060–1077.
- Argyres NS, Bercovitz J, Mayer KJ (2007) Complementarity and evolution of contractual provisions: An empirical study of IT services contracts. *Organ. Sci.* 18(1):3–19.
- Baker G, Gibbons R, Murphy KJ (1994) Subjective performance measures in optimal incentive contracts. *Quart. J. Econom.* 109(4):1125–1156.
- Bernheim BD, Whinston MD (1998a) Exclusive dealing. *J. Political Econom.* 106(1):64–103.
- Bernheim BD, Whinston MD (1998b) Incomplete contracts and strategic ambiguity. *Amer. Econom. Rev.* 88(4):902–932.
- Bose S, Ozdenoren E, Pape A (2006) Optimal auctions with ambiguity. *Theoret. Econom.* 1(4):411–438.
- Bull C (1987) The existence of self-enforcing implicit contracts. *Quart. J. Econom.* 102(1):147–160.
- Casadesus-Masanel R, Klibanoff P, Ozdenoren E (2000) Maxmin expected utility over savage acts with a set of priors. *J. Econom. Theory* 92(1):35–65.
- Chamberlain G (2000) Econometric applications of maxmin expected utility. *J. Appl. Econometrics* 15(6):625–644.
- Corts KS, Singh J (2004) The effect of repeated interaction on contract choice: Evidence from offshore drilling. *J. Law, Econom., Organ.* 20(1):230–260.
- De Castro L, Yannelis NC (2011) Ambiguity aversion solves the conflict between efficiency and incentive compatibility. School of Economics Discussion Paper EDP-1106, University of Manchester, Oxford, UK.
- Ellsberg D (1961) Risk, ambiguity, and the savage axioms. *Quart. J. Econom.* 75(4):643–669.
- Epstein LG, Zin SE (1989) Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica* 57(4):937–969.
- Ghirardato P, Marinacci M (2002) Ambiguity made precise: A comparative foundation. *J. Econom. Theory* 102(2):251–289.
- Gilboa I, Schmeidler D (1989) Maxmin expected utility with non-unique prior. *J. Math. Econom.* 18(2):141–153.
- Gilboa I, Postlewaite AW, Schmeidler D (2008) Probability and uncertainty in economic modeling. *J. Econom. Perspect.* 22(3):173–188.
- Granovetter M (1985) Economic action and social structure: The problem of embeddedness. *Amer. J. Sociol.* 91(3):481–510.
- Grossman SJ, Hart OD (1986) The costs and benefits of ownership: A theory of vertical and lateral integration. *J. Political Econom.* 94(4):691–719.
- Gulati R (1995) Does familiarity breed trust? The implications of repeated ties for contractual choice in alliances. *Acad. Management J.* 38(1):85–112.
- Gulati R, Singh H (1998) The architecture of cooperation: Managing coordination costs and appropriation concerns in strategic alliances. *Admin. Sci. Quart.* 43(4):781–784.
- Kellner C (2015) Tournaments as response to ambiguity aversion in incentive contracts. *J. Econom. Surveys* 159(Part A):627–655.

- Kelsey D, Quiggin J (1992) Theories of choice under ignorance and uncertainty. *J. Econom. Surveys* 6(2):133–153.
- Khalil F, Lawarrée J (1995) Input versus output monitoring: Who is the residual claimant? *J. Econom. Theory* 66(1):139–157.
- Klibanoff P, Marinacci M, Mukerji S (2005) A smooth model of decision making under ambiguity. *Econometrica* 73(6):1849–1892.
- Knight FH (1921) *Risk, Uncertainty, and Profit* (Hart, Schaffner & Marx, Boston).
- Kotowski MH (2012) Engineered ambiguity in the principal-agent problem. Working paper, Harvard University, Cambridge, MA.
- Kuhn HW (1953) Extensive games and the problem of information. Kuhn HW, Tucker WW, eds. *Contributions to the Theory of Games II* (Princeton University Press, Princeton, NJ), 193–216.
- Kvaløy O, Olsen TE (2009) Endogenous verifiability and relational contracting. *Amer. Econom. Rev.* 99(5):2193–2208.
- Legros P, Matsushima H (1991) Efficiency in partnerships. *J. Econom. Theory* 55(2):296–322.
- Levin J (2003) Relational incentive contracts. *Amer. Econom. Rev.* 93(3):835–857.
- Lopomo G, Rigotti L, Shannon C (2014) Uncertainty in mechanism design. Working paper, Duke University, Durham, NC.
- Maccheroni F, Marinacci M, Rustichini A (2006) Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74(6):1447–1498.
- MacLeod WB, Malcomson JM (1993) Investments, holdup, and the form of market contracts. *Amer. Econom. Rev.* 83(4):811–837.
- Mayer KJ, Argyres NS (2004) Learning to contract: Evidence from the personal computer industry. *Organ. Sci.* 15(4):394–410.
- Mukerji S (1997) Understanding the nonadditive probability decision model. *Econom. Theory* 46(1):23–46.
- Mukerji S (1998) Ambiguity aversion and incompleteness of contractual form. *Amer. Econom. Rev.* 88(5):1207–1231.
- Poppo L, Zenger T (2002) Do formal contracts and relational governance function as substitutes or complements? *Strategic Management J.* 23(8):707–725.
- Radner R (1981) Monitoring cooperative agreements in a repeated principal-agent relationship. *Econometrica* 49(5):1127–1148.
- Ryall MD, Sampson RC (2009) Formal contracts in the presence of relational enforcement mechanisms: Evidence from technology development projects. *Management Sci.* 55(6):906–925.
- Tadelis S (2002) Complexity, flexibility, and the make-or-buy decision. *Amer. Econom. Rev.* 92(2):433–437.
- Uzzi B (1997) Social structure and competition in interfirm networks: The paradox of embeddedness. *Admin. Sci. Quart.* 42(1):35–67.
- Vanneste BS, Puranam P (2010) Repeated interactions and contractual detail: Identifying the learning effect. *Organ. Sci.* 21(1):186–201.
- Williamson OE (1975) *Markets and Hierarchies, Analysis and Antitrust Implications: A Study in the Economics of Internal Organization* (Free Press, New York).