Exploring Relations Between Decision Analysis and Game Theory

Jules H. van Binsbergen, Leslie M. Marx
The Fuqua School of Business, Duke University, 1 Towerview Drive, Durham, North Carolina 27708
{jules.vanbinsbergen@duke.edu, marx@duke.edu}

Many authors, including Cavusoglu and Raghunathan (2004. Configuration of detection software: A comparison of decision and game theory approaches. *Decision Anal.* 1(3) 131–148.) in this journal, have argued that proper modeling of the strategic interaction between players requires a game-theoretic approach as opposed to a decision-theoretic approach. We argue in this paper, however, that there are many environments in which decision analysis can deal with strategic interactions just as well, and we present equivalence results for such environments. These equivalence results allow the prescriptive decision analyst to use the standard tools that a sound decision analysis requires, including decision trees and sensitivity analysis, even when confronted with strategic settings. We further present two technical comments on the Cavusoglu and Raghunathan (2004) paper.

1. Introduction

In their recent paper published in this journal, Cavusoglu and Raghunathan (CR) (2004) contrast a decision analysis approach with a game-theoretic approach to the problem of configuring detection software.[1] They argue that proper modeling for this problem would require a game-theoretic approach and that a firm applying the game-theoretic approach would achieve higher gains than a firm applying a decision analysis approach. Underlying these arguments is an intellectual divide between the fields of decision analysis and game theory. The commonly made distinction between the two fields is that decision analysis focuses on one decision maker (DM) facing an uncertain environment, whereas game theory also studies the strategic interaction between decision makers.

In contrast, in this comment we argue that there are many situations in which decision analysis can deal with strategic interactions as well. Many sequential games, such as the two-stage games considered in this paper (e.g., the Stackelberg leader-follower game), can be analyzed using a decision analysis approach, as long as the decision theorist chooses the inputs to the decision problem carefully. For this class of games,

sophisticated decision analysis is equivalent to game theory. In two-player, two-stage games in which the follower has a unique best response, we show that as long as the chance nodes of the leader in the decision analysis are chosen to represent the strategy of the follower, a decision analysis is equivalent to a game-theoretic analysis. However, we also provide an example in which a decision analysis and a game-theoretic analysis differ in meaningful ways. It is in these cases that one might make the argument for going beyond the confines of decision analysis.

In addition to commenting on the conceptual basis for a distinction between game theory and decision analysis, we make two technical points related to the CR paper. First, we show that the reasoning presented by CR for their solution to the sequential game is not correct. Second, we argue that once the authors allow for a discretization of the strategy space, then one can easily tackle the problem considered in their paper using decision analysis, although one would need to use a more sophisticated response function than the naive one assumed in CR's decision analysis approach.

We begin by defining what we mean by decision analysis and what we mean by game theory. Decision analysis can be defined as a structured way of thinking about how actions taken in a decision envi-

---

[1] The same issue contained another application of game theory within a decision analysis by Lippman and McCardle (2004).

ronment influence the final outcome. As such, we can distinguish three main features of the decision situation: (i) the set of alternatives, (ii) the chance and unknown events that can affect the outcome, and (iii) the outcome itself. Decision analysis then constructs models that are mathematical representations of the relationships within and between these three features of the decision situation. A fourth main feature, then, is an algorithm for solving the constructed model, usually called the solution concept.[2]

Once the model is solved, it allows the DM to infer the implications of each course of action that he might take, so that he can better understand the relationship between his actions and his objectives. Note that these four key features of a decision analysis, i.e., (i) the set of alternatives, (ii) the chance and unknown events, (iii) the results, and (iv) the solution concept, have straightforward counterparts in game theory—namely, (i) the strategy set, (ii) the moves of nature, (iii) the payoff mapping, and (iv) the equilibrium concept. The difference between decision analysis and game theory can be summarized by the following elements of a game that a decision analysis does not have: (v) the other players and (vi) the dependence of the payoffs on the actions of the other players. However, we argue in this paper that having multiple players does not necessarily imply the need for a game-theoretic analysis. There are many situations in which a decision analysis is equivalent to a game-theoretic analysis, because in many situations it is possible to model the payoff dependence on the other players' actions as chance nodes.

For the remainder of this paper, the definitions of game theory and decision analysis described above are too broad. Instead, we will focus on the definitions given in the sections below. In these definitions, we assume that some basic concepts are understood, such as the definitions of a decision tree and an extensive-form and a normal-form game. Further, we define a solution to the decision tree as a selection of actions from the decision nodes that maximize the decision maker's expected payoff.

The comment proceeds as follows. In §2, we state the required definitions and then provide results for

[2] These definitions are paraphrased versions of the definitions formulated by the Decision Analysis Society.

environments with sequential or simultaneous strategic interaction. We assess how well pure decision analysis (i.e., decision analysis that does not allow randomization over actions) and extended decision analysis (i.e., decision analysis that allows for randomization) are able to address strategic situations. In §3 we offer our more technical comments on the CR paper. In §4, we conclude.

2. Decision Analysis vs. Game Theory

2.1. Definitions

We begin by defining what we mean by pure decision analysis and then we define an extended decision analysis. Whereas pure decision analysis does not allow a strategic mixing over strategies, extended decision analysis does allow such mixed strategies. These mixed strategies are, in fact, optimal strategies as long as the decision maker is randomizing over actions that each maximize his or her expected payoff. Specifically in the context of auditing and detection, which is the example introduced in CR, such mixtures over actions play an important role.

DEFINITION 1. A *pure decision analysis* is a solved decision tree that provides a selection of *pure* actions from the decision nodes that maximize the decision maker's expected payoff or utility.

Note that a pure decision analysis does not let a decision maker randomize over actions. In fact, when a decision maker is playing against nature, there is no value added from randomizing over actions. When two actions lead to the same outcome, the decision maker is simply indifferent between the two. The DM can then randomly select one of those actions, but there is no reason for applying one randomization mechanism versus another. This randomization is therefore of no interest to the pure decision analyst.

Second, we define what we mean by a game-theoretic analysis.

DEFINITION 2. A *game-theoretic analysis* is a game, either in extensive or normal form, together with an equilibrium strategy profile.

As described in the introduction, one of the main characteristics of a game is the presence of multiple players. We can study the behavior of multiple players in a decision analysis by simply considering multiple trees. We focus on games with two players, although the analysis extends to a finite number of

players. In the presence of two players, we consider two trees, one for each player.

DEFINITION 3. A *paired pure decision analysis* is a pair of solved decision trees, where the decision nodes in one tree correspond to the (degenerate) chance nodes in the other tree, and where the (degenerate) probabilities on the chance nodes in one tree correspond to the payoff-maximizing pure decisions in the other tree.

DEFINITION 4. A paired pure decision analysis is *equivalent* to a game-theoretic analysis if the selected actions for Player 1 in the first decision analysis match the player's equilibrium strategy in the game-theoretic analysis, and similarly for the selected actions for Player 2 in the second decision analysis.

A pure decision analysis is intended to analyze a DM's optimal actions when playing against nature. Therefore, there is no value added from playing randomized (mixed) strategies. We can extend pure decision analysis by allowing for mixed strategies, which we call extended decision analysis. The main motivation for introducing such mixed strategies is that in strategic settings a solution in pure strategies may not exist.

DEFINITION 5. An *extended decision analysis* is a solved decision tree that provides a selection of (possibly randomized) actions from the decision nodes that maximize the decision maker's expected payoff or utility.

We can extend the definitions of a paired decision analysis and of equivalence in the obvious way.

DEFINITION 6. A *paired extended decision analysis* is a pair of solved decision trees, where the decision nodes in one tree correspond to the chance nodes in the other tree, and where the probabilities on the chance nodes in one tree correspond to the payoff-maximizing decisions in the other tree.

DEFINITION 7. A paired extended decision analysis is *equivalent* to a game-theoretic analysis if the probabilities associated with selected actions for Player 1 in the first decision analysis match the player's equilibrium strategy in the game-theoretic analysis, and similarly for the selected actions for Player 2 in the second decision analysis.

2.2. Sequential Strategic Interaction

We now investigate how well a pure decision analysis is able to model the strategic interaction between players in a two-player leader-follower game, which is defined below.

DEFINITION 8. A *two-player leader-follower game* is a pair (A, π), where A_1 and A_2 are the action sets and $\pi_1: A_1 \times A_2 \to \mathbb{R}$ and $\pi_2: A_1 \times A_2 \to \mathbb{R}$ are the payoff functions of the two players. The follower has a unique best response for every possible action of the leader if $\forall a_1 \in A_1$, $\arg\max_{a_2 \in A_2} \pi_2(a_1, a_2)$ is a (nonempty) singleton set.

One popular equilibrium concept for this game is the subgame-perfect equilibrium, which we define below. In what follows, given set S, we let $\Delta(S)$ denote the set of probability distributions over the elements of S.

DEFINITION 9. A *subgame-perfect equilibrium* (SPE) of the two-player leader-follower game (A, π), where the follower has a unique best response for every possible action of the leader, is (σ_1^*, a_2^*), where $\sigma_1^* \in \Delta(A_1)$ and $a_2^*: A_1 \to A_2$, such that $\forall a_1 \in A_1$, $a_2^*(a_1) \in \arg\max_{a_2 \in A_2} \pi_2(a_1, a_2)$ and $\forall a_1 \in \{A_1 \mid \sigma_1^*(a_1) > 0\}$, $a_1 \in \arg\max_{a_1 \in A_1} \pi_1(a_1, a_2(a_1))$.

We can now state our first proposition.

PROPOSITION 1. *Given a two-player leader-follower game in which the follower has a unique best response for every possible action of the leader, for every game-theoretic analysis involving a pure-strategy SPE, there exists an equivalent paired pure decision analysis.*

PROOF. Take a pure-strategy SPE of the extensive-form game as given. It must specify, for each action $a_1 \in A_1$ of the leader, the unique action $a_2^{BR}(a_1) \in A_2$ of the follower; and it must specify, as the equilibrium strategy for the leader, an action $a_1^{BR} \in \arg\max_{a_1 \in A_1} \pi_1(a_1, a_2(a_1))$. We now construct an equivalent decision analysis. For the first decision tree, begin with a decision node for Player 1, with a branch for each action in A_1, and with each branch followed by a chance node that has a branch for each action in A_2. For all $a_1 \in A_1$, for the chance node attached to decision node branch a_1, let the probability be one on branch $a_2^{BR}(a_1)$ and zero on all other branches. For the decision node, select the branch a_1^{BR}. Given the definitions of a_1^{BR}, this is a (correctly) solved decision tree. For the second decision tree, begin with a chance node for Player 1, with the branch for action a_1^{BR} receiving probability one and all other branches receiving probability zero. At the end of each branch, attach a

decision node with a branch for each action in A_2. For the decision node attached to chance node branch a_1, select branch $a_2^{BR}(a_1)$. Given the definition of a_2^{BR}, this is a (correctly) solved decision tree. By construction, the paired pure decision analysis is equivalent to the game-theoretic analysis. □

As an application of Proposition 1, consider the usual Stackelberg leader-follower game. To begin, view the leader as the decision maker and view the follower's behavior as probabilistic, but use the follower's best-response function to determine the conditional probabilities (conditional on the leader's choice) associated with his available actions. For any two-player sequential move game in which the follower has a unique best response for every possible action of the leader, a decision analysis approach with this type of sophisticated modeling on the leader's uncertainty regarding the follower's action is equivalent to a game-theoretic approach.

To give a more concrete example of the Stackelberg game, assume two firms, linear demand $P = a - bQ$, where $Q = q_1 + q_2$, and constant marginal cost c. Assume $a > c \geq 0$ and $b > 0$. In the first stage the leader, firm 1, chooses its quantity, and then in the second stage the follower, firm 2, Chooses its quantity. The payoff for firm i is $(a - bQ - c)q_i$.

In any SPE of the game, given quantity q_1 for the leader, the follower chooses q_2 to solve

$$\max_{q_2}(a - bq_1 - bq_2 - c)q_2.$$

Thus, the follower chooses $q_2^r(q_1) = (a - c)/2b - q_1/2$ in the subgame defined by a quantity choice of q_1 by the leader. Given this reaction function, the leader maximizes its payoff by choosing q_1 to solve

$$\max_{q_1}(a - bq_1 - bq_2^r(q_1) - c)q_1.$$

Thus, in the unique SPE of the game, the leader chooses $q_1 = (a - c)/2b$ and, on the equilibrium path, the follower chooses $q_2 = (a - c)/4b$. Off the equilibrium path, the follower chooses according to the reaction function $q_2^r(q_1)$.

To construct the paired decision analysis for this game, first construct a tree for the leader in which the decision node for the leader is followed by a chance node representing the follower's actions. Following choice q_1 by the leader, assign a probability of one to

action $q_2^r(q_1)$ by the follower. In this tree, the leader maximizes its payoff with choice $q_1 = (a - c)/2b$. Second, construct a tree for the follower in which the initial node is a chance node, representing the leader's actions, and then each branch for the chance node is followed by a decision node for the follower. Assign a probability of one to the branch representing $q_1 = (a - c)/2b$. In this tree, the follower maximizes its payoff with choice $q_2^r(q_1) = (a - c)/2b - q_1/2$ following the chance node branch representing $q_1 = (a - c)/2b$.

To give an example with finite action spaces, consider a game of product differentiation. There are two firms, a leader and a follower. The leader moves first, deciding between two possible product configurations, A and B. Then the follower responds with a product configuration decision of its own, choosing between configurations C and D. The payoffs are as follows, where $a \in [0, 30]$:

	Follower	
Leader	C	D
A	10, 10	20, 20
B	30, a	10, $30 - a$

In this game, the follower's best reply to A is D, and the follower's best reply to B is C if $a > 15$, and D otherwise (with indifference if $a = 15$). Thus, in the SPE, if $a < 15$, the leader chooses A, anticipating a payoff of 20 from A and a payoff of 10 from B; and if $a > 15$, the leader chooses B, anticipating a payoff of 20 from A and a payoff of 30 from B. Thus, the SPE outcome is (A, D) if $a < 15$ and (B, C) if $a > 15$. (If $a = 15$, then there is an SPE in which the leader chooses A, the follower chooses D, and the follower chooses D off the equilibrium path after the leader chooses B; and there is an SPE in which the leader chooses B, the follower chooses C, and the follower chooses D off the equilibrium path after the leader chooses A.)

Once again, the game can be represented as a paired decision analysis. In the leader's decision tree, the follower choice is modeled as a chance node with a probability of one of the follower's best reply. In the follower's decision tree, the leader's choice is modeled as a chance node with probability one on the leader's payoff-maximizing strategy.

Note that in this game, as the parameter a increases from below 15 to above 15, the follower's best reply to B changes and, as a result, the leader's optimal strategy changes. Thus, even though the parameter affects only the follower's profit, changes in the parameter can cause a shift in the leader's optimal strategy because the leader anticipates the effect of the parameter on the follower's best reply.

These results show that thoughtful decision analysis can be used in place of game-theoretic analysis for certain types of sequential strategic interaction. In the next section, we consider simultaneous strategic interaction.

2.3. Simultaneous Strategic Interaction

Now we investigate how well a pure decision analysis is able to model the strategic interaction between players in a two-player simultaneous-move game, which we define below.

DEFINITION 10. A *two-player simultaneous game* is a pair (A, π), where A_1 and A_2 are the action sets and $\pi_1 \colon A_1 \times A_2 \to \mathbb{R}$ and $\pi_2 \colon A_1 \times A_2 \to \mathbb{R}$ are the payoff functions of the two players.

DEFINITION 11. A *Nash equilibrium* of the two-player simultaneous game (A, π) is (σ_1^*, σ_2^*), where $\sigma_1^* \in \Delta(A_1)$ and $\sigma_2^* \in \Delta(A_2)$, such that

$$\sigma_2^* \in \arg\max_{\sigma_2 \in \Delta(A_2)} \pi_2(\sigma_1^*, \sigma_2)$$

and

$$\sigma_1^* \in \arg\max_{\sigma_1 \in \Delta(A_1)} \pi_1(\sigma_1, \sigma_2^*).$$

In what follows we refer to a strategy in a game-theoretic analysis as *strictly mixed* if it places positive probability on at least two actions, i.e., if it is not a pure strategy.

As the following proposition shows, we now have another example of equivalence between decision analysis and game-theoretic analysis.

PROPOSITION 2. *Given a game-theoretic analysis consisting of a two-player simultaneous-move game and a pure-strategy Nash equilibrium, there exists an equivalent paired pure decision analysis.*

PROOF. Take a pure-strategy Nash equilibrium of the simultaneous-move game as given. As in the proof of Proposition 1, one can choose actions in the paired decision analysis to match the equilibrium actions in the game-theoretic analysis. □

In contrast, when the Nash equilibrium of the simultaneous-move game is in mixed strategies, and each player in a decision analysis can only select one action, there is no pure paired decision analysis that is consistent with the Nash equilibrium. This suggests a flaw in the standard definition of a decision analysis, but as stated in the following proposition, this flaw is easily corrected by allowing a DM to randomize over actions that give it the same expected payoff or utility.

PROPOSITION 3. *Given a game-theoretic analysis consisting of a two-player simultaneous-move game and a strictly mixed Nash equilibrium, there exists a paired extended decision analysis that is equivalent to the Nash equilibrium of the game, although there does not exist an equivalent paired pure decision analysis.*

PROOF. Take the Nash equilibrium of the simultaneous-move game as given. As in the proofs of the above results, choose probabilities in the extended paired decision analysis to match the equilibrium probabilities. □

Although Proposition 3 shows that a paired extended decision analysis exists that is equivalent to a game-theoretic analysis for a two-player simultaneous-move game with a unique Nash equilibrium in completely mixed strategies, to find the Nash equilibrium, one must go beyond a one-step iteration of the paired decision trees.

To illustrate the use of decision analysis in simple simultaneous move games, consider the inspection game (Fudenberg and Tirole 1993, Example 1.7). In that game the firm chooses either to inspect or not inspect a worker, who chooses either to shirk or work. Working costs the user g and produces output with value v for the firm. Inspection costs h to the firm. The firm pays the user w unless the user is caught shirking. The payoffs are as follows, where $g > h > 0$ and $w > g$:

Firm	Worker	
	S	W
I	$-h, 0$	$v - w - h, w - g$
NI	$-w, w$	$v - w, w - g$

In the unique Nash equilibrium of the game, the firm mixes between inspecting and not, with probability g/w on inspection, and the worker mixes between shirking and working, with probability h/w on shirking.

In games such as this (simultaneous-move games with a unique Nash equilibrium in completely mixed strategies), the game-theoretic approach requires that one solve the following system of equations, where x denotes the probability with which the worker shirks and y denotes the probability with which the firm inspects:

$$-hx - (v - w - h)(1 - x) = -wx + (v - w)(1 - x)$$

$$w(1 - y) = w - g.$$

One could identify the Nash equilibrium by setting up two decision trees, one with the firm as the decision maker and one with the worker as the decision maker, as shown in Figure 1, and then iterating until the two decision trees were consistent in the sense that the behavior anticipated by the decision maker in one tree was, in fact, optimal for the decision maker in the other tree, and vice versa. If such a process converges, then it identifies a Nash equilibrium.

In Figure 1, the upper tree depicts the tree of the worker who chooses between shirking and working, or some randomization between the two with a probability x on shirking. Similarly the second tree shows the decision problem of the firm who chooses to either inspect, or not inspect, or to mix with a probability y on inspecting. As described above, in the Nash equilibrium, the two trees should be consistent with one another. This is guaranteed by the strategic use of the variables x and y, which appear in one tree as the decision variable and in the other as a subjective probability. When $x = h/w$ the expected value for each of the firm's alternatives is equal, and therefore the firm may as well mix over inspecting and not inspecting. Equivalently, when $y = g/w$ the worker may as well mix over shirking and not shirking. Therefore, for these values of x and y, the two trees are consistent.

For any other values of x and y the trees will be inconsistent. Suppose, for example, that the firm's subjective probability of dealing with a shirking worker equals $x > h/w$. Then it is optimal for the firm always to inspect. However, if the firm always

Figure 1   Decision Tree Representation of the Inspection Game
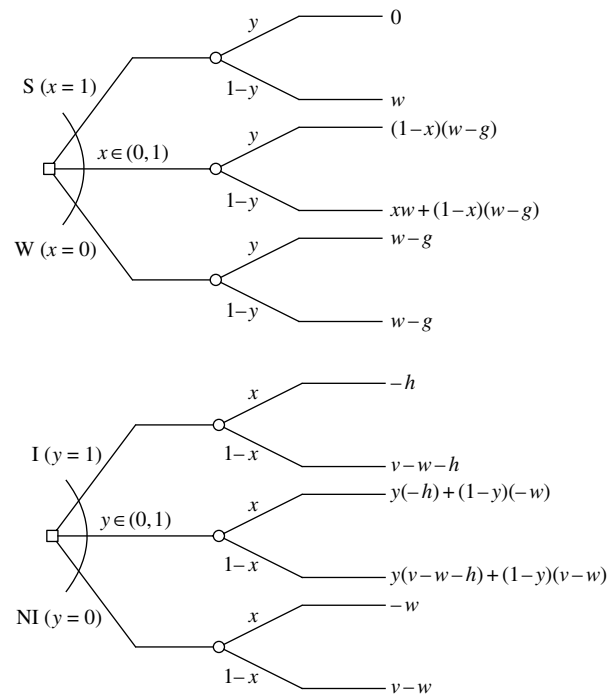
inspects, consistency of the two trees requires that the worker's subjective probability of facing an inspecting firm be $y = 1$. This means that the agent should never shirk, which (once more appealing to the consistency between the trees) implies that the firm should have a subjective probability of dealing with a shirking worker of $x = 0$, which contradicts the initially assumed subjective probability $x > h/w$.

As mentioned above, to identify the Nash equilibrium in the two decision trees of Figure 1, one could iterate until the two decision trees are consistent. However, this iterative process is not part of a standard decision analysis. In other words, there is no single-pass approach, such as backwards induction or rolling back the tree, that delivers the solution to this problem. It is problems like this one, in which the simultaneity is the key issue, that allow us to distinguish the decision analysis and game-theoretic approaches.[3]

---

[3] As noted by Fudenberg and Tirole (1993, p. 18), one can easily compute the optimal contract, i.e., the w that maximizes the principal's (i.e., the firm's) expected payoff. This optimal wage is given by $w = \sqrt{hv}$. Further, the game substantially changes when the

One thing that is part of standard decision analysis is sensitivity analysis to see how the recommended decision changes with changes in various parameters. In the context of a paired decision analysis, sensitivity analysis would be conducted with respect to some parameter by changing the value of that parameter in both of the trees in the paired decision analysis. A paired decision analysis specifies a particular solution for each tree (multiple solutions are possible if the DM is indifferent between alternatives at a decision node), and once a parameter value is changed, one can easily check whether the prior solution continues to be a solution in the revised decision trees. If it is not, then the decision, and corresponding equilibrium in an equivalent game-theoretic analysis, is not robust to that parameter change. As described above, in some cases an iterative process allows one to identify new solutions to the two trees such that they once again form a paired decision analysis, where the selections from the decision nodes in one tree correspond to the probabilities on the chance nodes in the other tree. A new solution identified in this way suggests one possible change in the outcome that might result from the parameter change under consideration. However, as noted above, there may be no finite iterative process that can identify a new solution. In this case, the sensitivity analysis may be best performed in the context of the game-theoretic analysis, where one can identify the new equilibrium of the game.

3. Solving the Detection Software Game

The remainder of our paper makes two main points regarding CR. First, we argue that the reasoning that they present for their solution to the sequential game is not quite right. Our second point is that once the authors allow for a discretization of the strategy space (their $\varepsilon > 0$), one can easily tackle the problem considered in their paper using decision analysis. To do this,

we need to apply the method we explained in the previous sections. Our method proposes a more sophisticated response function than the naive one assumed in CR's decision analysis approach.

3.1. Equilibrium in CR's Sequential Game

In the sequential game considered in CR, there are two players, the firm and the user. The idea behind the game is that the firm must set up detection software and a strategy for conducting audits, while the user decides whether or not to commit fraud. In the first stage, the firm chooses the configuration point and the inspection strategy. The configuration point represents the trade-off between false positives and false negatives and is given by (P_D, P_F), where P_D is the probability of detection conditional on there being fraud, and P_F is the probability of a false positive. The inspection strategy consists of two probabilities (ρ_1, ρ_2), where ρ_1 is the probability of investigation when the program signals a fraud and ρ_2 is the probability of investigation when the program signals no fraud. The damage from an undetected fraud for the firm is d and the cost of manual investigation is c. The user's pure strategies are to commit fraud or not. The user receives a benefit of μ for committing fraud if it is not detected. If the fraud is detected, the user incurs a penalty of β, leading to a net benefit of $\mu - \beta$. The probability that a user commits fraud is denoted by ψ.

To ease notation, let $s_1 \in S_1 \equiv \{(P_D, P_F, \rho_1, \rho_2) \mid P_D, P_F, \rho_1, \rho_2 \in [0, 1]\}$ denote the firm's strategy and $s_2 \in \{F, N\}$ denote the user's strategy. Let $\pi_1(s_1, s_2)$ be the firm's payoff given the players' strategies, and let $\pi_2(s_1, s_2)$ be the user's payoff given the players' strategies. The payoffs for pure strategies, given the values of (P_D, P_F), are shown in the following table:

Firm	User	
	F	NF
$\rho_1 = 1, \rho_2 = 1$	$(-c, \mu - \beta)$	$(-c, 0)$
$\rho_1 = 1, \rho_2 = 0$	$(-d - (c-d)P_D, \mu - P_D\beta)$	$(-cP_F, 0)$
$\rho_1 = 0, \rho_2 = 1$	$(-c + (c-d)P_D, \mu - (1-P_D)\beta)$	$(-c(1-P_F), 0)$
$\rho_1 = 0, \rho_2 = 0$	$(-d, \mu)$	$(0, 0)$

In the second stage of the game, the user takes the firm's strategy as given and chooses s_2 to solve $\max_{s_2 \in \{F, N\}} \pi_2(s_1, s_2)$. If $\pi_2(s_1, F) > \pi_2(s_1, N)$, then

---

principal can "commit" to an inspection level. In this case, the principal is better off. Consider the game where the principal plays first and commits to an inspection level y. The agent (i.e., the worker), after observing y, chooses whether to shirk. For a given value of w ($> g$) the principal can choose $y = g/w + \varepsilon$, where ε is positive and small. The agent then works with a probability of one. Note that this commitment essentially turns the simultaneous game into a sequential game, which is the subject of §2.2.

$s_2^*(s_1) = F$; if $\pi_2(s_1, F) < \pi_2(s_1, N)$, then $s_2^*(s_1) = N$; and if $\pi_2(s_1, F) = \pi_2(s_1, N)$, then the user is indifferent between F and N and so any mixture between F and N is a best response. Let $\sigma_2^*(s_1)$ denote the user's equilibrium strategy.

In the first stage of the game, the firm takes the user's strategy as given and chooses s_1 to solve $\max_{s_1 \in S_1} \pi_1(s_1, \sigma_2^*(s_1))$. CR identify a partition of S_1, $\{S_1^F, S_1^N, S_1^M\}$, where the user's unique best response for any $s_1 \in S_1^F$ is fraud, his unique best response for any $s_1 \in S_1^N$ is no fraud, and for any $s_1 \in S_1^M$, he is indifferent between fraud and no fraud.

CR dismiss the case of $s_1 \in S_1^F$. Presumably they have in mind an assumption under which the firm strictly prefers to choose some $s_1 \in S_1^N \cup S_1^M$ over any $s_1 \in S_1^F$. Clearly, the optimal strategy for the firm in S_1^F is the least-cost solution of never investigating, i.e., $\rho_1 = \rho_2 = 0$. In this case, the firm's payoff is $-d$, where d is the damage from undetected fraud. CR's Case 1 considers the case in which $s_1 \in S_1^N$. A correct analysis of this case shows that because the set S_1^N is open, the solution to $\max_{s_1 \in S_1^N} \pi_1(s_1, N)$ does not exist. Thus, there can be no equilibrium in which $s_1 \in S_1^N$. In contrast, CR discretize the set S_1^N and find the limit of the firm's best reply as the unit of discretization approaches zero.

It is important for the calculation of the solution that CR identify an upper bound for $\max_{s_1 \in S_1^N} \pi_1(s_1, N)$, which we can write as $\sup_{s_1 \in S_1^N} \pi_1(s_1, N)$. Next, CR's Case 2 considers the case in which $s_1 \in S_1^M$. CR dismiss this case as never giving the firm a higher payoff than in their Case 1. However, an optimal strategy for the firm does not exist for Case 1. A correct analysis would ask whether there exists $\psi \in [0, 1]$ such that there is a solution to $\max_{s_1 \in S_1^M} \psi\pi_1(s_1, F) + (1 - \psi)\pi_1(s_1, N)$ and such that $\max_{s_1 \in S_1^M} \psi\pi_1(s_1, F) + (1 - \psi)\pi_1(s_1, N) \geq \sup_{s_1 \in S_1^N} \pi_1(s_1, N)$. If such a ψ can be found, then we have identified an equilibrium of the game. Furthermore, as long as we can rule out equilibria involving $s_1 \in S_1^F$, the set of all such ψ characterize the set of all equilibria of the game.

Because S_1^M is compact and $\pi_1(s_1, F)$ and $\pi_1(s_1, N)$ are continuous in s_1, a solution exists for any ψ. However, one can show that for $\psi > 0$, $\max_{s_1 \in S_1^M} \psi\pi_1(s_1, F) + (1 - \psi)\pi_1(s_1, N) < \sup_{s_1 \in S_1^N} \pi_1(s_1, N)$. (This argument appears indirectly in CR.) Thus, there can be no equilibrium involving $\psi > 0$.

Figure 2   Decision Tree for the Firm in the Detection Software Game
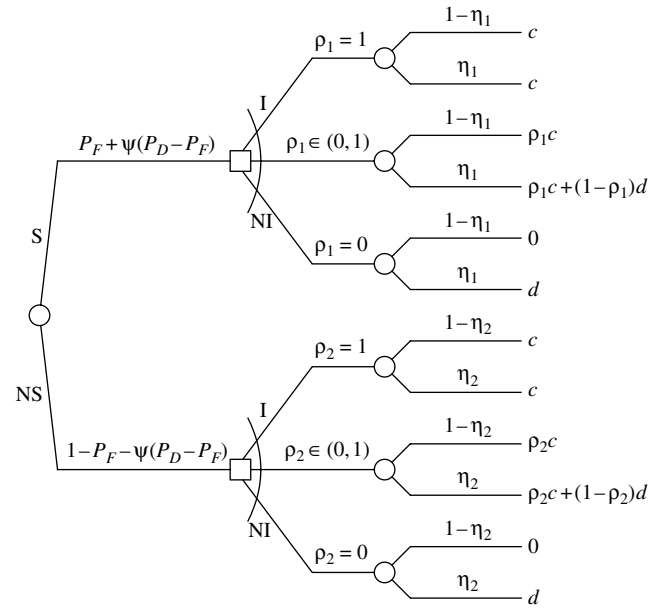
Note. The probabilities η_1 and η_2 are given, respectively, by $\eta_1 = P_D\psi/(P_D\psi + P_F(1 - \psi))$ and $\eta_2 = (1 - P_D)\psi/((1 - P_D)\psi + (1 - P_F)(1 - \psi))$.

Thus, we conclude that in any equilibrium $\psi = 0$, i.e., the user commits fraud with probability zero. Assuming that $\pi_1(s_1, N)$ has a unique maximizer on S_1^M, then the unique equilibrium is for the firm to choose s_1 to solve $\max_{s_1 \in S_1^M} \pi_1(s_1, N)$ and the user not to commit fraud. (If π_1 does not have a unique maximizer on S_1^M, then there exist equilibria in which the firm uses a mixed strategy.) In particular, the equilibrium value of s_1 is not an element of S_1^N as implied by the analysis presented in CR.[4] Although in equilibrium both players choose pure strategies, in equilibrium the user is indifferent between committing fraud and not. In addition, CR's arguments do not clearly establish the uniqueness of the equilibrium.

3.2. Extended Decision Analysis in CR's Sequential Game

We now apply our extended decision analysis approach to CR's detection software application. First, we present the two decision trees. Figures 2 and 3

---

[4] On p. 141, CR state that "the equilibrium solution of the sequential game when $\rho_1 P_D + \rho_2(1 - P_D) > \mu/\beta$ is as follows. $\rho_1^* = 1$, $\rho_2^* = 0$, $P_D^* = \mu/\beta, \ldots$." Although this suggests that the game has an equilibrium in this case, one can show that the condition $\rho_1 P_D + \rho_2(1 - P_D) > \mu/\beta$ is not satisfied at the equilibrium values.

Figure 3   Decision Tree for the User in the Detection Software Game
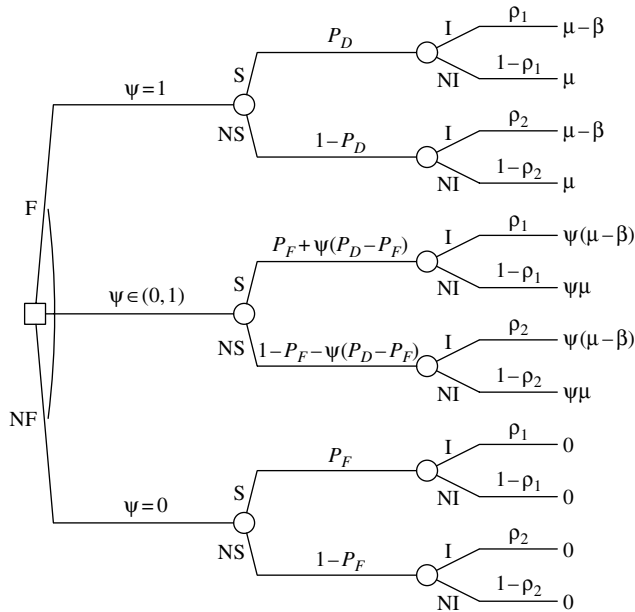
below show the decision trees of, respectively, the firm and the user. The firm chooses a mixed strategy where, conditional on a signal, it investigates with probability ρ_1, and conditional on no signal it investigates with probability ρ_2. The user chooses a mixed strategy over committing fraud or no fraud, i.e., commit fraud (F) with probability ψ and no fraud (NF) with probability $1-\psi$. The symbols S and NS indicate whether there is a signal (S) or no signal (NS) from the detection software. The detection software gives a false positive (i.e., it detects a fraud when there is none) with probability P_F and gives a false negative (i.e., it does not detect a fraud when there is) with probability $1-P_D$.

In the sequential game, the firm rolls back the tree of the user to derive the probability of fraud ψ as a function of the probabilities ρ_1 and ρ_2, substitutes this probability into its own decision tree and optimizes over ρ_1 and ρ_2.

4.   Conclusion

In a response to CR, which was published in this journal, we formulate three points, one conceptual and two technical. Conceptually, we argue in this paper

that there are many situations in which decision analysis can deal with strategic interactions just as well as game theory can. When the equilibrium of the game is in pure strategies, it is relatively straightforward to find decision trees that are consistent with the equilibrium of the game. However, when the equilibrium of the game involves mixed strategies, we need to extend decision analysis for it to be able to model strategic situations.

One might argue that in most decision settings, some player has to move first. As long as the decision made by that first mover is observed by the second mover, then allowing for randomization over actions is sufficient to prove equivalence between game theory and an extended decision analysis, at least one with a sophisticated view towards modeling expectations about the second mover's behavior. In this environment, game theory adds value to the decision analysis through its suggestion on how the decision maker should form its beliefs about the behavior of the environment, i.e., about the actions of players whose behavior is modeled as uncertain.

We further make two technical points related to the CR paper. First, the reasoning presented by CR for their solution to the sequential game is not quite correct. One can also complete their argument by establishing uniqueness of the equilibrium. Second, we argue that once the authors allow for a discretization of the strategy space, then one can easily tackle the problem considered in the paper using decision analysis. To do this, one would need to use a more sophisticated response function than the naive one assumed in CR's decision analysis approach.

References

Cavusoglu, H., S. Raghunathan. 2004. Configuration of detection software: A comparison of decision and game theory approaches. *Decision Anal.* **1**(3) 131–148.

Fudenberg, D., J. Tirole. 1993. *Game Theory*. MIT Press, Cambridge, MA.

Lippman, S. A., K. F. McCardle. 2004. Sex, lies, and the Hillblom estate. *Decision Anal.* **1**(3) 149–166.