Game Theory and the Law

Douglas G. Baird
Robert H. Gertner
Randal C. Picker

Harvard University Press
Cambridge, Massachusetts, and London, England

Backwards induction and subgame perfection. Selten (1975) pioneered the concept of subgame perfection and other refinements of the Nash equilibrium solution concept. Kreps (1990a) also has a good introduction to subgame perfection. Fudenberg and Tirole (1991a) explores the basic principles at work in the extensive form game. Dynamic consistency is discussed in Kydland and Prescott (1977).

CHAPTER 3

Information Revelation, Disclosure Laws, and Renegotiation

In the first two chapters, we looked at situations in which everyone was completely informed, except possibly about what decisions the other player had made. We now examine games of incomplete information, situations in which one player possesses knowledge that the other does not. This informational asymmetry itself can affect the way each player behaves; and legal rules can play a large role in determining how parties share information with each other. Indeed, many important legal reforms have focused on information and whether and how it is conveyed. Laws, for example, may mandate disclosure of information. Those who acquire more than 5 percent of the stock of a publicly traded firm must disclose their interest to other investors.[1] A company that intends to close a plant may have to give advance notice of the closing to its employees.[2] In the case of real estate sales, the doctrine of *caveat emptor* is giving way to laws requiring sellers to disclose whether the basement leaks or the neighbors are noisy. In other cases, the government limits the transfer of information to prevent discrimination or protect rights of privacy. Laws exist, for example, that make it illegal for an employer to inquire whether an applicant is disabled. Prospective students at a federally funded educational institution may not be asked their marital status. In this chapter and the next, we explore the kinds of effects that rules governing information may have on the way people interact with one another.

The following is an example of a problem arising from asymmetric information. A buyer knows something about a piece of land that the seller does not. The seller is a farmer and the would-be buyer is a geologist. The knowledge is whether the land has valuable minerals on it. The person who lacks knowledge, the farmer, in our example, must act notwithstanding the uncertainty. That person, however, does not act arbitrarily. First, a person can draw on previous experience and weigh the different possibilities. The farmer might not know whether there is oil or other minerals on the land, but nevertheless might have some sense of the probabilities. A farmer in Texas might think that there is one chance in ten that there is oil on the land, but a farmer in Illinois might put the chance at only one in a hundred. Second, the person who lacks information can draw inferences from the way another person acts. A farmer who starts with the belief that the land is unlikely to have oil on it might think it more likely that the land contains oil if an oil company geologist comes and asks if it is for sale.

When one person has information that another does not, this asymmetry itself affects how both parties behave. Farmers will want to know the identity of their prospective buyers. Potential buyers will, to the extent possible, conceal information (such as their training in geology) that tends to increase the price. Equally important from our perspective, the kinds of legal rules that are in place affect the kinds of inferences that parties can draw. In this chapter, we are concerned with how two parties behave when the informed party has the ability to convey crucial information to the other if that party chooses to do so. We then go on to examine situations in which both parties are equally well informed, but neither has the ability to convey to a court what is known. We postpone until the next chapter the case in which again only one player is informed, but that party has no ability to convey the information to the other directly. Anything the other player learns is learned by drawing inferences from the actions that the first player takes. Before we can confront any of these problems, however, we must introduce a new solution concept.

Incorporating Beliefs into the Solution Concept

In this chapter, we use the extensive form game to model the interactions between parties when one has information that the other does not. To do this, we need to develop a new solution concept, known as the *perfect Bayesian equilibrium.* This solution concept builds on an idea that we already encountered in our discussion of iterated dominance in Chapter 1. When we examined games in which we used the repeated elimination of dominated strategies, we had to posit the beliefs that each player would have. We found a solution by positing that each player believes both that the other player will not play dominated strategies and that the other player shares these same beliefs. When we

encounter situations in which parties are incompletely informed, we must also identify ideas that we can use to predict the beliefs (and thus the actions) of the players. These ideas are largely based on two principles. First, rational players should change their beliefs in light of the actions that other players take; and second, they should act in a way that is consistent with their beliefs. Moreover, in an equilibrium, the beliefs of the players should be consistent not only with their own actions, but also with those of other players.

In the previous chapter, we developed the concept of subgame perfection, the idea that the actions of the players were Nash, not only in the game as a whole, but also in every subgame. We are using the idea of "perfection" in a parallel sense here. One looks not at whether actions are optimal in subgames, but rather at whether actions are optimal given the beliefs of the players. A perfect Bayesian equilibrium is "perfect" in the sense that the actions are optimal given not only the actions of the players, but also the players' beliefs. A proposed solution is suspect if it requires one player to have beliefs that are inconsistent with the actions that player takes or the actions that other players take. A solution to a game should not assume that a party harbors such beliefs, just as a solution should not assume that a player consistently takes a course of action that is less than optimal given the actions of the other players.

We can illustrate the intuition behind these ideas by recalling the movie *The Maltese Falcon.* One of the principal figures in the movie is Kasper Gutman, played by Sydney Greenstreet. He spends seventeen years tracking down a gem-encrusted statue of a falcon that the knights of Malta had once offered in tribute to the king of Spain. He finally finds it in Istanbul in the hands of a Russian general named Kemidov. The statue has been covered with black enamel and appears to be a curiosity of only modest value. Greenstreet tries to buy the apparently worthless statue, but Kemidov refuses to sell it. Two of Greenstreet's confederates (played by Peter Lorre and Mary Astor) then try to steal it from the general. At the end of the movie, they discover that the statue they stole was only an imitation that the Russian general had substituted for the original after Greenstreet had offered to purchase it.

Greenstreet, of course, had not told the Russian general that the statue was gem-encrusted, but the general had inferred from Greenstreet's eagerness to buy the statue that it was of great value. Greenstreet made the mistake of bidding too much. As Peter Lorre says to Greenstreet in exasperation in one of the film's best moments, "You!

It's you who bungled it! You and your stupid attempt to buy it! Kemidov found out how valuable it was. No wonder we had such an easy time stealing it."

If we were to model the interaction between Greenstreet and Kemidov as a game, Greenstreet's strategy space might be the offers he could make. Kemidov's strategy space is whether to accept or reject any offer. A solution for this game, however, should take into account Kemidov's assessment of the chance that the falcon he possesses is the genuine, gem-encrusted Maltese Falcon and how he changes that assessment in light of the offer that Greenstreet makes. Kemidov will accept only an offer that exceeds the value to him of keeping the falcon, given his beliefs about the chances that the statue is gem-encrusted. Moreover, the beliefs that Kemidov has about the value of the falcon must be updated in light of the offer that Greenstreet makes. For example, if Greenstreet appeared and offered $100,000 for a statue that Kemidov had previously thought was worth only $10, Kemidov would reevaluate his original assessment of the worth of the statue. He would infer not only that it was worth more than $10, but maybe that it was worth even more than $100,000. After all, Greenstreet would make such an offer only if he knew something about the statue that Kemidov did not. Greenstreet has every incentive to offer Kemidov less than the statue's true value.

Our model of the game would have to begin with Kemidov's initial assessment of the likelihood that the statue was something other than an uninteresting ceramic figure of a bird. Kemidov starts with the belief that the chance that the statue is valuable is quite low. (We know this because he does not think it worthwhile to inspect the statue more closely or have it appraised before Greenstreet comes on the scene.) In addition to this background assumption, we want to specify how Kemidov should update his beliefs in response to Greenstreet's offer. There are essentially two kinds of offers that Greenstreet can make: offers that he would make for the statue if it were gem-encrusted, and offers that he would make regardless of whether the statue were gem-encrusted or a simple ceramic figure.

Let us assume for the moment that, if the statue were gem-encrusted, Kemidov would place a higher value on it than Greenstreet would. In other words, if Kemidov knew that the falcon were gem-encrusted, he would not sell it. What can we say about the likely outcome of this game? Given that Kemidov would not sell the falcon if it were gem-encrusted, Greenstreet would be able to buy it from him only if he made an offer that he would be willing to make (and Kemidov knew he was willing to make) even if the statue were exactly what it appeared to be. The likely course of play cannot be one in which Greenstreet makes an offer that he would make (and that Kemidov knows he would make) only if the statue were gem-encrusted.

In any equilibrium, Kemidov's beliefs have to be updated in light of Greenstreet's offer, and Kemidov must act optimally, given those updated beliefs. If the offer is one that Greenstreet would make only if the statue were gem-encrusted, and if Kemidov knew this, Kemidov should no longer believe that the statue is very likely a ceramic figure of little value. Greenstreet "bungled it," in Lorre's words, because he offered too much and thus unintentionally conveyed information about the value of the statue.

The equilibrium in a game in which a player has private information tends to be one of two general kinds. First, there is a solution in which the outcome is the same regardless of the information the other player has. (Greenstreet makes an offer that he is willing to make regardless of whether the statue is gem-encrusted, and Kemidov is willing to take such an offer.) Such a solution is a *pooling equilibrium*. The offer in such a case communicates no information—Kemidov assesses the offer in light of his own initial beliefs about the falcon and the chance that it is anything other than it appears.

The second possibility is one in which a player is better off taking an action, even though the action communicates information. If, contrary to what we have been assuming, Greenstreet puts a higher value on the falcon than Kemidov if it is gem-encrusted, but not otherwise, the outcome might be one in which Greenstreet makes an offer only if the falcon is gem-encrusted, and the offer would be accepted. Such a solution is a *separating equilibrium*.

The Perfect Bayesian Equilibrium Solution Concept

We want to formalize the basic ideas that we discussed in the last section. We do this by setting out the perfect Bayesian equilibrium concept in the context of an extensive form game. We use another variation on *Valley Lea*, the case involving the dairy and the food processor that we discussed in the last chapter. The word "Bayesian" in the name incorporates the idea that uninformed players put probabilities on different events and then update them using Bayes's rule when other players take actions that convey information. (In probability theory, Bayes's rule provides a means to capture formally the way rational people should update their beliefs in the wake of new information.)

The word "perfect" reflects the idea that beliefs and actions have to be consistent with each other.

The intuition behind this solution concept, as we suggested in our example drawn from *The Maltese Falcon*, is straightforward. A player who is uncertain about what another player has done nevertheless has beliefs, based, perhaps, on previous experience. In addition, these beliefs are updated in light of new information. A solution to a game should take these beliefs and a player's ability to update them into account. As stated, the beliefs and the strategies of the players should be consistent with each other.

A proposed solution to a game is suspect if it depends on one of the players having beliefs that are inconsistent with the actions that the players take in equilibrium, or if it requires players to take actions that are inconsistent with their beliefs. We can test whether a proposed equilibrium is a perfect Bayesian equilibrium in much the same way we tested for a Nash equilibrium. We ask whether, in the proposed equilibrium, a player's actions are a best response, given that player's beliefs and the actions and beliefs of the other players.

Assume that Dairy must decide whether to use dry milk itself or sell it to Processor. Whether it uses the milk itself or sells it, Dairy is exposed to potential tort liability if the milk is tainted. The benefits of selling the dry milk to Processor therefore turn both on how much testing Dairy does and how much testing Processor does. Dairy moves first and has three choices. It can decide to process the milk itself and enjoy a profit of $3. If Dairy decides to process the milk itself, Processor does not get to move, and it receives a payoff of $0. If Dairy sells the milk, it must either test low or test high. Processor in turn must test low or high. If both test high, the tests are redundant. They will both spend so much on testing that each will enjoy profits of only $1. If Dairy tests high and Processor tests low, Dairy still enjoys profits of only $1, but Processor saves on the costs of testing and earns profits of $3. If Dairy tests low and Processor tests high, their positions are reversed—Dairy enjoys profits of $5, but Processor enjoys profits of only $1. If both test low, the ensuing tort liability reduces both their profits to $0. Figure 3.1 illustrates the extensive form of the game when Processor knows how much testing Dairy has already done.

In this game, Processor knows what kind of test Dairy has run when it moves. This game has multiple Nash equilibria. One arises from the strategy combination in which Dairy sells and tests low, and in which Processor tests high in the event that Dairy tests low, and tests low in the event that Dairy tests high. This solution is Nash because both

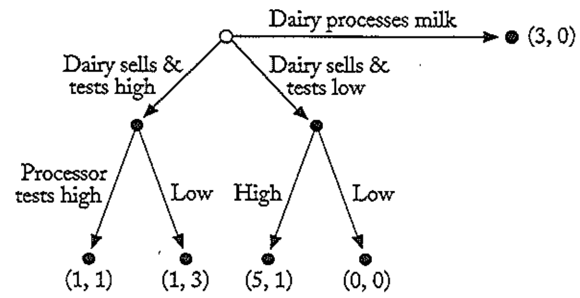

Figure 3.1 Sale v. production in the firm (Processor knows whether Dairy tested high or low). *Payoffs:* Dairy, Processor.

Dairy and Processor are playing a best response, given the strategy of the other. It is, of course, easy to see that Dairy can do no better. Under this strategy combination, Dairy enjoys a payoff of $5, its highest possible payoff. Seeing that Processor will not want to deviate is similarly straightforward. When Dairy tests low, Processor cannot improve its position by deviating from the strategy of testing high. It is better off testing high (and enjoying $1) than testing low (and receiving $0).

Another strategy combination that is Nash is the one in which Dairy processes the milk and in which Processor tests low when Dairy sells the milk. Given that Processor moves low, Dairy is better off processing the milk itself. Because Dairy is going to process the milk itself, Processor cannot improve its lot by deviating from testing low—and instead testing high—in the event that it receives the milk.

As we saw in the last chapter, we can use backwards induction to solve this game instead of the Nash equilibrium concept. We focus on the player who moves last—in this case, Processor. If given a chance to move, Processor will find it in its self-interest to test low if Dairy tests high (and thereby receive $3 instead of $1), but to test high if Dairy tests low (and thereby receive $1 instead of $0). Hence, when Dairy decides on its move, it can predict how Processor will respond to its strategies of testing high and testing low respectively. With this prediction in hand, Dairy will choose to sell and test low (and receive a payoff of $5), rather than sell and test high (and receive a payoff of $1) or process itself (and receive a payoff of $3).

We are able to use backwards induction in this game because Processor knows what kind of test Dairy ran when it moves. (The extensive form in Figure 3.1 tells us this because the left- and right-hand nodes at Processor's moves are not connected with a dashed line, indicating that they are in separate information sets.) Solving the game, however, is less straightforward if Processor does not have this knowledge. Consider the same game again, except that Processor does not know what kind of test Dairy ran. This new game is set out in Figure 3.2. The dashed line connecting the two nodes at which Processor might find itself on its move shows that Processor does not know whether it is at the left-hand node or the right-hand node when it moves.

We cannot use backwards induction to solve this game. The players make their testing decisions independently of each other, if they test at all. In addition, this game has multiple Nash equilibria. Hence, we cannot rely simply on the Nash solution concept to solve this game either. We cannot use the idea of subgame perfection. There is no proper subgame other than the game as a whole. There is no information set that contains a single node, other than the initial node. One cannot isolate part of the game and posit the strategy choices of the parties once the game reaches that node.

Because knowledge of Processor's position is not built into the game, we have to focus on the beliefs that it has about its position. When we incorporate beliefs into the solution concept, we can eliminate some Nash equilibria as solutions to the game. We can make some plausible assumptions about these beliefs, about how actions of the players affect these beliefs, and about how they in turn affect the actions of the players. If, for example, Processor believes that there is a 50–50 chance that Dairy tests high, it should test low. Testing low gives Processor an expected payoff of $1.5, whereas testing high gives it a payoff of only $1.[3]

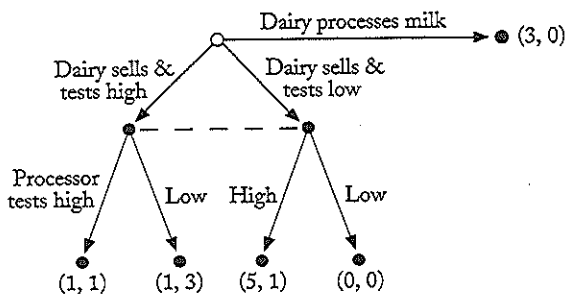

Figure 3.2 Sale v. production in the firm (Processor does not know Dairy's test). *Payoffs:* Dairy, Processor.

Any equilibrium in which Processor believes that Dairy is as likely to test high as to test low is one in which Processor must test low. Any other course is not rational for Processor given its beliefs. The Nash equilibrium concept requires that a player's strategy be optimal given the strategies of the other players. Once we incorporate beliefs into our solution concept, we also require that a player's strategy be consistent with that player's assessment of the probability of being at each node in the information set. These two ideas together are known as the requirement of *sequential rationality*.

As with the case of *The Maltese Falcon*, we must ensure that a player's assessment of the chances of being at any given node is consistent with the equilibrium action of the other players. Just as we do not think that players will embrace an equilibrium in which one player adopts a strategy that is suboptimal given the strategy of another, we should think that, in equilibrium, a player's beliefs are going to be consistent with the actions of the other player. Any equilibrium, for example, in which Processor believes that Dairy is as likely to test low as to test high, must be an equilibrium in which Dairy is in fact equally likely to test low as to test high in the event that it sells the milk. In equilibrium, the strategy choices of each player should be the best responses to the strategy choices of the others. The beliefs of the players should also be consistent with the strategy choices the other players have made.

In the game in Figure 3.2, there are multiple perfect Bayesian equilibria. Consider the case in which Dairy processes the milk itself, Processor tests low, and Processor believes that, if Dairy sells the milk, there is a 40 percent chance that Dairy tests high. This combination of strategies and beliefs forms an equilibrium. Dairy is behaving optimally given Processor's strategy of testing low. Processor's strategy is sequentially rational, given its belief that there is a 40 percent chance that Dairy will test high.[4]

There is nothing magical about Processor's belief that Dairy has a 40 percent chance of testing high. We would reach the same conclusion if Processor believed that there was a 35 percent or a 45 percent chance that Dairy would test high. Indeed, a perfect Bayesian equilibrium exists whenever Dairy processes the milk, Processor tests low, and Processor believes that there is more than a 33 percent chance that Dairy will test high. The belief that the chance of testing high is 40 percent, like the belief that it is 35 percent or 45 percent, is not inconsistent with Dairy's actions because Dairy never sells the milk in the proposed equilibrium. Processor never has any action that it can use to update its starting assessment, and this solution concept puts no restrictions on that starting assessment. This combination of actions, however, does depend on Processor's believing that there is at least a 33 percent chance that Dairy will test high. If Processor believed that the chance that Dairy would test high was less than 33 percent, Processor's best response, given this belief, would be to test high rather than low.

Another perfect Bayesian equilibrium in this game is one in which Dairy tests low, Processor tests high, and Processor believes that, in any case in which Dairy sells the milk, it will test low. Dairy's strategy is optimal because it is receiving the highest possible payoff. Processor is acting in a way that is sequentially rational given its belief that Dairy has tested low and given that Dairy has, in fact, tested low. Finally, Processor's belief is consistent with Dairy's equilibrium strategy of testing low.

Bayes's rule requires a player to update beliefs only in the wake of actions that take place in equilibrium. If an action is not taken in equilibrium, there is no new information that a player can use to update beliefs. Bayes's rule therefore does not place any constraints on a player's beliefs about actions that are taken off the equilibrium path. In the last chapter, we refined the Nash equilibrium concept in order to eliminate strategy combinations that rest on threats that are not credible. We can refine the perfect Bayesian equilibrium concept in the same way to ensure that a player's beliefs are plausible, even when those beliefs concern actions that never occur in equilibrium.

Consider again the equilibrium in which Dairy processes the milk itself and Processor believes that, were Dairy to sell the milk, there is again a 40 percent chance that Dairy will test high. Bayes's rule does not constrain Processor's belief because the case in which Dairy sells the milk is off the equilibrium path. Once an action is off the equilibrium path, nothing in the perfect Bayesian equilibrium solution concept constrains Processor's beliefs. Nevertheless, it seems implausible that Processor believes that Dairy tests high when it sells milk. Dairy always does better by processing the milk itself than it does by selling the milk and testing high. Dairy's strategy of processing its own milk dominates the strategy of selling and testing high. Processor is not likely to have a belief that requires Dairy to play a dominated strategy.

Processor's beliefs are likely to be molded by both Bayes's rule and the principle that people do not play dominated strategies. Processor should therefore infer that whenever Dairy sells the milk, it will test low. Once we posit that Processor has this belief, we can solve this game. The only perfect Bayesian equilibrium that survives this refinement is the one in which Dairy sells the milk and tests low, Processor tests high, and Processor believes that Dairy tests low when it sells the milk.

Verifiable Information, Voluntary Disclosure, and the Unraveling Result

The perfect Bayesian equilibrium with refinements is a tool that we can use to analyze how parties might draw inferences from the actions of other players. In this section, we want to examine these issues in the context of laws governing the disclosure of information, such as a law that requires a seller to disclose all known defects in a product. At the outset, we must recognize that there are different kinds of information. Some information is *verifiable;* that is, it can be readily checked once it is revealed. For example, the combination to a safe is verifiable information. The combination either opens the safe or it does not. Other information is *nonverifiable.* An employer wants to know whether a guard who was hired was vigilant. The employer might be able to draw inferences from some events, such as whether a thief was successful or was caught, but such information may not be available and, even if it is available, may not be reliable. Even a lazy guard may catch a thief, and a thief may outwit even the most vigilant guard.

Legal rules governing information have to take into account whether verifiable or nonverifiable information is at issue. When we analyze laws that require disclosure of information, we must confine ourselves to information that can be verified after it is disclosed. Moreover, we must limit ourselves to cases in which a court can determine whether a player possesses the relevant information. Courts cannot sanction parties for breaching a duty to disclose information if they have no way of telling whether a disclosure is truthful. Similarly, a court cannot sanction a party for failing to reveal information if it has no way of determining whether a party possesses the information in the first place.

At the outset, we want to focus on information that, once revealed, can be verified, both by the other players in the game and by a court that must subsequently apply the legal rule. We then examine situations in which the issue is slightly more complicated. We explore first an important principle in the economics of information known as the *unraveling result.* Consider a seller with a box of apples that has been sealed. The box can hold as many as 100 apples. The seller knows the number of apples in the box. The buyer does not know the number of apples in the box and has no way of counting them before the sale. The buyer, however, does know that the seller knows the number of apples in the box. After the buyer purchases the box and takes delivery, the buyer can count the apples. If the seller lied about the number of apples in the box, the buyer can sue the seller at no cost and recover damages. The buyer has no action against the seller who remains silent. Once we make these assumptions, we can conclude that all sellers will accurately and voluntarily disclose the number of apples in the box, no matter how few there are.

Focus on a seller with 100 apples. This seller will disclose. Because the information is verifiable and because of the legal remedy, the seller's disclosure will be believed and the seller will be able to charge the price for 100 apples. If the seller does not reveal the number of apples in the box, there is no way that the seller can persuade the buyer to pay the price of 100 apples. As long as the buyer believes that there is some chance that there are fewer than 100 apples in the box, the buyer will be unwilling to pay for 100. Hence, the seller with 100 apples discloses.

Now consider whether an equilibrium exists in which the seller has 99 apples and says nothing about how many apples are in the box. The seller would find remaining silent a good idea only if the buyer believed that there was some chance that there were 100 apples. The buyer, however, will not hold this belief in equilibrium because, as we have just seen, in equilibrium the seller with 100 apples discloses. If the seller remains silent, the buyer will infer that the seller has 99 apples or fewer. By remaining silent, however, the seller with 99 apples becomes lumped with sellers with very few apples.

This seller, knowing that the buyer will never pay the price of 100 apples, prefers to reveal that there are 99 and receive that price rather than one that reflects the chance that there may be fewer than 99 apples in the box. Because again the information is verifiable and the buyer has an effective legal remedy, the buyer will know that the seller is telling the truth. The seller with 98 apples goes through the same reasoning process, as does the seller with 97 apples, and so forth. This is the unraveling result. Silence cannot be sustained because high-value sellers will distinguish themselves from low-value sellers through voluntary disclosure. In the end, all sellers disclose their private information.

Whenever we examine any laws governing disclosure of information, we need to be aware of the possibility of unraveling. Cases involving the Fifth Amendment's privilege against self-incrimination illustrate this point. The Fifth Amendment to the United States Constitution provides that "[n]o person . . . shall be compelled in any criminal case to be a witness against himself." The privilege against self-incrimination is one of the most basic rights of criminal defendants. This right would be a limited one, however, if the failure to testify were to lead to an inference against the defendant. It was not until 1965, in *Griffin v. State of California*,[5] however, that the Supreme Court held that the self-incrimination privilege barred comment on the failure to testify. It was not until 1981, in *Carter v. Kentucky*,[6] that the Court held that the right required an instruction to the jury, on the defendant's request, that no inference be drawn from the failure to testify:

> A trial judge has a powerful tool at his disposal to protect the constitutional privilege—the jury instruction—and he has an affirmative constitutional obligation to use that tool when a defendant seeks its employment. No judge can prevent jurors from speculating about why a defendant stands mute in the face of a criminal accusation, but a judge can, and must, if requested to do so, use the unique power of the jury instruction to reduce that speculation to a minimum.[7]

Given the logic of unraveling—that someone with information will disclose it, rather than be subject to the inference that arises from the failure to disclose it when one can do so—the privilege against self-incrimination becomes meaningless unless steps are taken to prevent the adverse inference from being drawn. As the Court recognizes, however, there may be insurmountable limits on the ability of the legal system to prevent inferences from being drawn from silence. A jury instruction may make the problem worse because it may alert the jurors to their ability to draw inferences from a defendant's failure to testify. To be sure, we can infer that this jury instruction aids those who do not testify if defense lawyers regularly ask for such an instruction. Nevertheless, we may doubt, given the ability of the jury to draw inferences notwithstanding the instruction, whether the jury instruction is in fact "a powerful tool."

Because the person who holds favorable verifiable information has an incentive to reveal it, the allocation of the right or duty to inquire or disclose should not affect whether verifiable information is revealed. The idea that verifiable information may be disclosed voluntarily has important implications. It calls into question two standard legal approaches to revelation of information—inquiry limits and disclosure duties. We examine each in turn.

Inquiry limits attempt to prevent decisionmakers from obtaining information thought to be an inappropriate basis for making a decision. For example, the Americans with Disabilities Act bars an employer from asking whether an applicant has a disability and also bars preoffer medical tests.[8] A similar restriction applies regarding inquiries of applicants to educational institutions.[9] There are also limits on the kinds of inquiries that can be made of an applicant to rent or purchase a dwelling.[10]

Prohibiting questions about disabilities is a prominent form of inquiry limits, but others are common as well. Regulations implementing Title IX of the Education Amendments of 1972 forbid inquiry into the marital status of an applicant for employment at or admission to a school receiving federal funding.[11] Many states bar preemployment inquiries into religious or political affiliations for prospective public school teachers or for all prospective state employees.[12] Illinois bars inquiries into whether a prospective employee has filed a worker's compensation claim.[13] Inquiry limits are especially common under rules of evidence. The general protection for privileged matters—usually matters between attorney and client, physician and patient, or spouses, for example—is a form of inquiry limit, as are rules forbidding questions regarding a victim's prior sexual history in rape trials.[14]

Inquiry limits, however, may be ineffective unless there is some mechanism that prevents voluntary disclosure of the information. Barring an employer from asking whether an applicant has a disability might not affect whether the employer learns of the disability. When applicants who are not disabled also know the legal rule, they can disclose this information when it is verifiable.

As soon as the healthiest applicants disclose the results of their medical tests, the slightly less healthy ones may follow suit. The inability of the employer to require a medical test before making an offer may be irrelevant if applicants know that their employer wants the information, but cannot ask for it. They can simply volunteer it. Information problems can turn our usual intuitions upside down. We generally think that parties must know the legal rule for it to be effective. In this context, a legal rule barring inquiry might work best if the applicants did not know the legal rule. Knowledge of the legal rule might itself suggest the importance of the information to the employer.

The unraveling problem may give us some guidance about how inquiry limits might be made effective. Assume, for example, that the information is such that, if a job applicant volunteered it, the employer could not readily ascertain whether the applicant was telling the truth, but might be able to make that determination after the fact. (Applicants, for example, might volunteer that they never drank alcohol and had never had a drinking problem. Although the employer might have no ability to verify this information at the time the applicants made these statements, the employer might well learn that such statements were false at some subsequent time.) In such a case, a law that forbids inquiring might work effectively only if the law either provided some means to prevent applicants from volunteering the information or provided some means of keeping the employer from responding adversely upon discovering that an applicant had lied.

A rule that prevents the employer from retaliating against applicants who lied is another example of how a legal rule can change the way individuals behave even though it attaches consequences to actions that do not appear in equilibrium. Because applicants' assertions that they never drank would no longer be credible, applicants would no longer make them; or, if made, such assertions would no longer convey information. Once the employer cannot retaliate against the liar, the lies will not be made in the first place. A law forbidding employers from retaliating against workers who lied might be necessary to ensure that a law limiting an employer's ability to inquire into information operated effectively. It would also ensure that, after applicants were hired, they could disclose information, thus allowing the employer to accommodate their disabilities.

These problems with inquiry limits appear elsewhere as well. A law forbidding an inquiry into political affiliation seems to serve little purpose, as long as applicants for city jobs know that those hiring want employees from a particular political party. Applicants from that party will make their sympathies known, and those doing the hiring can draw an inference from silence. Rules limiting the transfer of verifiable information should be two-sided. In other contexts, two-sided bars are used. For example, it is unethical both for an attorney to condition an offer of settlement of a suit against a client on opposing counsel's agreement not to sue that client again[15] and for an attorney to accept such an offer.[16] As a general matter, however, statutes with inquiry limits are one-sided, and, in cases of verifiable information, such rules are potentially defective.

Laws requiring disclosure of information raise similar problems. An important issue in labor law concerns the employer's duty to convey information to the union. The Supreme Court held in *NLRB v. Truitt Manufacturing Co.* that an employer who claims financial hardship as a reason for not increasing pay must be willing to back up its claim.[17] Failure to provide the relevant information may support a finding that the employer has not bargained in good faith. As the Court noted, if a claim of financial hardship "is important enough to present in the give and take of bargaining, it is important enough to require some proof of its accuracy."

Strikes may arise as a result of private information. A union may be willing to incur the costs of a strike as a way of distinguishing between the high-profit firms that can afford a higher wage and the low-profit firms that cannot. If the high-profit firm loses more when there is a strike, it may be willing to settle earlier at a higher wage. The union could then infer that those firms that refuse to settle are not losing much from the strike because they are low profit.[18] If the legal rule had the effect of revealing the types of all firms at the outset, there would be no need for a strike.

Whether we can justify the rule of *Truitt* on this basis, however, is not obvious. As long as the private information was verifiable, low-profit firms would voluntarily disclose the information. High-profit firms that claimed low profits, but did not support the claim, would not be believed. The legal rule of *Truitt* would not affect the amount of information conveyed. With or without the rule, low-profit firms would disclose information and high-profit firms would not. *Truitt* would affect what high-profit firms said in the bargaining. They would have to be careful not to claim financial hardship, but this in itself would not make any difference. Unsupported statements of the high-profit firms would not be believed regardless of the legal rule. Justifying *Truitt* seems to require in the first instance confronting the question of why the unraveling result may not hold. (We might, for example, ask whether it is costly for low-profit firms to disclose financial information.)

There are other ways in which the unraveling principle informs our understanding of disclosure laws. For example, the Securities and Exchange Commission (SEC) has recently retreated somewhat from the strict disclosure requirements for issues that are not large and not widely held, in part because of the way the unraveling principle works in the absence of a legal rule. Sophisticated investors will infer bad news from an issuer's failure to disclose relevant financial information. Therefore, the incentives of the issuer would appropriately trade off the costs of information acquisition with the value of information.[19]

The unraveling result may also help to justify some rules that forbid the disclosure of information. Banks and bank examiners are banned from discussing bank examination reports publicly. Bank examiners gather information about a bank's condition and then use this information in their negotiations between a bank and regulatory authorities to fix problems before the risk of failure becomes severe. If banks were free to make reports public, banks would, on average, be made worse off because of unraveling. Without a law forbidding disclosure, the bank that receives a good examination report will want to make it public; it will raise the bank's stock price and reduce the probability of a run. Unraveling will begin, however, and depositors will in turn infer that any bank that does not reveal its report received a bad one. Although a bank in a relatively weak financial condition might have been able to work through its problems, it will not survive when depositors infer that there is a small risk of failure and then rush to take their money out. Disclosure, in other words, gives only a small benefit to good banks, does nothing one way or the other for bad banks (which will fail anyway), but harms banks that fall in the middle. Because of the unraveling principle, the law works only if limits are placed on a bank's ability to talk about a report, regardless of whether it is favorable.

Disclosure Laws and the Limits of Unraveling

The unraveling result depends on the ability of one player to infer the other player's information from that player's silence. Players are able to reveal only the information that they acquire. This linkage between acquisition and revelation of information may itself prevent unraveling. Unraveling may not occur (or will not be complete) if there is a chance that a player has never acquired the relevant information. In such a case, one will not be able to tell whether players are silent because they do not have the relevant information or because they have the information but do not wish to reveal it. The inability to draw inferences from silence in these cases might seem to justify some legal rules. Mandatory disclosure laws, however, may not work in this context. A law that mandates disclosure requires a court to distinguish players who withhold information from those who do not have it. If we need such a law only when the uninformed player cannot draw such a distinction, we must explain why the court is able to do something that a player cannot. The question, in other words, is whether a court can subsequently determine whether a player who claimed ignorance was telling the truth and inflict a penalty large enough to ensure that no one would have an incentive to lie.

Instead of a case in which both players know that one possesses the combination to a safe, we have a case in which one player may or may