

3

Utilizing Strategic Ignorance in Negotiations

Sarah Auster and Jason Dana

Abstract

This chapter analyzes the role of information in strategic decision-making settings. It considers several situations in which it could be individually advantageous to deliberately ignore information, particularly when this ignorance can be signaled to the other parties in the decision, and introduces purely psychological reasons why a negotiating party might want to ignore information. In some situations, information actually constrains the action set available to the individual. Examples involve inadvertently leaking private information to the other side, knowledge triggering one's own moral constraints, and knowledge biasing the individual in ways that will harm the negotiation. Even if information acquisition is completely private, a behavioral agent will sometimes negotiate better by deliberately avoiding information.

Introduction

In the game of American football, there is a surprise play called the “naked bootleg.” Often, when a team is very close to the goal line, the quarterback takes the ball and hands it off to a back who has a running start toward the defensive line. That player runs behind a group of blockers and tries to bash his way through the defense to score. In the naked bootleg, everything is set up to look like a hand off, except that after the snap, the quarterback fakes and keeps the ball, rolling out in the opposite direction to the blockers, thus exposing himself as “naked.” The entire play is predicated on deception: If the defense correctly reads that the quarterback is keeping the ball, the play will most likely fail.

In a 2013 matchup between the Denver Broncos and Dallas Cowboys, Peyton Manning (quarterback for the Broncos) executed a perfect example of the play to score a touchdown—his first rushing touchdown in 62 games

(Petchesky 2013). In an interview following the game, Manning revealed an interesting bit of strategy: he had told the offense that he would hand the ball off to the running back (i.e., they were informed about the wrong play). He told only the running back that he would not be handing off the ball, as Manning did not want the back to panic and try to rip the ball away when the handoff failed to materialize.

Why would a team want its own players ignorant as to what it is doing? In this case, the obvious answer is that if they did not truly believe that they were blocking for a traditional running play, they might not be able to fool the defense into thinking it was a traditional running play. Knowledge of the real play was useless to the rest of the offensive players because they would execute as if it were a traditional running play in either event. To the extent that they were psychologically incapable of ignoring useless knowledge, or incapable of preventing themselves from conveying it to the other team, that knowledge could be harmful if it made their performance less convincing. Ignorance made them better able to perform in a strategic interaction. Rational players in this situation should choose to be ignorant.

In this chapter, we explore the value of deliberate ignorance in negotiations and related areas. Specifically, we suggest that there are situations in which knowledge effectively limits an agent's action space, and thus the agent is better off without this knowledge. A negotiation is defined as a decision-making process by which two or more parties agree how to allocate scarce resources. Thus, the situations that we consider here are characterized by strategic interaction between parties who cannot have everything they want. This definition allows that parties may be asymmetrically informed. It does not include as negotiations, however, situations in which one party can unilaterally determine the outcome for all parties (see Dana et al. 2007), or where one's own payoff is not determined by the joint decisions of the self and others. While some theoretical research has documented situations in which ignorance brings strategic advantage in bargaining, less research has been conducted into whether people do use ignorance in such situations. Even less research has been conducted into information avoidance that, for purely psychological reasons, could lead to better outcomes in a negotiation.

We begin by summarizing some reasons why a standard decision theoretic agent would exercise deliberate ignorance in a negotiation, and descriptive research into whether people actually avoid information in these situations. Then, we speculate on psychological reasons why a person in a negotiation would benefit from ignorance apart from reasons given by standard decision models. To be clear, the subject matter will not be information that one defers for the purpose of also keeping others from receiving it. Rather, these are situations in which the decision maker is made worse off by having the information for their own use.

Strategic versus Individual Decision Making

Recent papers have catalogued a variety of situations in which an individual decision maker benefits by ignoring or even refusing costless information (Golman et al. 2017; Hertwig and Engel 2016; Kadane et al. 2008). This effort adds to our understanding of decision making on its face because standard decision theoretic models, like those which form the foundations of modern economic analysis, assume that information has a nonnegative value. The same is not true, however, of standard game theoretic models. When decision makers' payoffs are determined by their own behavior as well as by the behavior of others, information can be detrimental. For example, if Player 1 knows something that Player 2 does not, and Player 2 understands this, an *information asymmetry* exists between players. Such asymmetries can lead to deficient outcomes such that even the player "advantaged" with more information (Player 1) may prefer not to possess the information, or at least prefer that Player 2 not know that Player 1 has more information.

To see how information can damage outcomes in a strategic interaction, imagine playing a single round of a game with one other person. You announce first "red" or "black" after which the other player announces "red" or "black." A card is then turned over from a deck of playing cards. If both players have announced the correct color, both receive a payoff of 1. If neither has announced the correct color, both receive a payoff of 0. If one player has announced the correct color and the other has not, the player with the correct color gets 50 while the other player gets 0. If both parties understand the game, you might expect your opponent to announce the opposite color that you select, which would greatly increase the expected payoffs for both of you. (Imagine a payoff much larger than 50 if your intuition doesn't match.) Now imagine that before announcing your color, you get to look at the card that will be flipped, but your opponent cannot. What would you do? There is no personal gain in knowingly announcing the wrong color; you would simply receive zero. Knowing that, your opponent might reason that it is best to announce the same color that you do. This extra bit of knowledge has now effectively ruined your ability to coordinate with your opponent in a way that is good (in expectation) for both parties. You would be better off if you did not view the card, or if your opponent thought you did not view the card.

The possibility for knowledge to ruin coordination is not limited to situations of asymmetric information. In social dilemmas, defection is no longer a clear equilibrium if it is common knowledge that there will be repeated play. Once it is common knowledge that the players are in their final round of play, however, there is no longer a reason to cooperate, and players are collectively disadvantaged. There are likewise many reasons why publicly ignoring information in a negotiation could be individually beneficial. Ignoring information can serve diverse functions: from strengthening one's bargaining position to

solving the “hold-up” problem to combatting adverse selection and solving principal-agent problems.

Normative Reasons for Deliberate Ignorance in a Negotiation

Why would a standard decision theoretic agent refuse some information during negotiations? Perhaps the first writer to suggest the use of deliberate ignorance as a tool in bargaining was Thomas Schelling (1956). Schelling considered bargaining problems in which one party wishes to convince another of something; for instance, a buyer who wishes to convey that she will not pay a seller more than X (i.e., her *reservation price* is X). Because such preferences are private, and parties are known to bluff in search of a good deal, there is a need to make such commitments credible. Deliberate ignorance is one method through which a bargaining party could communicate a credible commitment. For example, labor union leaders could publicly avoid meeting with their membership to signal to management that there is no intention to end the strike without a better offer.

In situations where it is difficult to write complete contracts, there is a well-known problem called the *holdup problem*. When one party has made a prior commitment to a relationship with another party, the latter can “hold up” the former for the value of that commitment. For example, if an automobile manufacturer developed an exclusive relationship with a firm that provided certain automobile parts for production, the parts supplier could change its prices in times of increased demand and the manufacturer would be put in a poor position to negotiate. The possibility of a holdup can lead to underinvestment in relationships that would otherwise be profitable.

The holdup problem is solved if the vulnerable party, the manufacturer in the above example, can keep its information (e.g., sales projections) private. Then, they could not be held up for the value of the surplus created by the agreement. Rogerson (1992) shows that a variety of solutions to the holdup problem exist when information is asymmetric and suggests that the party that has the bargaining advantage should precommit to allowing the vulnerable party to keep its information private. By doing so, the vulnerable party can trust that they will not be exploited and invest more. This idea was extended by Lau (2008) and Hermalin and Katz (2009), who developed specific conditions under which the party that would hold bargaining power does best to avoid learning the information that gives them power. The holdup problem is thus one in which too much bargaining power is, ironically, harmful because it can cause valuable deals to break down. Remaining deliberately ignorant is, therefore, a mechanism that can be used to cede some power, if that ignorance can be credibly communicated to the other party.

There are other contracting situations that provide an interesting context in which deliberately avoiding information can lead to better outcomes. Crémer

(1995) modeled the value of “arm’s length” relationships in principal-agent problems with moral hazard and renegotiation. The principal chooses only to observe the agent’s production (in this case, they are a firm and a supplier) and not to seek additional information about the causes of subpar production, thus not allowing for “excuses.” Ignorance in this case credibly commits the principal to not accepting excuses, even if they are good. These arm’s length relationships may be beneficial as they create better incentives and raise overall production. Roesler and Szentes (2017) analyzed the role of information in a bilateral trade problem between a monopolistic seller and a privately informed buyer. Their work shows that the optimal amount of information for the buyer is partial because if the seller knows the buyer is fully informed, the seller finds it optimal to charge a higher price.

Adverse selection (Akerlof 1970) introduces another situation in which parties would benefit by deliberately avoiding information, provided the choice to remain uninformed is publicly observable. Adverse selection occurs when sellers have private information about the quality of the goods they own. In the classic example, a used car may be of high quality (a “peach”) or low quality (a “lemon”), with buyers being uncertain of which they are getting. Buyers will thus be willing to pay a price that is an average of the different qualities. If sellers know which car they are holding, however, they will only find it profitable to sell lemons, driving peaches from the market and, ultimately, causing the market to break down: a buyer should not want to buy a car that a seller is willing to let go.

The same problem arises in insurance markets if clients have private information about their risk type. Berkman (this volume) considers the problem of genetic testing. If people knew their own genetic information and insurers did not, it could cause a breakdown in health insurance markets because the people wanting certain kinds of insurance would only be the people who were particularly likely to incur payouts (see Rothschild and Stiglitz 1976). If individuals are required to report to insurers whether they received genetic testing, not getting tested offers an advantage, even when the tests are free, because insurers would understand that the insurance is particularly attractive to people who tested positive and would avoid covering them. It is thus clear that if parties were uninformed on key dimensions, adverse selection problems can be circumvented and profitable trades allowed to happen. This solution, however, is only possible if the decision to obtain information can be observed by the insurance firms.

Descriptive Results on Deliberate Ignorance in a Negotiation

Do people make strategic use of ignorance in bargaining where the opportunities arise? The limited evidence from laboratory experiments suggests that they

do. Perhaps the most important evidence comes from Conrads and Irlenbusch (2013) and Poulsen and Roos (2010).

Conrads and Irlenbusch assigned subjects to play one of a variety of take-it-or-leave-it bargaining games where one player, the proposer, chooses an offer and the other, the responder, can accept the offer or reject it and leave both players with nothing. One of these offers was always better for the proposer, but whether it was better or worse for the responder was sometimes left uncertain to the proposer. When the proposer did not know the responder's possible outcomes, almost no offers were rejected. There would be no point in the responder punishing the proposer at a cost to self when the proposer did not even know which offer was fair to the responder. More interesting were cases where the proposer could choose to reveal this information. If the proposer's decision to reveal the information was public, several proposers chose not to reveal, and unfair offers were rejected at a lower rate. This result confirms Schelling's original conjecture that communicating deliberate ignorance could strengthen bargaining position. Here, it allowed the proposer to seek a favorable outcome while blunting any inference of intention to be unfair. As further evidence that this behavior was strategic, Conrads and Irlenbusch also ran a condition in which the proposer's decision to remain ignorant would be private and unknown to the responder, thus destroying the strategic aspect of ignorance. Few proposers in this condition decided to remain ignorant, confirming that indeed their ignorance was strategically deliberate in the public reveal condition.

Poulsen and Roos (2010) experimented with a similar Nash bargaining game. Pairs of players made demands for shares of a resource pie. If their demands did not exceed 100% of the pie, they got their demands. If they exceeded 100%, both players got nothing. One player was allowed to make the demand first, essentially transferring the game to a take-it-or-leave-it ultimatum that gives the first mover a strategic advantage (if the first mover demands more than half, the second mover will have to demand less or else leave with nothing). Poulsen and Roos allowed second movers to choose not to see the first mover's offer and to make this choice public before the first mover chose. Doing this would essentially transform the game to a simultaneous choice, where the focal equilibrium is a fifty-fifty split. After some practice, over 80% of second movers employed deliberate ignorance to enhance their bargaining power. Like Conrads and Irlenbusch (2013), Poulsen and Roos (2010) also used a private ignorance condition where the second mover could choose not to see the first mover's demand, but the first mover would not know it. When the decision to reveal was private, over 80% of second movers wanted to see the first mover's demand, confirming again that the deliberate choice of ignorance in the public reveal condition was strategic.

Although they do not directly investigate deliberate ignorance, experiments by Sloof et al. (2007) suggest that people would benefit from ignorance in a real holdup problem. They ran a laboratory experiment in which subjects assigned to the role of buyers could choose to pay a cost that made a seller's

goods more valuable, which thus created more surplus. The seller, however, set the price of the goods and the buyer had to transact if the exchange was profitable. The investment was thus at risk for a holdup; if sellers know about the investment, they can raise the price to capture all the excess benefit. The experiment manipulated whether the decision to invest was public or private. When the decision to invest was public, buyers anticipated holdup and invested less often, destroying welfare and lowering even sellers' earnings. Sellers were not given the opportunity to blind themselves to the investment in this study. The results played out, however, such that sellers would indeed benefit if they could publicly signal ignorance, as theory suggests.

Psychological Reasons for Deliberate Ignorance in a Negotiation

Whether individuals benefit from deliberate ignorance in negotiations for purely *psychological* reasons—that is, when information does not normatively change the structure of the bargaining task—is, at least empirically, a frontier topic. We are aware of little evidence that bears directly on this question. Owing to the ideas and results above, however, we can speculate that deliberate ignorance could be valuable in several ways when we admit a richer psychology on the part of the negotiating players. Specifically, we identify three reasons why information can actually serve to limit an agent's action space, and thus make the agent worse off:

1. Knowledge could unintentionally “leak” to another party.
2. Knowledge could invoke moral image constraints.
3. Knowledge could lead to self-serving bias in the agent's interpretation of what is fair, and fairness constrains actions.

Leaking Knowledge

As the chapter's opening example about calling a naked bootleg play in football suggests, deliberate ignorance could be useful in conveying a credible commitment to a course of action. As noted earlier, bargaining situations often entail the parties trying to convince each other of matters that privately held preference. Signals that one is “serious” become important in the process. As Schelling (1956) notes, “if a man knocks at a door and says that he will stab himself on the porch unless given \$10, he is more likely to get the \$10 if his eyes are bloodshot.” Short of making one's eyes bloodshot, how does one convey commitment to an outcome when there is otherwise reason to doubt the commitment?

Frank (1988, 2011) suggests that people evolve moral emotions to solve problems such as cooperation in one-shot social dilemmas or even bargaining from poor power in a take-it-or-leave-it ultimatum. He recalls a humorous

experience from a concert he attended in his youth: A dog walked up and, seeing a man who was in a drug-induced stupor, urinated on him. Would this dog have attempted that with any of the more alert concertgoers? Probably not, as it would be deterred by the possibility of receiving a swift kick. Similarly, how does one know to whom one should and should not make small ultimatum offers? Absent moral emotions, we might conclude that no matter what the responder claims, they would take something rather than nothing. But if we anticipate the moral emotion of anger in the responder, we might fear making a small take-it-or-leave-it offer because anger makes it satisfying to punish the offer, even at a small cost to the self.

The psychological problem that negotiating agents face is that moral emotions may be difficult to fake. Simply put, people can be bad emotional liars and, therefore, accidentally pass information to the other party. Trivers has written extensively about the value of self-deception in convincing others (for a recent summary, see Trivers 2011b). It is difficult to convince others of something one does not believe, and Trivers argues that self-deception evolved in humans to solve the problem of credibly signaling to others. Therefore, people have incentives to manage their own beliefs. Rather than considering all information and forming Bayesian posterior beliefs, they might be better served to seek information selectively to manipulate their beliefs.

We argue specifically that when information cannot be easily ignored, there is a risk of signaling one's private information to the other party in subtle or even unwitting ways. Thus, information changes the set of feasible actions and can have a negative value. Returning to the naked bootleg story, it is intuitively appealing that the outcome may be better if players do not know what play they are running, in part because the players need to react exactly the same, regardless of whether an actual goal-line running play or a naked bootleg is called. The reason for *lying* to them about the play, however, is best understood as preventing knowledge from being involuntarily signaled or leaked to the other team. Similarly, in bargaining situations, a potential buyer might be concerned about appearing too eager to buy the object and prefer not to learn the exact valuation before engaging in the negotiation. Unintentionally conveying one's beliefs and preferences can limit the possible outcomes in a negotiation. Deliberate ignorance through choosing one's informational signals can thus be an important self-disciplining device for negotiating with others.

Ironically, moral emotions themselves can limit the individual's action set in a beneficial way that negotiators may try to signal. A subject of lore in negotiations is that some people will have a fake angry phone call before speaking with someone in a potential negotiation. By showing the other party that one is angry, there is the thought that they will be afraid to push you too far toward your reservation value, as you might be "crazy" enough to reject a profitable but unfair offer if in a state of anger. All of which brings up an interesting question: Why should the man on the porch with bloodshot eyes or the negotiator who has lost control of his or her emotions get a better outcome than a

“rational” bargainer? Subjects with economics training give lower offers in ultimatum games, but they also accept lower offers (Carter and Irons 1991). This finding shows a perverse effect of “understanding” the game: if one were proposing an ultimatum to a student with philosophy training and a keen interest in social justice, it would be unwise to make a low offer. The true “economist” could then wind up worse off because she fears rejection from “stubborn” or “irrational” counterparts, yet can be relied upon to accept small offers.

Moral Image Constraints

People are powerfully constrained by moral image, yet they also seek ways to avoid this constraint so that they can be more self-interested (Dana et al. 2011). For example, a burgeoning literature on social preferences demonstrates that people do not like to appear unfair, either to themselves or to others. Though not a literal constraint, subjects in economic experiments apparently feel constrained by moral image. Even when they have total bargaining power, subjects will often share an experimental surplus to appear fair, but become more selfish when image concerns can be avoided through ignorance (Thunström et al. 2016; Van der Weele 2012). Dana et al. (2007) demonstrate this pattern using a modification of a simple dictator game. When subjects were allowed to choose between \$6 for themselves and \$1 to an anonymous other subject, or \$5 for both, more than three quarters chose the “fair” \$5 for both options. In a subsequent manipulation, however, subjects could choose between \$6 and \$5 for themselves while remaining uncertain about the impact it had on the other subject. In this manipulation, the payoffs were either conflicting, as described above, or aligned such that \$6 for the dictator gave \$5 to the other subject and \$5 for the dictator gave \$1 to the other subject, as decided by a coin flip prior to the experimental session. Close to half of dictators did not acquire the information of which game was being played, even though it required simply clicking a button on the computer screen. As a result, most chose \$6 for themselves, ultimately securing less than half the number of \$5–\$5 outcomes as when the outcomes were known. Apparently, people abide by fairness, but would happily rely on ignorance so as not to have to abide by fairness.

The result of this study, however, does not clarify whether the source of the moral constraint is looking unfairly toward one’s self (not revealing means you will never know you were unfair) or looking unfairly toward others (the other does not know whether you revealed, thus providing plausible deniability for being unfair). Further experiments by Dana et al. (2007) showed that players strategically took advantage of *plausible deniability*, where it existed, to be unfair without appearing so. That some players are willing to be unfair only when it does not appear so is interesting in this context because, again, the other subject is anonymous and cannot punish unfair behavior. These results suggest that even the imagined disapproval of the other party constrains behavior. Deliberate ignorance of the impacts of one’s behavior on others expands the

action space to options that people will not allow themselves to choose when they know that they will appear unfair.

Self-Serving Bias

Another psychological reason why a negotiator would benefit from deliberate ignorance is the failure to remain objective. Babcock et al. (1995) demonstrated that self-serving biases could lead to a breakdown in mutually valuable negotiations. Subjects read materials from a legal case either before or after they were assigned to the role (plaintiff or defendant) they would be negotiating. They also privately predicted how the judge on the case would rule and were paid for their accuracy. They then attempted to negotiate a settlement from a surplus provided by the experimenter. The longer they went without settling, the more the surplus shrank, before a neutral judge ultimately decided the case. Thus, they had an incentive to reach a settlement on their own.

When subjects were assigned their roles before reading the case materials, rather than after, they were more likely to reach a costly impasse, suggesting that they could not process information objectively once they knew which side they wanted it to favor. Buttressing this interpretation, the gap between the plaintiff's private prediction of the judge's award and the defendant's prediction was larger when the roles were learned before reading the case rather than afterward. Because there is no strategic advantage to inflating the private and incentivized prediction of the judge's award, it appears that subjects were unable to remain objective even if they wanted to be. Even though they were biased to favor their own side, the resulting bargaining failure was personally costly to the subjects, and thus they would have been better off had they been able to remain objective. In this situation, deliberate ignorance could have been beneficial (e.g., ignoring information about the roles while learning the facts).

Just as a teacher might use blind grading of exams or the philosopher implores us to get behind a "veil of ignorance" in evaluating distributional justice, the Babcock et al. (1995) studies demonstrate the benefits of ignorance in disciplining bias (see also MacCoun, this volume). Bias can cause failures to reach mutually beneficial agreements if the parties care about fairness. What is less clear is whether people appreciate these effects, and whether they would desire to use deliberate ignorance to negotiate more objectively.

Conclusion

Research into the phenomenon of deliberate ignorance is a somewhat new and intriguing field as applied to individual decision making, because standard analyses of decision making hold that ignorance should not work to the advantage of a decision maker. Negotiation adds a layer of complexity to the topic because it entails convincing others to agree to decisions, and little research

speaks directly to the empirical and psychological aspects of deliberate ignorance in negotiations. In this chapter, we have speculated that for a number of behavioral reasons, remaining deliberately ignorant can actually expand an agent's choice set in a negotiation.

