Yale ICF Working Paper No. 97-01
(April 1997)

"Common Knowledge As A Barrier To Negotiation"

Ian Ayres

And

Barry J. Nalebuff

COMMON KNOWLEDGE AS A BARRIER TO NEGOTIATION

by

IAN AYRES* and BARRY J. NALEBUFF**

May 4, 199710:30 am

File: D:\NALEBUFF\COMMON.PAP

*William K. Townsend Professor, Yale Law School. ayres@mail.law.yale.edu (e-mail)
http://elsinore.cis.yale.edu/lawweb/faculty/ayrespub.htm

**Milton Steinbach Professor, Yale School of Management. barry.nalebuff@yale.edu.
http://mayet.som.yale.edu/coopetition. Conversations with Adam Brandenburger, Jennifer Brown, Clark Freshman, Eric Green and Kenji Yoshino were very helpful. We want them to know we appreciate that help and know, too, that we know they know how helpful they were.

TABLE OF CONTENTS

COMMON KNOWLEDGE AS A BARRIER TO NEGOTIATION

INTRODUCTION

When we communicate one thing, we often unavoidably send other messages. To start with a simple example, imagine that Ian says to Barry, "My mother's name is Karen." From Ian's communication, Barry learns more than just the underlying bit of information (mom's name). The communication also lets Barry know that (a) Ian knows his mom's name and (b) Ian knows that Barry knows Ian's mom's name.

What is less well understood is that *when we teach, we learn*. When Ian tells Barry about his mom, Ian learns several things. For example, by telling Barry, Ian knows that Barry knows Karen is the name of Ian's mom. Direct communication of a fact can potentially create two infinite series of knowledge. If we symbolize the underlying bit of information (Karen is Ian's mom) as , then Ian's communication might create the following hierarchy of beliefs:

TABLE 1: POTENTIAL HIERARCHY OF BELIEFS

Ian's Knowledge	*Barry's Knowledge*
1a. Ian knows	1b. Barry knows
2a. Ian knows Barry knows	2b. Barry knows Ian knows
3a. Ian knows that Barry knows that Ian knows	3b. Barry knows that Ian knows that Barry knows
and so on . . .	and so on . . .

[1] Of course, for this to happen, it must be apparent to both the speaker and listener that the listener heard and understood what the speaker was saying. In many contexts it will be unclear whether the speaker received the message. (Indeed, we will argue that such ambiguity may on net facilitate settlement.) But in other contexts -- say for example, if Barry's response confirms that he understands ("Oh really, my mother's named Karen too) -- it will be clear as a practical matter that common knowledge is created.

If Ian's communication succeeds in creating both of these infinite series, economists would say that the underlying fact is "common knowledge." Table 1 makes clear how directly communicating a simple fact can produce other types of knowledge. Ian begins with what we will call "first-order" information (1a) and wants to convey this to Barry (1b). But in doing so Ian may be teaching Barry "higher-order" information as well (2b, 3b, 4b, etc.). Moreover, by teaching Barry, Ian may unavoidably acquire higher-order information himself (2a, 3a, 4a, etc.).

Now for something like "my mom's name is Karen," it doesn't matter that "I know you know" and "you know I know you know" etc., but in other cases it does. Indeed, the purpose of this article is to show that during negotiations parties will often want to communicate first-order

---

2 The term is usually defined more compactly: "Something is common knowledge if it is known to each player, and in addition, each player knows that the player has this knowledge; knows that the other person knows the player knows it; and so forth." Douglas G. Baird, Robert H. Gertner, and Randal C. Picker, Game Theory and the Law 304 (1994).

For pioneering work on common knowledge, see Adam Brandenburger, Knowledge and Equilibrium in Games, 6 J. Econ. Persp. 83 (1992). See also John Geanakoplos, Common Knowledge, 6 J. Econ. Persp. 53 (1992); Paul Milgrom & Nancy Stokey, Information, Trade and Common Knowledge, 26 J. Econ. Theory 17 (1982); Aumann, Agreeing to Disagree, 4 Annals of Statistics 1236 (1976);

In both common parlance and the common law, the term "common knowledge" merely means that an underlying piece of information is generally known. See, e.g., Garrison v. Heublein, Inc., 673 F.2d 189, 192 (7th Cir. 1982) (vodka was not unreasonably dangerous because the dangers involved in the use of alcohol are "common knowledge"); Caldwell v. Knight, 92 Ga. App. 747, 89 S.E.2d 900 (1955) ("common knowledge" doctrine does not require expert evidence when the reasonableness of the defendant's conduct involves matters within the collective common knowledge of the jury). In terms of Table 1, this definition only requires "first-order" knowledge (1a and 1b) and not the higher-order knowledge required under the game-theoretic definition.

3 Notice that these two infinite series are linked together in a kind of "double-helix" structure. In its strongest sense, (3a) implies (2b) which in turn implies (1a); and, (3b) implies (2a) which implies (1b). Higher-order knowledge implies certain forms of lower-order knowledge (but not vice versa). However, if Barry and Ian merely "believe" certain things -- and hence may be mistaken -- then a higher-order belief (say, 3a) would not necessarily imply a lower-order belief (2b).

information, but will not want to communicate to others (or teach themselves) higher-order information. We will show that common knowledge can itself be a barrier to efficient negotiation. Even when the communication of first-order knowledge facilitates agreement, the acquisition of higher-order knowledge at times can cause negotiations to unravel.

This is where mediation comes in. Mediators can break the link between communicating first-order information and communicating higher-order pieces of information. If a mediator caucuses first with Ian and then tells Barry , Barry will have acquired the first-order knowledge but he may not know that Ian knows (2b) and Ian may not know that Barry knows (2a) and so forth. The use of caucus mediation can usefully prevent from becoming common knowledge between the Ian and Barry. Caucus mediation can communicate (first-order knowledge) without creating common (higher-order) knowledge among the parties.

In this article, we will work through a series of examples showing how mediators can facilitate agreement by *preventing* the creation of common knowledge. In each example, the communication of first-order information () lubricates trade, but the communication of higher-order information would impede trade. Direct communication at times is unhelpful because the listener also learns higher-order information (Barry knows that Ian knows) and at other times it is

---

4 In caucus mediation, it's true that:

(a) when Ian tell the mediator, may become common knowledge between Ian and the mediator;and

(b) when the mediator tells Barry, may become common knowledge between the mediator and Barry.

But making common knowledge between the mediator and each of the individual parties may not be as destructive to negotiation (as making common knowledge between the individual parties themselves) -- because the mediator will usually not have as long an ongoing relationship with the individual parties as they do with each other.

5 Our paper extends the thesis of Jennifer Brown and Ian Ayres that the uniquely mediative role for mediators is to control the flow of information between the parties. Jennifer G. Brown & Ian Ayres, *Economic Rationales For Mediation*, 80 VIRGINIA LAW REVIEW 323 (1994).

unhelpful because the speaker herself learns higher-order information (Ian knows that Barry knows) that s/he would prefer not to know.

Our examples are organized around the "value claiming"/"value creation" dichotomy introduced by David Lax and James Sebenius. Communicating first-order information can be valuable in either type of negotiation. In distributive negotiations, it is often useful for one side to communicate to the other side "You are claiming too much, because . . . ." In integrative negotiations, it is often useful for one side to communicate "You could create more value if you would . . . ."

The problem is often the ellipses. First-order information often concerns *threats* and *insults*. In a distributive negotiation it can be both individually and jointly useful for one party to explain why the other side's BATNA is lower than she thinks. The problem is that the way I make your BATNA low is typically not very nice: "In the absence of an agreement, you will have a low payoff, because I will do something that will hurt you (break off relations, bring suit, etc.)." But higher-order information about threats can destroy value. If it becomes common knowledge that Ian has threatened Barry, then Barry may be more reluctant to trust Ian. And Ian (the threatener) may be less likely to trust Barry: after all, how can Ian trust Barry if Ian knows that Barry knows that Ian is willing to exploit weakness?

Common knowledge can also inhibit integrative bargaining. To tell the other side how he or she could increase the joint gains from trade, often entails revealing existing weaknesses about

---

6 David A. Lax & James K. Sebenius, The Manager as Negotiator (1986). In distributive negotiations -- where the parties are merely trying to determine whether there are gains from trade and how to divide them -- value claiming predominates as parties rivalrously claim different shares of the gains from trade. In contrast, integrative negotiations present opportunities for the parties to create additional value by varying the terms of trade -- so the parties must engage in a mixture of "value claiming" and "value creating" activity.

7 BATNA refers to "Best Alternative To Negotiated Agreement". See infra note 23. A negotiator wants his adversary to have a weak BATNA. And more than that, he want his adversary to believe that she has a weak BATNA. Because, ultimately, his adversary's perceptions will determine her willingness to agree.

the other side. To create additional value, you sometimes need to implicitly insult the other side. As with threats, insults that become common knowledge can destroy value. An employer might want to tell an employee that they could have a (jointly) more valuable relationship is the employee stopped stealing. But directly communicating this information can make both an employer and the employee less willing to continue the relationship. How can the employee continue to work with someone who knows that she has been stealing. And how can the employer continue to hire someone who know the employer countenance stealing.

Caucus mediation is not the only way that savvy negotiators try to solve these "common knowledge" problems. A benefit of many existing, indirect modes of speech is that they might communicate first-order information without creating the problems of common knowledge. There is a tradeoff here: by communicating indirectly we reduce the chance that the listener will understand the first-order information, but in return we reduce the discomfort of creating common knowledge. Caucus mediation can serve the same function: the mediator can speak directly about the first-order information but preserve uncertainty about whether the other side knows that this is being communicated.

It is often difficult to prevent second-order information from being communicated along with first-order information. If I tell a mediator a particular threat, I have a pretty good idea that the other side will be told. But it is less certain whether the other side will know that I know and (even less certain I will know that the other side knows that I know). Since higher-order knowledge is contingent on a chain of beliefs, uncertainty at lower levels may become exponentially magnified at

---

8 Some mediators claim that they will only pass on information with a speaker's consent, but in practice mediators indirectly disclose information garnered from one side during caucusing. Jonathan M. Myman, The Model Mediator Confidentiality Rule: A Commentary, 12 Seaton Hall Legis. J. 17 1988). See also Brown & Ayres, supra note 5, at 391-92.

higher levels. Even when mediators can't prevent the communication of second-order information, they can be successful at creating uncertainty about the more attenuated forms of higher-order information. And we will show that controlling third- and fourth-order information can at times beneficially affect negotiations.

The remainder of the paper is organized in three parts. The first shows how negotiations can be crucially affected by beliefs about higher-order information. The second part discusses a series of examples concerning "value claiming" talk. And the last part discusses examples concerning "value creating" talk.

I. THE IMPACT OF HIGHER-ORDER INFORMATION ON STRATEGIC BEHAVIOR

A problem with most game-theoretic models is that they make one of two extreme assumptions about information. They either assume that particular pieces of

---

9 If there is an 80% chance that first-order knowledge is communicated, and an 80% chance that the sender knows about the communication, and an 80% chance that the receiver knows that the sender knows -- then there is only a 51% chance of third-order knowledge (.8 x .8 x .8).

10 Professors Susskind and Cruikshank, for example, have described how mediators can preserve uncertainty in caucus mediation:

> A skilled intermediary can, in private meetings with the other participants explore whether they would be willing to give up Y and Z in exchange for X. This might be phrased, "What if I could get them to give up X? Would you trade Y and Z?" Of course, the neutral already knows that such a trade is possible. He or she must phrase the question, though, in a what-if format to protect the confidentiality of the information secured earlier.

Lawrence Susskind & Jeffrey Cruikshank, Breaking the Impasse: Consensual Approaches to Resolving Public Disputes 147 (1987).

Brown and Ayres have suggested that mediators should be more forthright about how they indirectly disclose caucus information, Brown & Ayres, supra note 5, at 327 & 392, but our analysis suggests that in a caucus mediation, the disputants may want the mediator to simultaneously (a) reveal caucus information *and* (b) say that she is not revealing. The fiction of mediator confidentiality may be a usefully device to preserve uncertainty about higher-order information.

information () are private information or they assume that particular pieces of information are common knowledge. But as our previous discussion makes clear there are literally an infinite number of intermediate assumptions that might be made. This section will show that assuming common knowledge is not just an innocuous proxy for situations where both parties know a particular fact. Even when both sides know a particular piece of first-order information, the parties' beliefs about higher-order information can crucially affect how they behave and the ensuing equilibrium. For now we will not concern ourselves with how players acquire information, we seek only to establish that higher-order information matters.

Example No. 1: How Big Is the Strike Fund?

To illustrate the importance of higher-order beliefs, we'll begin with a story about a labor negotiation occurring before and possibly during a strike. For the sake of argument, imagine that both labor and management are intransigent -- having made incompatible take-it-or-leave-it offers concerning job security. The only way to resolve the dispute will be for one of the two side to cave and accept the other's position. The question is who will concede.

Let us further assume that the cost of the strike is $10,000 a day for labor and $30,000 a day for management. However, management is better capitalized:

---

11 In subsequent section, we will then explore how mediation and indirect speech can allow a party to affect the other side's (and her own) higher-order information set.

12 A dichotomous outcome might be caused by some indivisible contract term -- such as job security.

Management's strike fund is $3,000,000 and Labor's strike fund is $700,000. Accordingly, Management can afford to hold out for 100 days while labor's strike fund will be exhausted at the end of 70 days. (We assume that when the funding runs out, a party has to concede and the strike ends.)

Although Management can afford to wait 100 days, it is not in management's economic interest to do so. Management figures that giving into the union's demands will cost it (in present discounted value) a total of $1 million. Hence if the strike were to last more than 33 days, it would be better for management to simply give in today. For labor, the strike is much more valuable. The present value of the job security it has demanded are worth $600,000. Consequently, Labor would be willing to strike for 60 days if that is what it took to win. Should it anticipate that the strike would take more than 60 days, then it's not worth fighting -- even if Labor thought that it would ultimately win.

Given this scenario, who do we expect to win the conflict? At first glance, it looks like labor is willing to fight longer than management and so it should be the victor. Anticipating this, management might as well give in right away. But as you

---

13 At this writing American Airline and its pilots are negotiating in the shadow of similar strike consequences. Tom Stieghorst, Deregulation, Unions: Airline Industry's Rocky Ride, Sun-Sentinel (Fort Lauderdale, FL) at 14A (February 9, 1997) ("American has $ 1 billion in profits socked away").

14 Even if management knew that it could ultimately win in 34 days, the strike would cost management $1,020,000 (34 x $30,000) and the wage concession would only save management $1,000,000.

15 Job security might cost Management more than it benefits Labor if it makes production less efficient.

might suspect, the parties' optimal behavior is more subtle and depends critically on beliefs about beliefs. To keep things as simple as possible, we assume that Management's and Labor's daily strike costs ($30,000 and $10,000, respectively), Management's strike fund ($3,000,000) and Management's and Labor's cost of conceding ($1,000,000 and $600,000, respectively) are all common knowledge. We focus solely on the players' beliefs about the size of Labor's strike fund ($700,000) to show that higher-order beliefs about the fund can crucially affect the length of the strike.

To understand the play of the game we need to consider what would happen at all possible points in the game. Let's begin at the end and analyze the player's optimal behavior at the end of the game and then work backward. The game cannot last longer than 70 days because Labor would need to concede after this point. In the event that the players reach the 70th day, what should management do? If management knows Labor's budget constraint, then it shouldn't concede. By holding out a single additional day (costing $30,000), management can assure itself that Labor will concede (producing a $1,000,000 benefit). Even though holding out for 70 days produces a net loss for management, from the perspective of the 70th day the costs of past strike days are "sunk" and therefore should not affect management's decision on whether to continue.

Indeed, this logic suggests that if Management knows the size of Labor's strike fund,

---

16 For simplicity, we also assume no discounting of future profits.

17 (-$1,200,000 = $1,000,000 - 70 x $30,000)

Management's dominant strategy is not to concede after day 37. Upon day 37, Management will have to wait no more than 33 days in order to be assured of winning the strike and the incremental cost of holding out from this point on is worthwhile in order to win the strike. Of course, we haven't yet determined whether management will want to holdout until day 37, but if it does, management will continue to hold out. For sake of notation, call the fact that Labor has only a 70 day strike fund . Our first conclusion therefore is that: Management should stay in the game from period 37 onwards if it knows .

If Labor knows that Management know , then at any point after day 37, it should quit immediately. That is, the consequence of Management knowing is that Management won't be giving in, and once Labor understands this (by knowing that Management knows), Labor might as well save the otherwise wasted strike costs and give in today.

What if management knows that labor knows that management knows ? Then management shouldn't quit, if it ever reaches day 4 of the strike. The reason is that if Management can hold on until day 37--which is only 33 days away--Management knows it will then win right away. Management knows that on day 37 Labor will see the handwriting on the wall and then concede immediately. Therefore Management's holding out after day 4 is worthwhile, since the wait will at most be 33 days until victory.

To make the next step in the argument requires that labor know that management knows that labor knows that management knows . In that even, labor should concede right away at day 4. They can predict at day 4 that the game will be over on day 37 with them conceding, so they might as well concede now.

Of course, if management knows that labor knows that management knows that labor knows that management knows then they should surely not concede on even the 1st day. All they have to do is wait until day 4 and victory will be theirs. Four days is a small price to pay.

---

18 See Lucian Bebchuck, A New Theory concerning the Credibility and Success of Threats to Sue, 25 J. Legal Studies 1 (1996).

The final step of the argument arises if we can say that labor knows that management knows that labor knows that management knows that labor knows that management knows L. It's a mouthful. If true, labor might as well give in right away. The alternative is to strike for 4 days and then concede.

It seems that by using backward induction we have "solved" the game. The predicted result is that labor will concede immediately. But that conclusion requires some very strong assumptions about higher-order knowledge. In particular, we can only confidently predict that Labor will immediately concede if we are willing to assume the existence of sixth-order information. It might very well be the case that both management and labor know L. Labor might even know that management knows L. But, management isn't sure of this fact. Hence it can't count on labor giving in on day 37. Since it's not worth fight 37 days to win, management could in full rationality

---

19 We can restate this example in terms of Table 1, by relabelling Ian = Labor; Barry = Management; and = = first-order information that Labor's strike fund is $700,000. Our examples shows that the following types of information are sufficient to produce the following equilibrium behavior:

Knowledge

Equilibrium Behavior

(1b) Management knows
Management will hold out from day 37 on
(2a) Labor knows Management knows
Labor concedes on day 37
(3b) Management knows Labor knows Management knows
Management will hold out from day 4 on
(4a) Labor knows Management knows Labor knows Management knows
Labor concedes on day 4
(5b) Management knows Labor knows Management knows Labor knows Management knows
Management will hold out from day 1 on
(6a) Labor knows Management knows Labor knows Management knows Labor knows Management knows
Labor concedes on day 1

This example exploits the "double helix" structure of iterated knowledge, discussed supra note 3, in that we assume (6a) implies (5b), which implies (4a), etc.

decide to concede right away. In fact, it could even decide to concede on any day between 1 and 36.

Example No. 2: Rationality and the Prisoners' Dilemma

Critics of economics often attack the assumptions that individuals are rational. But in many contexts the unexamined assumption that rationality is common knowledge can have profound effects on how players behave. Assuming that the rationality of each is the players is common knowledge is much stronger than merely assuming each player is rational. We might not want to abandon the notion of rationality, but we may want to relax the notion of common knowledge of rationality. Doing so allows for markedly different equilibrium behaviors.

To explore how higher-order knowledge about rationality can affect a game, let's look at a prisoners dilemma game that is repeated a finite number of times. Our point here is not to show how higher-order knowledge about rationality affects negotiations, but to show more generally how higher-order knowledge can affect equilibrium behavior.

First let's analyze the case where there is only one iteration. If either person is rational, she should defect. Period. End of discussion. And the only requirement is rationality. It doesn't matter whether you think the other person is rational or not. Defecting is the dominant strategy if you are rational.

Now, let's look at a two-iteration version. Defection in the first period is now only a dominant strategy if I am rational *and* if I believe that you are rational. If I think that you might not

---

20 In this stylized example, Management would have every reason to communicate that it knows the size of Labor's strike fund in order to make L common knowledge. But in later examples both parties will be better off if higher-order information is not produced.

21 Lax and Sebenius have argued explicitly that the negotiators' dilemma of whether to create or claim value is analogous to the prisoners' dilemma of whether to cooperate or defect. David Lax and James Sebenius, The Manager As Negotiator 29-41 158-166 91-105 (1986).

be rational, it might make sense for me to cooperate in the first period in order to induce you to cooperate in the second-period.

Finally, let's look at a three iteration version. What is a sufficient assumption for me to defect in the first period? Well (from the last paragraph) I know that my opponent will defect next period if (at the beginning of next period) she believes I'm rational. Because if she thinks that I am rational (even after seeing me cooperate in the first period), she will believe in the third period that I will defect and therefore she will defect in the second period. Since she will defect in the second period, I might as well defect in the first. It is my belief about her belief of my rationality that is crucial. And so, if I believe that she is rational and if I believe that she will think that I am rational even if she sees me cooperate in the first period, then I will want to defect in the first period.

The take home lesson is that in a finitely repeated prisoner's dilemma game, the unique equilibrium is defect-defect in all rounds if the players have sufficiently high-order beliefs about rationality. The order of beliefs about beliefs about rationality has to be commensurate with the number of repeat plays. We don't have to assume that the player's are irrational to produce cooperation. It may be sufficient -- and more reasonable -- to simply assume that rationality is not common knowledge (and does not extend to order levels higher than the number of rounds). This "resolution" of the finitely repeated prisoner's dilemma shows how assumptions concerning seemingly attenuated, higher-order knowledge can dramatically affect the play of the game.

II. "VALUE CLAIMING" TALK

We turn our attention now to how communication during negotiations can affect both first- and higher-order information. In a negotiation, it is useful for each party to know the maximum and minimum price at which a deal might be struck. ADR theorists tend to refer to these extremes as each party's BATNA (Best Alternative To Negotiated Agreement); economists tend to use the

---

22 Cooperation can also be induced in infinitely repeated games if the future is sufficiently important (i.e., the discount rate is sufficiently small). Robert Axelrod, The Evolution of Cooperation (1984).

term "reservation price." But whatever their name, misinformation about these extremes can impede negotiation. For example, a seller who is misinformed about the buyer's reservation price may demand a price that exceeds the maximum amount a buyer is willing to pay. Poor information about BATNAs can impede trade because a party may mistakenly demand more gains from trade than actually exist. To prevent this inefficiency, it will be both individually and socially useful for negotiators to credibly communicate information about both their own and the other side's reservation price. These communications take the generic form:

(1) Your BATNA is weaker than you think, and

(2) My BATNA is stronger than you think.

But these coolly rational communications are often backed up by hot threats that can impede settlement if the threats become common knowledge.

A. *Don't Claim So Much Because Your BATNA Isn't As Strong As You Think*

Each of the four examples in this section involve threats. Remember our primary goal is to show that first-order communication of threats can facilitate settlement, but higher-order communication (i.e. the threat becoming "common knowledge") can impede settlement. If we

---

23 We prefer the term "value added." See Adam M. Brandenburger & Barry J. Nalebuff, Co-opetition (1996) ADR theorists might not equate BATNA with a reservation price, because they are reluctant to monetize the "best alternative." But when a BATNA is easily monetizable the two concepts are identical. For example, if Yale is negotiating to buy electricity from a local supplier and Yale knows that it can buy as much as it wants out of state at 3 cents/kw, then we could say that this competitive offer represents either Yale's BATNA or its reservation price.

24 These are not the only examples of "value claiming" talk. For example, if you are reluctant to negotiate because you mistakenly believe that there are no gains from trade, I might want to tell you, "My BATNA is lower than you think." In example 4, we will also discuss a lying example, where one party wants to communicate "I know (what you already know) your BATNA isn't as strong you are claiming it is." See infra at 18. And as we discuss at the end of this section, "Value claiming" talk can relate to any primitive of a distributive bargaining model. See infra at 15.

succeed in establishing this, it is straightforward to show that mediators can help achieve these competing objectives -- passing on first-order information, while filtering the higher-order knowledge.

Example No. 3: A Threat of Divorce

Imagine a husband and wife who are having trouble in their relationship. The husband might say "Honey, I want this relationship to work but if you can't make these changes then I want a divorce." Some spouses would respond to this type of "threat" by saying: "You got it." The husband might reply, "I didn't want the divorce, I wanted the changes." But by this point, it may be too late.

The problem is that the husband's threat communicated not only the consequences of maintaining the status quo but higher-order information as well. The wife now knows that the husband is contemplating a divorce and the husband knows that the wife knows this. The wife might not be willing to see a marriage counselor or try to work something out if she thinks her husband has one foot out the door. The wife can't pretend not to know and the husband can't pretend that the wife doesn't know.

The higher-order information can make each side more reluctant to stay in the marriage. The wife is reluctant to stay in a marriage when the husband knows she is knuckling under to a threat. And even if she were willing, the husband may think that given her current knowledge, she wouldn't be really committed to change and thus he isn't willing to make the commitment. The

---

25 We know of an engagement threat that is all to close to this example. After wedding invitations had been sent, the bride-to-be raised the possibility of postponing the ceremony. The groom-to-be responded in no uncertain terms: "if you postpone the wedding, I will never marry you." This statement caused the engagement to quickly unravel. The bride-to-be couldn't marry someone who would make such an explicit threat.

26 In terms of table 1, K is the fact that the husband is contemplating divorce. The husband's threat not only communicates K (1b) but higher-order information.

higher-order information that the husband learns when he threatens can make him less willing to stay in the marriage. Common knowledge of the threat destroys the possibility of mutual trust and commitment.

Savvy mediators (marriage counselors) can help the couple see the consequences of their actions -- without creating common knowledge. In essence, it is useful for the wife to understand that her BATNA isn't as strong as she thinks it is. She may mistakenly think that she can keep the status quo and preserve her marriage. The marriage counselor (after talking privately with the husband) might ask the wife: "What do you think will happen if you stay on the current course?" or (if that doesn't do the trick) "Do you think your marriage might end if nothing changes?". Of course, the wife might infer that even the first, more indirect question has its roots in a threat that the husband communicated to the mediator. But if done well, the mediator can convey the message but retain sufficient uncertainty on both the husband and the wife's side about higher-order information.

Example No. 4: An Opportunistic Threat

The disruptive effects of direct threats apply equally well to business settings. Imagine a negotiation between two corporations concerning a contract renewal to supply certain goods on January 30. The supplier might try to negotiate a higher price by saying on January 15: "Look, we are the only company in a position to produce the goods right now for you, so if you want your supply on January 30, here's the price and you better sign this contract or your factory is going to get shut down." The buyer might sign the contract because she has no alternative. As in the divorce example, the supplier in essence is trying to tell the other side that it has a weaker BATNA than it realizes: If it doesn't agree to pay the higher price, it will be left with no goods at all. But the buyer's BATNA is weaker because of the seller's opportunism. The seller is threatening to do something it may have a legal right to do but which violates any norm of a "partnership" orientation.

Directly communicating this threat might cause the relationship to end because the threat

becomes common knowledge between the two companies. The threat might force the buyer ultimately to break off the relationship because the buyer realizes that the seller has already actively considered opportunism. Once the buyer knows that the seller doesn't have a "partnership" orientation, the buyer might well choose not to continue a long-term cooperative arrangement.

Moreover, what the seller learns by making the threat may now make the seller herself more reluctant to continue trading with the buyer. Now that the seller knows that the buyer knows that the seller can't be trusted, the seller has less reason to trust the buyer. An unintended consequence of the seller's threat is that the seller now realizes that he can no longer trust the buyer. In retrospect, the seller may wish he had not made the threat, but there is an irreversibility to his threat: how can the buyer ever not know what she now knows. The knowledge ir irreversible.

Once again, a mediator could communicate to the buyer the consequence of no agreement without the threat becoming common knowledge. The mediator doesn't need to directly suggest the consequence: "The supplier might break off relations if you don't pay the higher price." It will often be sufficient to ask the buyer: "What do you think will happen if you don't reach this agreement? What do you think the supplier might do? What would you do if you were the seller?" This type of mediation can be effective precisely because it obscures the higher-order information. Even if the buyer comes to realize that the seller may stop shipment in the absence of renewal, the seller doesn't know that the buyer has come to realize this. The buyer may say, "It's one thing for me to pay a higher price to secure a much needed shipment, but it's another thing for me to

---

27 Going back to Table 1, the underlying bit of information () that the seller (Ian) is trying to communicate to the buyer (Barry) is the possibility that the seller might not ship the current order if the buyer doesn't agree to a price increase. The seller's direct threat communicates not only this bit of information to the buyer (1b), but the buyer also learns that the seller knows (2a).

28 In terms of Table 1, the seller's threat teaches the seller: the seller knows that the buyer knows (2b).

continue dealing with someone who knows that I agreed to pay the higher price because of an opportunistic threat." The mediation by preventing the latter can help to preserve relations.

Sometimes the role of the mediator can be played by the negotiators, themselves. One of the things that good negotiators often do is turn themselves into mediators. They get the other side to imagine what will happen if we don't reach an agreement: "Think about what would you do if you were in my shoes?" This process of getting the other side to imagine can be even more effective if the person imagining doesn't reveal what he has thought. To be effective, the seller doesn't need to know what it is that the buyer is thinking and the buyer doesn't have to know that the seller knows and so on.

Of course, this example also highlights a potential dark side of mediators' desire to preserve relationships -- because the example can also be seen as a method to facilitate seller opportunism. Potential victims might want to commit to direct communication. If the parties understand that opportunism would become common knowledge, they might act more civilly. While all of our examples are constructed to show how common knowledge can impede settlement, there are many contexts where *creating* common knowledge is a goal of at least one of the parties. The powerful urge to have your "day in court" is in part a desire to have the other side's guilt (or your own innocence) become not only common knowledge between the litigants but among the public generally. Or even in private negotiations, it may be important for me to obtain an explicit acknowledgment of your prior wrong doing before parties can proceed to a settlement. Mediators in trying to preserve relationships at times overlook the legitimate desires of individual negotiators to directly confront their adversaries and thereby make certain information

---

29 It is not enough that I know you acted badly and you know you acted badly. I may want explicit acknowledgment that "you know I know" and "I know you know I know."

common knowledge.

Example No. 5: A Gratuitous Threat

In this example, we turn to an experimental anecdote showing how a direct threat can cause a negotiation to unravel. The Harvard Negotiation Project filmed a simulation of an "unscripted and unrehearsed" negotiation, entitled "The Hacker/Starr Negotiation." In this simulation, Hacker and Starr were partners in a computer software joint venture. Hacker's job was to come up with ideas and write the computer code; Starr was a dentist who provided capital. The dispute concerned the rights to a screen saver program written by Hacker. Hacker claimed that he owned the program because Starr had vetoed the idea (repeatedly) and because Starr had written the program on his own time. Starr claims that the joint venture owned the program because it was written with a company computer and because the partnership agreement was broadly worded to give the partnership rights to any software written by Hacker.

The first portion of the mediation was a tremendous success. The parties overcame the acrimony and rancor from a previous meeting and agreed in principle to continue their partnership with a warm sense of good-feeling. They then turned to discuss how to divide the profits from the screen saver program. Hacker (who had been carefully trained/coached by his negotiation adviser, Roger Fisher) tried to turn the discussion toward brainstorming about possible options to resolve the remaining dispute. Hacker began by suggesting the possibility of either arbitration or a 5% Starr -- 95% Hacker split of the profits, and the following conversation ensued:

---

30 Trina Grillo, The Mediation Alternative: Process Dangers for Women, 100 Yale L.J. 1545 (1991). Hence in victim/offender mediation when mediators strongly counsel criminal victims to "own their own feelings" (i.e., not directly blame the criminals), they may be destroying value. Jennifer Gerarda Brown, The Use of Mediation to Resolve Criminal Cases: a Procedural Critique, 43 Emory L.J. 1247 (1994).

31 The Hacker/Starr Negotiation, Harvard Negotiation Project (1985). The two negotiators had previously met neither their clients nor each other. Hacker and Starr were actors who were told the basic simulated facts giving rise to the dispute and were only instructed to do "the best they could for themselves."

Hacker: Arbitration is a good possibility. . . . And also, I've already offered 5-95 and that's another option. But maybe there're more [options]. Name more if you think of one.
Starr: Well, one of the options is litigation that I could sue you.
Hacker: Sue me?
Starr: Well, you wanted options.
Hacker: So, I will put that down. "Sue Allen." For what?
Starr: [Reading from paper] First of all, you breached your contract. You disregarded the Board of Directors. Managed poorly. You did not devote full time to the company. . . . You caused a decline in the revenue of the company.

Starr's insistence on explicitly stating his willingness to sue and explicitly listing all of Hacker's bad acts almost ruined the deal. In the film, Hacker became visibly despondent -- slumping back in his chair and lowering his head.

Starr's threat was gratuitous. It was already obvious that a suit was an option, and Hacker had been asking for only non-litigation options. The parties had already agreed to the much more important aspects of the deal (a new management structure, a new capital infusion). Starr once again is trying to say: "Your BATNA is not as strong as you think." But the communication was gratuitous because Hacker had already indicated that if he were in Starr's shoes he would feel slighted (and want to pursue litigation). It's as if the dentist couldn't help but say: "Good thing that you are finally being reasonable, because here are all the ways I was going to let you have it if you weren't." Starr's statements almost kill the deal because Hacker suddenly feels very uncomfortable continuing a long-term partnership now that Starr has coolly described how close he had been planning to screwing Hacker six ways till Sunday. Hacker knows the Starr is not currently planning to sue him (because they have in principle decided to rededicate themselves to the joint venture), but the very knowledge that Starr is the type of person who would have done this tells Hacker that maybe Starr is not the type of person with which he wants to have a long-term relationship. And more importantly for the purposes of this paper, Hacker may have more difficulty maintaining his concessions (e.g., the offer of arbitration) once it becomes common knowledge that this is a byproduct of an explicit litigation threat.

Example No. 6: A Threatening Lie -- An Example of Second-Order

Communication

We end this section with a discussion of how a negotiator might respond when it catches the other side lying in negotiation. Imagine that you are a supplier for a large corporate buyer. The purchasing agent for this corporation comes to you and says: "If you don't lower your price by 15%, I will buy from Acme who has offered me a 15% lower price." By some means, you find out that the purchasing agent is lying. Now what should you do?

Your initial impulse might be to directly confront the purchasing agent and say "you're lying." The purchasing agent has (falsely) claimed that she has a strong BATNA. But unlike our earlier examples, the response is not "Your BATNA isn't as strong as you think" -- because the buyer already knows that her BATNA is weaker than she has claimed. Instead the impulse is to say in effect, "I know what you already know (that your BATNA is weaker than you've claimed)." Your impulse then is not to communicate first-order information, but second-order information.

The problem now is that conveying this second-order information -- I know that you know you are lying -- may embarrass the purchasing agent in ways that make it impossible for her to continue to buy from you. The purchasing agent may be willing to lie but be ashamed about having been caught lying. Even if she isn't ashamed, for the sake of your reputation, you may now have to end the relationship: you can't let others (even liars) know that you're the type of person who does business with someone who is known to be a liar.

Instead of making the purchasing agent's lie common knowledge, it might be more

---

32 Acme might have independently revealed their true offer or perhaps you were able to read upside-down Acme's actual quoted price inadvertently left on the purchasing agent's desk.

33 In terms of Table 1, supra at 1, Ian is the purchasing agent; Barry is the seller and the first- order information is that the claimed Acme price discount is false. Because he is not a pathological liar, Ian already knows that the claimed discount is false (1a). Because Barry has independently learned of Acme's true quote, Barry also knows that Ian's claim is false (1b). But Barry's impulse is to let Ian know that Barry knows the claim is false (2b).

productive to respond with a "white lie" of your own: "Well, Acme may be able to cut its price, but I think that the quality of our product is better and I hope that in spite of the price difference, you'll still be willing to continue buying from us." Even if (it is common knowledge that) there are no real quality differences between your and Acme's goods, your white lie might prevent the purchasing agent's lie from becoming common knowledge. If the purchasing agent comes back and ultimately buys from you, she doesn't know that you know she's a liar.

Of course, if the goods in question are really homogeneous commodities, the purchasing agent may guess that you have caught her in a lie. But even then the white lie might facilitate trade because you won't be sure whether she received the message. It's embarrassing enough for her to have to buy from someone who knows she lied, but it's even more embarrassing to buy from someone who knows she knows that she has been caught. Thus even if the white lie communicates second-order information (and it might not), it can still facilitate trade by impeding the communication of even higher-order information. It might be useful to let the purchasing agent know that you are on to her, but keep yourself in the dark about whether she ever got the message.

Here again a mediator once could help implement a similar out come. It is much less embarrassing for a mediator to confront (perhaps indirectly) the purchasing agent about the lie because the purchasing agent will not have to deal with the mediator on a long term basis. The lie may become common knowledge between the mediator and the purchasing agent, but this may be much less troubling than making the lie common knowledge between the seller and the purchasing agent.

B. *Don't Claim So Much Because My BATNA Is Stronger Than You Think*

---

34 For an insightful discussion on the ethics of "white lies," see Gerald B. Wetlaufer, The Ethics of Lying in Negotiations, 75 Iowa L. Rev. 1219 (1990).

35 The purchasing agent may be particularly averse to public knowledge of her lie because she fears punishment from her superiors in the corporation.

It is often effective to educate the other side about the strength of your BATNA. Instead of saying "Don't claim so much because your BATNA is weaker than you think" you might want to say "Don't claim so much because my BATNA is stronger than you think." Indeed, this was the purpose behind the purchasing agent's claim (albeit false) in the last example: My BATNA is stronger than you think it is (because I can get a great price from Acme if I don't buy from you) so you better claim less (i.e. cut your price).

Yet the reason why a speaker's BATNA is unexpectedly strong can be just as off-putting as the reason why a listener's BATNA may be unexpectedly weak. Accordingly, there may be times when a negotiator will want to communicate first-order information explaining why her BATNA is unexpectedly strong, but will not want this explanation to become common knowledge. In such situations, as before, caucus mediation and indirect communications might be of use.

Example No. 7: Threatening Artificial Scarcity

One of the ways to increase your BATNA is to induce (or threaten to induce) artificial scarcity. Imagine that Nintendo is negotiating with Toys "R" Us (Toys) over how to split $100 profits on each of its video machines for a particular Christmas shopping season. Nintendo might begin the negotiation by revealing that there is an unexpectedly strong demand for its machines. Toys might respond that Nintendo still needs to use Toys as a retail outlet if Nintendo wants to sell 1 million units this season (Nintendo's total productive capacity). Accordingly, Toys could argue, "We both need each other and therefore we should split the $100 profits per machine."

In turn, Nintendo might want to threaten to reduce it productive capacity to increase the strength of its own BATNA. In essence, Nintendo would be saying, "While we need Toys "R" Us to sell 1 million units, if you don't agree to give us more of the profits, we will only produce 750,000 machines. We won't need you (or any other particular retailer) if we reduce our production. In that case, you'll be scrambling to get any units that you can and you'll be happy to

---

36 A more true-to-life narrative about Nintendo's marketing success is included in Adam M. Brandenburger & Barry J. Nalebuff, Co-opetition (1996).

give us most of the profits. So you (Toys "R" Us) should agree to let us have most of the profits, or we will create artificial scarcity which will mean we won't need to deal with you at all."

As should be obvious by now, the problem for Nintendo is that while it might want to communicate this first-order information to Toys, it might not want this threat to become common knowledge. Toys might be reluctant to continue a long term relationship with a seller who knows that it has caved in to a threat. Moreover, making the threat common-knowledge may reduce Nintendo's ability to trust Toys future marketing efforts -- after all, what kind of a person is willing to continue dealing with someone who opportunistically threatens to reduce quantity?

The larger point of this example is that negotiators will at times want to communicate information not about what the speaker will do to reduce the listener's payoffs in the absence of agreement, but what the speaker will do to increase her own payoffs in the absence of agreement. Fisher and Ury explicitly detail a variety of ways that negotiators can strengthen their own BATNA or weaken those of the other side. To be effective, these BATNA changes must not only be communicated, but be communicated in a way that does not queer the deal. Direct communication that makes these threatened BATNA changes common knowledge can ultimately be counter productive -- hence creating the possibility for caucus mediation and indirect communication to facilitate trade.

While we have focused almost exclusively on BATNA talk, negotiators might also want to communicate about any of the other underlying variables that affect distributive negotiations. For example, game-theoretic models of bargaining often turn on the parties' cost of bargaining or

---

37 Studies in experimental economics indicate that people will incur some loss in order to punish others for unfair behavior. Kahneman, Knetsch & Thaler, Fairness as a Constraint on Profit Seeking: Entitlements in the Market, 76 Am. Econ. Rev. 721 (1986).

38 Roger Fisher & William Ury, Getting To Yes: Negotiating Agreement Without Giving In 183-84 (1991).

degree of rationality. And Jennifer Gerarda Brown in this symposium has suggested "satiation" and "optimism" theories of hope that might affect distributive negotiations. One can easily imagine a negotiator saying: "Don't claim so much, because:

(a) I'm less rational than you think;

(b) My bargaining costs are lower than you think;

(c) Your bargaining costs are higher than you think;

(d) I'm more optimistic than you think, etc.

To make these claims credible, the speaker will often have to explain why the world is different than the listener thinks. Before providing direct explanations, the wise negotiator will consider the ramifications of communicating higher-order information.

III. "VALUE CREATION" TALK

The common knowledge problem can also be a barrier to integrative bargaining. In integrative settings, parties have more to bargain about than simply the price because the joint gains from trade can be increased if non-price terms are chosen wisely. Under such circumstances, it will often be mutually advantageous for one side to communicate information of the kind "We'd have larger gains from trade, if you would (agree to) do X." The problem is that in ongoing relationships emphasizing that I would like "X" often contains an implicit insult that you have been producing "not X" in the past. Communicating the first-order information (containing this implicit insult) facilitates integrative bargaining, but communicating higher-order information can embarrass and erode the trust of both sides to the negotiation.

Example No. 8: Misfeasance (herein Body Odor)

Imagine that you are negotiating an extended employment agreement with your employer.

---

39 See, e.g., Ian Ayres, *Further Evidence of Discrimination in New Car Negotiations and Estimates of Its Cause*, 94 MICHIGAN LAW REVIEW 109 (1995).

40 Jennifer Gerarda Brown, The Role of Hope in Negotiations (this symposium).

You enjoy your job, but you would be happier if your boss used (more effective?) deodorant. You know that she needs it, but she does not. What should you do?

You could tell her directly. But by now it should be obvious that to do so would convey not only the first-order information of interest but also higher-order information: "I know that you need deodorant" and "You know I know you need deodorant". The implicit fact that "your body odor has bothered me in the past" can be insulting and/or embarrassing -- especially if it becomes common knowledge. It's easier to correct the problem and preserve your relationship if the boss isn't sure that you have identified her odor as a problem or at least if you aren't sure whether the boss knows you know.

Or consider the communication problem that Ian encountered with his child's day care center. He wanted to communicate that he would get more value out of the center if one of the care givers would give more detailed comments about his child's progress. But Ian was worried that if his criticism became common knowledge (between this provider and him) that everyone would feel uncomfortable. Ian's solution was to communicate using negative pregnants and what mediators call "flip-sides." In his evaluation he stressed how much he appreciated the comments of the two other care givers (negative pregnant) and instead of saying that he disliked the generic comments of the third, he emphasized that he would appreciate individualized comments about his child (flip sides). While the care giver probably understood Ian's first-order message, Ian does not know for sure that she knows his criticism and moreover, she does not know whether Ian knows that she knows it.

These are both (rather trivial) examples of "misfeasance" -- because the listener doesn't realize that she is doing something that reduces joint value. And as in our previous examples, the

---

41 Mediation Training Materials, The Settlement Institute (1992). When a mediator uses "flip sides," she paraphrases a party's negative statement by transforming it into a positive statement about what the party desires in the future. For example, if A says "B is always late," the mediator might translate "A, you would like B to be more prompt."

speaker would like to tell someone something they don't know without letting them know that you know it (or without letting yourself know that they know you know it).

In these settings, a suggestion box or an anonymous note may have done the trick. Even if the recipient can probably figure out who sent the note, a sender will be uncertain about whether the recipient knows. Anonymous notes and suggestion box can be useful because they may prevent higher-order or common knowledge from being created.

In other situations, a speaker might want to use indirect and/or ambiguous statements -- taking a chance that the listener will not understand the first-order message, but in return preserving a kind of "plausible deniability" about higher-order information. In more serious contexts, a speaker might even want to enlist the help of a mediator to send the message.

Example No. 9: Malfeasance (herein Adultery and Theft) -- An Example of Second-Order Communication

In contrast to the unintentional misfeasance of the last section, it is sometimes useful for negotiators to communicate information about malfeasance. Imagine for example that a husband is having an affair. The husband already knows that he is being unfaithful. But unbeknownst to him, his wife discovers his infidelity. The wife might want to tell her husband that they could have a better relationship if he would be faithful. In essence, she would be communicating second-order information: "I know what you already know -- that you are having an affair."

---

42 We used a coin flip to determine the gender of the adulterer.

43 Unlike our earlier divorce example, the wife would not be threatening to divorce her husband if he did not stop ("Your BATNA is weak because . . ."). Rather she might just be engaging in integrative bargaining ("We would have a more mutually beneficial relationship, if you would . . .").

44 The malfeasance examples in this section are similar to the false threat example discussed above. The seller was tempted to tell the purchasing agent about his lie (another type of malfeasance): "I know what you already know -- that you have

But as in our previous examples, making the affair common knowledge might cause the relationship to unravel. As an initial matter, the adultering husband may be too embarrassed to continue. He might be too ashamed to live with someone who knows that he had (or is having) an affair. It's one thing to have acted badly but another to be caught. The husband might be able to control how much he "guilt trips" himself but he can't control how much his wife will now guilt trip him.

Moreover, common knowledge of the adultery might make the faithful wife less willing to continue the relationship. She has to live with the fact that her husband knows she knows and that she didn't have the nerve to "do anything" about it: What kind of person would just sit there and let her husband have an affair and continue living with him? And so, the need to show that she is not a wimp might force her to call it off.

Direct confrontation changes the situation from "the husband knows of the adultery" and "the wife knows of the adultery" to "the husband knows the wife knows" and "the wife knows the husband knows the wife knows." And these higher-order pieces of information can have destructive consequences.

Indirect communication via a mediator or ambiguous statements by the wife can at the very least let the husband know that his affair might become known. Or even if the mediator effectively

---

acted badly."

45 We can go even further. The unfaithful husband might be willing to continue in a marriage where his wife knows that he has been unfaithful. But maybe he can't continue in a marriage where his wife knows that he knows that she knows of his infidelity -- and did nothing about it.

46 This suggests that if you learn that a friend has been the victim of adultery, you should not disclose the adultery in a way that makes the malfeasance common knowledge between the spouses. At most you should disclose the adultery to the faithful spouse and let him or her decide whether making the information common knowledge is worthwhile. A similar argument might hold for other forms of malfeasance (such as employee theft) except for third-party interests in deterring bad acts generally.

lets the husband know that his wife probably knows of the affair, the mediation can help preserve the marriage by not letting the wife know that the husband knows she knows, and so on.

While we have used adultery as our example, the same basic story could be told with regard to other types malfeasance. If an employer finds out that her employee has been stealing (or intentionally shirking, etc.), she might want to signal that she knows of the wrongdoing without finding out whether the employee learns of the signal.

Or imagine a parent who has discovered that her teenager has tried marijuana, but doesn't want the teenager to know she knows. If the teenager's use becomes common knowledge, the parent would be forced to punish or condone the behavior. Punishing would be hypocritical and condoning would send the wrong message. Hence the best outcome is for the parent to be informed (to monitor if the experimentation becomes a problem), but not have the teenager know she knows.

It is easy to see that directly telling someone "you could do better in the future" often implicitly communicates an insult, "you've been bad in the past." But even when this linkage is unavoidable, we have tried to show that the direct statement might impede settlement because it might also communicate "I know you've been bad" and "You know I know."

CONCLUSION

Critics might argue that only a couple of economists would go to such great lengths to show that threats and insults can impede negotiation. But in fairness we have argued something more subtle. If these threats and insults merely impeded trade, one could simply counsel -- like the Vaudevillian doctor -- to stop making such statements. Instead we have tried to illustrate types of information that negotiators will want to communicate to facilitate trade -- but which can become destructive if the information becomes "common knowledge."

One of the classic tough guy responses to intimidation is the query: "Is that a threat or a promise?" This question implicitly arises during negotiations whenever one party describes what

---

47 "Doctor, my arm hurts whenever I lift it like this." "Don't lift it like that."

she will do in the absence of agreement. While it is often difficult to distinguish threats from merely describing the consequences of failing to reach agreement, our analysis suggests that higher-order communication might be a necessary element of what it means to threaten (or to insult). Telling the other side about (how you will affect) her BATNA is likely to be less threatening if she does not know you know (or if you do not know she knows you know).

Mediators often work to increase the objectivity of negotiators. The mediator may serve as an "agent of reality" or engage in "reality checking": making sure that the each party knows both her own and the other side's BATNA. Yet the goal of increasing objectivity could be applied to higher-order information as well. In a trade negotiation, a mediator or a negotiation consultant might want to increase the seller's objectivity about how the buyer perceives the seller. Besides making the seller understand what her BATNA is, it can often be useful for the seller to understand what the buyer thinks the seller's BATNA is. Even if the seller's BATNA is actually weak, she may have more bargaining power because the buyer thinks her BATNA is strong -- and she can bluff her way to a favorable agreement. Understanding the other's perceptions of you can be particularly difficult. The other side often doesn't know as much about you as you do. And its impossible not to know what you know. (This is why you can't play chess against yourself.)

Acquiring objectivity about others (and others' perceptions of you) is difficult, but does not necessarily require mediation. Mediators are not the only ones who can serve as agents of reality:

---

48 Communicating second-order information also seems to be a necessary component of successful "irony." If Ian ironically tells Barry, "That paper was *really good*." Part of Ian intent is not only for Barry to know that the paper is bad, but for them to share in the common knowledge joke that Ian intended just the opposite of what he said.

49 Game-theorists succeed at doing this in highly reductive models. A player who knows her reservation price imagines how the other side knowing only the distribution from which the reservation price is drawn would play.

50 You can't forget the lines of attack you are planning as white, when you try to imagine black's defense.

Negotiation consultants, lawyers, and friends can enhance someone's objectivity about herself, about her perception of the other side and about the other side's perception of her. In theory "reality checks" could be done by each negotiator without the help of a mediator or the other side.

The most difficult challenge is when a negotiator has private information that could make the other side more objective. This article has attempted to show how communication between the parties could be used to increase each side's objectivity. If you have a strong BATNA but believe the other side's perception of you is mistaken, you will be tempted to say "My BATNA is stronger than you think" (or at other times "Your BATNA is weaker than you think"). Savvy negotiators will at times turn themselves into mediators -- and try to increase the other side's objectivity in ways that facilitate settlement.

Caucus mediation can play a central role in helping each side educate the other. Controlling the flow of information between the parties is the only uniquely mediative task. But more objectivity is not always good. Somewhat surprisingly, we have shown that mediators in many instances should not strive to create better objectivity about higher-order knowledge. It is useful for a seller to be able to communicate to a buyer why the buyer's BATNA is weaker than she thinks, but it can be destructive for this threatening "why" to become common knowledge between the negotiators.

Game-theorists often model private information as *the* but-for cause of inefficient distributive bargaining. In these simple bargaining models, if each side's BATNA were common

---

51 Of course, we are less likely to hear: "You're mistaken, my BATNA is weaker than you think it is." But see supra note 24.

52 Brown & Ayres, supra note 5. Caucus mediation is the most "alternative" of the ADR techniques. It is the only dispute resolution method where the parties don't speak face to face.

53 Ian Ayres & Eric Talley, *Solomonic Bargaining: Dividing A Legal Entitlement To Facilitate Coasean Trade*, 104 YALE LAW JOURNAL 1027 (1995).

knowledge, the parties would instantaneously (and costlessly) reach agreement. But we have argued that while the lack of first-order information can impede trade, the presence of higher-order information should be thought of as a barrier to negotiation, a transaction cost that might be avoided by ambiguous or indirect communication or by caucus mediation.

Our argument provides an economic justification not only for mediation but for a host of communication techniques that are already in use -- such as "negative pregnants" and "flipsides". Our goal however is to do more than merely translate current practice into economic jargon. Instead, we believe that more explicit attention to "common knowledge problems" can lead toward a better practice. In closing we focus on two take home lessons.

First, in choosing between direct vs. indirect communication, negotiators (and mediators) should focus on the tradeoff between first- and higher-order information. The cost of ambiguous or indirect statements is that the listener may not understand your primary (first-order) message, but the benefit is that the first-order information is less likely to become common knowledge. The optimal degree of ambiguity should often turn on the relative size of these benefits and costs. Particularly in on-going relationships the costs of making threats and insults common knowledge may outweigh the benefits of clear first-order communication. Mediation and indirect communication can provide a useful, intermediate choice between the extremes of silence or common knowledge.

Second, before speaking, negotiators should consider not only how their speech will affect the perceptions of the listeners, but how speaking will change their own perceptions (even if the listener says nothing in response). Speaking changes your own information set -- because now you know that the listener knows something (and you know that the listener knows that you know

---

54 See Martin J. Osborne & Ariel Rubinstein, A Course In Game Theory 117-31 (1994).

55 See Mark Suchman, Translation Costs: A Comment on Sociology and Economics, 74 Oregon L. Rev. 257 (1995).

she knows, etc). We all know the witticism, "I could tell you X, but then I would have to kill you." Almost always it means, "I would have to kill you because of your additional knowledge." But we have shown that sometimes the phrase might mean: "I would have to kill you because of *my* additional knowledge." Once I gain higher-order information by making a statement ("I know you're committing adultery" or "I will breach my contract if you don't agree to pay more"), I might be forced to break off our relationship myself. Savvy negotiators not only think about how their statements will change the other side's "frame," but how their statements will necessarily change their own "frame." And here we are not talking about a cognitive bias story where we might react differently to a problem if it is framed as a gain or a loss, but instead we have shown how statements irreversibly change the speaker's information set.

In diverse settings, humans find it useful to communicate without creating common knowledge. In international diplomacy, giving your adversary an option to "save face" can often be interpreted as a way of obscuring whether she has capitulated to your threat. And some homosexuals and lesbians want their sexual orientation to be an "open secret" -- meaning, they want their friends and relatives to know their gay, but they don't want their orientation to become common knowledge. Negotiators would do well to be sensitive to this subtle, but recurring human

---

56 See Erving Goffman, On Face-Work, in Interaction Ritual: Essays on Fact to Face Behavior 6-7 (1967) (on the importance of saving face).

57 In some families if the "open secret" of a child's orientation became common knowledge, parents would feel the need -- because of religious principals -- to break off their relationship or vocally criticize the child. D. A. Miller, The Novel and the Police (1988).

Understanding the potential utility of preventing common knowledge also provides a justification for President Clinton's "don't ask, don't tell" regulation of gays in the military. Judge Eugene H. Nickerson in Abel v. United States, 880 F. Supp. 968 (E.D. NY, 1995) criticized this regulation by arguing in part:

> To "accommodate" the privacy of heterosexuals presumably means, for example, to keep their naked bodies safe in the showers from the stares of homosexuals. But if indeed there are homosexuals who wish to peek at naked bodies, they might do so quite as readily when their orientation is a secret as

dynamic.

> when it is open. The only difference will be that heterosexuals will not know which of their servicemates are homosexuals, and heterosexuals will have reason to have a generalized suspicion of everyone in the showers, hardly a circumstance likely to increase "cohesion."
>
> To suggest to heterosexuals that the secrecy policy will "accommodate" their privacy interests is to attempt to mislead them. They are not dunces or ostriches. They can hardly be unaware that because of the passage of the Act homosexuals are serving with them.

Id at 978. But the regulation need not assume that heterosexuals are either "dunces or ostriches." It might be true that bigoted officers will feel pressure to retaliate against any soldier whose queer orientation becomes common knowledge. Even if bigoted heterosexuals know the identity of homosexuals in the shower, preserving plausible deniability about their knowledge might foster unit "cohesion." However, to see this potential benefit from the policy does not mean that the benefits outweigh the (civil rights) costs.