

Commentary

Lying about Reservation Prices in Business Negotiation: A Qualified Defense

Alan Strudler

University of Pennsylvania, USA

This essay offers a philosophical defense of deception about reservation prices in business negotiation. Its discussion is prompted by arguments that Charles N.C. Sherwood makes in a recent issue of *Business Ethics Quarterly* and develops ideas I put forward in an earlier issue of *Business Ethics Quarterly*. The essay argues that although reservation price deception cannot be justified by appeal to the consent of negotiating parties, it can be justified by appeal to a separate but related notion, assumption of risk, as long as the assumption of risk occurs in a suitably fair context.

Key Words: negotiation, deception, consent, fairness

One might think that business negotiation is like a game. It has its own distinctive rules. One of these rules allows players to lie to one another. A person who plays the negotiation game consents to its rules. And therefore, lying is morally acceptable in business negotiation. Let's call the preceding argument the *argument from consent*. Albert Carr (1968) wrote an essay defending something like it fifty years ago. His essay has been much discussed but its ideas largely rejected (e.g., Koehn 1997; Dees and Cramton 1991; Sinnicks 2022). Although the scope of the proposed license to deceive is unclear in Carr's essay, it is clear enough in much of the literature that the essay prompted, a literature that focuses on deception regarding reservation prices, that is, on the worst price that one is willing to accept in a negotiation (Carson 1993; Dees and Cramton 1991; Strudler 1995). This essay discusses and defends lying about reservation prices in negotiation.

Charles N.C. Sherwood (2021) presents the most recent challenge to the argument from consent.¹ He says that the argument fails because, among other reasons, people do not consent to negotiation lies. I will interpret Sherwood's criticism as modal in its

¹ Sherwood does not limit his discussion to reservation price deception but also discusses deception about more material or important aspects of a negotiated transaction. To that extent, his discussion maps Carr (1968) but is broader than Carson (1993), Dees and Cramton (1991), or Strudler (1995). My focus here remains deception about reservation prices.

This article has been updated since its original publication. Please see DOI: <https://doi.org/10.1017/beq.2023.21>.

Business Ethics Quarterly 33:4 (October 2023), pp. 763–776. DOI:10.1017/beq.2023.1

Published by Cambridge University Press on behalf of the Society for Business Ethics.

© The Author(s), 2023.

core. It provides that people cannot consent to lying because doing so would be self-defeating, undermining the effectiveness of relevant communication in negotiation. Sherwood does not completely explain the mechanics of this self-defeat; he does not fully explain how it is that consent to lying would undermine communication. But Kantian moral theorists can assist him. Like Sherwood, some Kantian moral theorists say that consent to lying is somehow self-defeating (O'Neill 1985; Korsgaard 1986). I believe that their analysis, though not about negotiation, lays out the mechanics of relevant self-defeat more fully than Sherwood does, improving our position for making sense of the argument from consent.

I will maintain that both Sherwood (2021) and Kantian moral theorists err: it is possible to consent to lying without damaging the process of communication or the prospect of successful lying, and so the difficulties in consent to lying cannot be used to show the presumptive wrongness of lying. I end up devoting considerable space to the argument from consent, even though Sherwood and the Kantians reject that argument. One might wonder why I bother. I believe that there is insight in the argument from consent, insight that its critics may miss, and that something resembling the argument from consent can be made to work. The insight in the argument from consent is that acceptable deception should accommodate deference to the will of the deceived. I will suggest, however, that that insight gets better reflected in an “assumption of risk” account that I propose than in the argument from consent. Ultimately, I aim to show that even though we have no reason to think that people generally consent to lying in negotiation, there is good reason to think that people generally assume the risk of lying about reservation prices in some negotiations, and so at least some lying can be acceptable in negotiation as it is in poker.

This essay proceeds as follows. In section 1, I explain and critique Sherwood's (2021) argument that consent to lying cannot be part of a meaningful negotiation process. In section 2, I explain how Kantian moral theory can be used to unpack and strengthen Sherwood's analysis. In section 3, I explain and justify my “assumption of risk” account of the casuistry of lying and explain its emergence from the consent account. In section 4, I consider why we need negotiation lies. In section 5, I consider an objection to my account.

1. SHERWOOD ON THE POSSIBILITY TO CONSENT TO LYING

The argument from consent provides that deception in negotiation is acceptable because negotiating parties consent to such deception, just as poker players consent to deception in their game.² Describing his main objection against that argument as a *reductio ad absurdum*, Sherwood (2021) purports to demonstrate that accepting the idea that we may consent to lying in negotiation has implications so absurd that we must reject the idea. Sherwood's argument cannot be purely empirical. He suggests more than that experimental or anecdotal evidence shows difficulties in consenting to lies. Instead, he suggests that there is something about consenting to lying that must fail, that would itself be *absurd*. That suggestion is conceptual, and it will form

² Arthur Applbaum (1999) provides an incisive account of the relevance of consent and fairness in negotiation lies, although he does not draw the conclusions that I do.

part of my focus.³ (Sherwood does also offer empirical arguments against lying in negotiation, grounded in social science research for his claims about the consequences of deception, but a consideration of the empirical arguments would take me far astray.)

What, then, is Sherwood's (2021) *reductio* argument against the idea that we consent to being lied to in negotiation? During negotiation, he observes, people do not say that they consent to being lied to. They do not expressly consent to lies. If consent occurs in negotiation, it must therefore be implicit or tacit consent. However, Sherwood argues, tacit consent to lies does not occur in negotiation. Whenever tacit consent occurs, it can be made explicit without altering the meaning of what one communicates or affecting the transaction to which consent is made, Sherwood maintains. But if one tries to make explicit consent to lying in negotiation, "then the dialogue of negotiation becomes all but emptied of meaningful content" (7). It follows, Sherwood concludes, that people do not consent to lying in meaningful negotiation; consent to lying would destroy such negotiation. Sherwood acknowledges that, apart from consent to lying, making consent explicit can serve to clarify communication, without in any way undermining it. An example may illustrate his point. Consider a comparatively uncontroversial case of tacit consent. When I order food off a menu, I tacitly consent to pay. But if a restaurateur is suspicious of me, perhaps because he has been stiffed by customers in the past, he might require me to sign a form requiring me to pay for the food I order. His doing so would be strange, but it would not destroy the possibility of our transaction or undermine our discussion. It would instead clarify our communication about that transaction. So it makes sense to recognize tacit consent in the case of purchasing items off a restaurant menu but not in the case of lying in negotiation, Sherwood might say.

What should we make of Sherwood's (2021) conceptual argument that consent to lying undermines the possibility of communication about the subject matter of the lie because lying is inconsistent with the possibility of consent? To show why I do not find it convincing, I offer a counterexample, the recently invented parlor game *Liar, Liar*, a game in which lying is required by the rules.⁴ In *Liar, Liar*, players must make false statements about their own lives to one another with the intent to deceive; a player acquires points through successful lying; the rules expressly require players to lie, and without lying, there is no game. Lying is, then, constitutive of the game. If you understand the rules of *Liar, Liar*, then you understand that you cannot win without successfully lying to others and that you cannot play without inviting others to lie to you. Because players choose to play in awareness of rules that make lying constitutive, it seems to follow that people unavoidably consent to being lied to. The

³ As Colin McGinn (2017, chap. 3) explains (using the terms *a priori* and *a posteriori* in place of the terms *conceptual* and *empirical*), "tradition insists that knowledge falls into two broad classes: a priori knowledge and a posteriori knowledge. These categories are conceived as exclusive and exhaustive: no piece of knowledge is both a priori and a posteriori, and any piece of knowledge is one or the other. Without going into detail, we can characterize a posteriori knowledge as knowledge acquired by means of the senses ("by experience") and a priori knowledge as knowledge not so acquired, but rather acquired 'by reason alone' or 'intuitively.'"

⁴ <https://liarliargame.com/>.

form of consent is interesting. It does not get expressed verbally. Presumably, few players mouth the words “I consent” when joining the game. But that only shows that consent requires no words (Dougherty 2021). In general, people can consent through action. We do it regularly. If I stand at the doorway to your office and ask whether I may enter, and you respond by waving me in, then you consent to my entering your office by your action. Similarly, when a person joins a game that cannot be played without lying, she consents to playing by the rules, from which it follows that she also consents to being lied to. The dynamics of the game would not be affected if consent were instead made explicit. Indeed, if, halfway through a game of Liar, Liar, a person complains about being victimized by a lie, saying that he never agreed to being lied to, his complaint would be absurd.

So there is a logically possible game, a game that gets played in fact, in which people apparently consent to lying. The game is coherent; people can consistently play it, although we may presume that the level of trust is not high, just as one might presume that the level of trust is not high in some successful business negotiations. Why, then, think that consent to lying is somehow self-defeating? If such consent is consistent with the game of Liar, Liar, why not think it also consistent with business negotiation? The most rigorous possible answer to the questions, the most plausible rational ground for Sherwood’s (2021) concerns, can be found in Kantian moral theory.

2. THE CONCEPTUAL ARGUMENT LAID BARE

Whereas Sherwood (2021) focuses on how consent to lying would undermine communication in the negotiation process generally, Kantian moral theorists focus more narrowly on how consent to lying undermines lying itself and would render the consent impossible. Of course, if consent to lying is impossible, requiring consent would render the “dialogue of negotiation . . . all but emptied of meaningful content,” as Sherwood (2021, 7) suggests. The Kantian argument relies on the idea that consent is valid only if a consenting party is relevantly informed: you cannot meaningfully consent to something that you do not understand (O’Neill 1985; Korsgaard 1986). Suppose, for example, that I contemplate lying to you. That would create a problem for the possibility that you may validly consent to my lie, because doing so would require me to inform you about the truth of the proposition about which I will lie. Consider a simple version of the argument in which your only source of information required for informed consent to a lie comes from me, the prospective liar. If I inform you about the truth of the relevant proposition, then I cannot lie to you about the truth of that proposition, because lying involves the intent to deceive, and I cannot deceive you about the truth of a proposition when I have already disclosed to you the truth value of that proposition, something that I must have done to get your consent. Suppose, for example, that it now rains outside. I ask for your consent for me to lie to you about whether it rains outside. You give your consent, but then I falsely say that it is not raining outside. Because I asked for your consent on lying, you know that I may be lying, and hence my lie is ineffective. The information I provide in getting your consent makes it impossible for me to deceive you, and I defeat my

attempt to lie by disclosing that information. Because it is impossible for me to deceive you, and it is impossible for you to consent to the impossible, it is impossible for you to consent to being lied to—or so goes the argument. Valid consent to lying is impossible, an idea that should elicit Sherwood's approval.

The Kantian argument explains the inherent difficulty in getting consent to lying in certain circumstances; it captures an appeal in Sherwood's (2021) position. Its transparency also reveals its limits. The Kantian argument at most shows that if one gets consent to lie to a person about a particular proposition, then one cannot lie to her about the truth of that proposition. But consent to lying need not involve such narrow scope. A somewhat concrete illustration of limits in applying the Kantian argument follows: imagine that I ask your consent that I lie to you about one thing that I will say in the next five minutes, explaining that I will say quite a few things. You can coherently consent, and I can successfully lie to you, because although you know that I lie, you do not know which of the statements I make is a lie. It follows that nothing in the Kantian argument shows that one cannot consent to being lied to about the truth of a member of a set of propositions, even if one cannot consent to any particular lie. It seems to be the kind of consent needed for negotiation lies, because one might consent to the idea that some of the things one will be told are lies, not that all things will be lies.

In *Liar, Liar*, by joining that game, a player signifies her consent to being deceived. Indeed, to the extent that the consent is voluntary and informed, and advances each player's interest in a good time, I can find nothing objectionable about the lying that occurs. There is nonetheless an inherent difference between the lying that occurs in *Liar, Liar* and that which occurs in negotiation: in the case of the former, but not the latter, presumably, the truth gets revealed at the end of the game. But I see no reason to think that this difference matters either morally or logically. Both cases involve manipulation. In each case, if lies are acceptable, one must explain why.

Sometimes, for example, in *Liar, Liar*, lying licensed by consent seems not only logically possible but also morally acceptable. Why not also think that consent to lying is logically possible outside games, as in negotiation? Logical possibility need not correspond to realistic possibility. Do people in fact license lying by their consent in negotiation? The factual question poses difficulties concerning which Sherwood (2021) provides some alert. No doubt explicitly expressing consent to lying in negotiation, uttering the words "I consent" would alter the dynamics of the process. But more generally, explicitly expressing consent is itself often strange. If I appear at Herman's office door, obviously hoping to enter so that I might chat with him, it would seem to be an act of chilling pomposity, or a joke, for Herman to say "I give my consent for you to enter." So ordinarily, he should not give explicit consent. Consent should be expressed less directly, perhaps with a smile, perhaps by a directive, for example, saying simply "please come in." As Quill Kukla (2021) explains, consent need not occur through any particular discrete act. When two old friends meet and hug one another, no distinctive act can be interpreted as conferring a consent to hug. It hardly follows that the hugging is nonconsensual. In the case of our two huggers, one can infer consent from the fact that the two

friends, in Kukla's words, "act in ways that express their own values, desires, and purposes" (273). Their consent, like much consent, must be extracted from a gestalt of interaction, and there is no mechanical way to do it. That complicates the inference to consent, demonstrating that it involves judgment and is nothing mechanical. The factual question of whether consent to lying in negotiation occurs may thus prove challenging.

Questions of whether consent to lying occurs in negotiation appear even more challenging in light of the fact that one may consent to engaging in an activity even if one does not consent to all aspects of that activity (Hersch 2020; Katz 1990). Suppose, for example, that, desperate for a ride, you ask Marcello to drive you in his car, thus unequivocally expressing your consent to be driven, while both realizing that sometimes he drives drunk and hoping that he is not now drunk. It turns out he is drunk. Nonetheless, when Marcello negligently crashes the car and injures you, you do not lose your ground for complaint against him, despite your consent to drive with him; you did not consent to his recklessness or drunkenness. The mere fact that one consents to an activity does not entail that one consents to the ugly aspects associated with that activity (Hersch 2020; Katz 1990). Similar reasoning is available regarding lies in negotiation. Many view the prospect of deception as an ugly aspect of negotiation; for example, going into a car dealership and buying a car is often considered an ugly aspect. Even if we consent to negotiating with the car dealer, it no more follows that we consent to her lies than that you consent to Marcello's drunk driving. We cannot, then, infer that you consent to being deceived from the mere fact that you play the car-buying game. If one hates some aspect of what one does but does it anyway because it is a practically unavoidable aspect of something that seems worthwhile, the scope of one's consent need not extend to that aspect.

Here is where we stand. Sherwood's (2021) *reductio* against consent to lying is unpersuasive. Kantian moral theory unpacks Sherwood's account, providing a basis on which to assess its merits, but provides no vindication. It seems, then, possible to consent to lies. Still, we lack reason to think that people in fact consent to negotiation lies. Nothing suggests that people want to be lied to in business negotiation or that they try to communicate their consent to lies. Yet consent requires some sort of communication of an intention to consent (Dougherty 2021). We cannot, therefore, count on the argument from consent to conclude that lying is acceptable in negotiation, because we lack assurance that such consent occurs.

We are not left in limbo regarding the acceptability of lying in negotiation, I will argue. The acceptability of lying does not depend on consent. Here, again, the game analogy is worth exploring. Perhaps lying in games is acceptable without the players' consent, and if so, perhaps lying in negotiation is similarly acceptable.

3. ASSUMING THE RISK OF LIES

Consider Puritanical Paolo, who watches with disgust a poker game in which people lie to each other. To prove to the putative miscreants that honesty is the best policy, Paolo joins the game, but only after announcing that the other players act wrongly by

lying, instructing the players that they should not lie to him and that he does not consent to their lying. Paolo notices that the players laugh dismissively at his instruction, but he joins the game anyway. Paolo does not consent to lies. One consents only if one intends to convey that one consents (Dougherty 2021; but for a contrary view, see Alexander, Hurd, and Westen 2016). Paolo plainly has no such intention. Yet it still seems acceptable to lie to him.⁵

A familiar concept explains the acceptability of some lying in poker and negotiation, as well as the acceptability of some tackling in American football, but does not rely on the concept of consent: the concept of assumption of risk (Dorfman 2014).⁶ When a person knowingly and voluntarily exposes herself to a risk, there is a presumption that she cannot reasonably complain when the risk materializes, even if she did not consent to the risk: she assumes the risk. I will suggest that a person assumes a relevant risk when two conditions are satisfied: first, that she knowingly and voluntarily exposes herself to that risk, and second, that it is fair to impose that risk on her. In Paolo's case, he understands the risk he takes by joining the game, and so he satisfies the first condition. Moreover, it would be unfair to expect the other poker players to change their game to accommodate Paolo, and so he satisfies the second condition. Assumption of risk thus differs from consent in a crucial respect. Only the latter inherently involves communicative intent. One can assume the risk privately. If you post a sign saying "DO NOT ENTER: DANGEROUS WILD BEASTS ROAM BEHIND THIS FENCE," and I, on a lark, secretly enter anyway, then I may assume the risk of getting attacked by the beasts. But for me to consent to the attack, I must somehow try to convey that to you.

Assumption of risk is most commonly discussed in legal contexts, but its roots, many argue, are in morality (Dorfman 2014; Katz 1990; Feinberg 1988). In relevant respects, the ideas of consent and assumption of risk occupy a moral continuum. Whether a person consents to a risk or simply assumes that risk, it is her choice that

⁵ Applbaum (1999) considers a case similar to Paolo's case, though involving a different game, American football. He asks us to suppose that a person joins a game of football in which tackling is allowed but wears a sign on his shirt saying that he does not consent to being tackled. Does Applbaum's football player consent to being tackled? Applbaum himself declines to answer that question. But it seems clear enough that the football player does not consent. At best, the football player equivocates on consent, but equivocal consent is no consent at all.

⁶ As Simons (2002) explains, there is a good deal of skepticism in legal circles about the viability of the concept of assumption of risk; he provides an objective review of such skepticism. (For a lucid defense of skepticism about assumption of risk, see Sugarman, 1997.) It is now perhaps the majority view among scholars and judges that the concept of assumption of risk should be, or already has been, replaced by concepts like comparative negligence. Simons (2002) plausibly defends the concept of assumption of risk against skepticism, in my view, and also argues that attempts to replace it implicitly invoke it. Dorfman (2014) explains that much of the skepticism about assumption of risk arises from the perception that it is a politically regressive doctrine, aimed at shifting liability away from corporations and toward consumers and workers. Properly understood, Dorfman argues, assumption of risk rejects such a shift and is instead a progressive doctrine aimed at achieving a fair balance between the powerful and the less powerful. Dorfman himself offers "a liberal-egalitarian account of the doctrine, arguing that the *fact* of making a choice (to assume a given risk) is not sufficient to *justify* the shifting of responsibility from the negligent injurer to the choosing victim. For it is also necessary that the latter must be acting under conditions that render thus shifting fair" (315).

determines the consequences she faces. Both the idea of consent and the idea of assumption of risk thus recognize the moral importance of respect for a person's right to choose or for his will. In negotiation, we cannot show respect for a person's will by considering whether he consents to lying—demonstration of such consent is typically infeasible—but we can consider whether he assumes the risk of lying.

I suggest that when a person joins a football game, knowingly and voluntarily exposing himself to the risk of being tackled, he presumptively assumes the risk of being tackled, even if he announces that he does not consent to being tackled. Similarly, when a person knowingly and voluntarily negotiates with a car salesperson, he presumptively assumes the risk of being lied to about reservation prices, even if he does not consent to the lies. Yet not all knowing and voluntary *exposure* to risk should be regarded as *assumption* of risk. Sometimes it seems unfair to impose risk on a person even if she knowingly and voluntarily exposes herself to it. Consider a factory worker, Ida, who knowingly and voluntarily exposes herself to a risk of injury by disabling an easily defeated safety device on a piece of dangerous equipment, and as a result, she injures herself. Courts have declined to invoke the concept of voluntary assumption of risk to bar the worker from recovering from injuries such as hers, instead declaring that equipment manufacturers should have made a more foolproof safety device (i.e., *Vanskike v. ACF Industries, Inc.* [1981]), a judgment that strikes me as reasonable. Still, if the management had gone to great lengths to install a safety device and it was only because of Ida's extraordinary deviousness and persistence that she was able to disable the safety device, it would be reasonable to say that Ida assumed the risk of injury. That suggests that beyond knowingly and voluntarily exposing oneself to risk, in order to assume a risk, it must be fair to assign that risk to the person who exposes herself to it. As Dorfman (2014) plausibly argues, considerations of fairness determine whether a person's intentional exposure to risk constitutes a voluntary assumption of risk. It is unfair, on one hand, to make Ida absorb the cost of her injury by treating her as assuming the risk when management could have feasibly prevented her injury by using a better safety device; if, on the other hand, management makes the device safe enough, Ida's action should be regarded as assuming the relevant risk. What constitutes "safe enough" is a difficult line-drawing problem that cannot be resolved here. That difficulty aside, what reflection on these cases shows, I maintain, is that in knowingly and voluntarily exposing oneself to risk, one can thereby assume that risk, including costs that arise when the risk materializes, but only if it is consistent with fairness.

Is the imposition of risk of deception unfair in the case of lies about reservation prices in negotiation? Typically not, I will argue in the next section. Still, it seems clear that lying about reservation prices is not always fair and that there is more than a single way in which unfairness can occur. A naive and gullible negotiator may not appreciate the prospect of reservation price lies, and so it may be unfair to regard him as knowingly assuming the relevant risk. In that case, the sophisticated negotiator risks wronging the naive negotiator, a risk that she should avoid and something that makes the job of a salesperson so morally treacherous. A seller's lies about material aspects of a deal may pose very different fairness problems. Such lies can undermine the prospect of buying what one actually thinks one is buying, making it virtually

impossible for one to get a fair deal and thus for the buyer to assume the risk of negotiation deception.⁷ In the case of the naive and gullible negotiator, it is unfair to treat him as assuming the risks of deception, because through no fault of his own, he did not understand the risks he faced. In the case of a buyer who would be deceived by material lies, it is unfair to treat him as assuming the risks of deception, because the costs associated with those risks fall disproportionately on the buyer. But not all reservation price lies are unfair.

4. WHY LIE ABOUT RESERVATION PRICES?

Reservation price information is crucial to successful negotiation. The information can lead one to know that there is a point in moving forward in negotiation: that a seller's reservation price, for example, is not so high that a buyer cannot afford to pay it. Also, information about reservation prices may help one understand how far one can press the other party to make concessions in price. Why shouldn't each party to a negotiation, then, be forthcoming about her own reservation price? Because one may not know that one's candor will be reciprocated, and if it is not, one may suffer a competitive disadvantage as one's bargaining foe presses for an outcome and one is at a loss about how hard to press back. Both parties to a typical negotiation suffer the reservation price problem; it is, in that sense, collective. The solution to the problem, I maintain, involves offering statements that are cautious enough to protect one's bargaining position but helpful enough to allow one's bargaining foe to make inferences that allow her to reach reasonable conclusions about one's bargaining position. Telling lies and making misleading statements, if done right, can put one's bargaining foe in a position to make inferences that allow her eventually to infer some reasonable approximation of the truth about one's reservation price. These lies put each party in position to move forward with the negotiation while having a reasonable prospect of getting adequate information about the other party's reservation price. The distribution of both burdens and the prospect of benefits is fair. Reservation price lies can form part of a fair solution to a bargaining problem, I will maintain.

When a person does not want simply to disclose his reservation price in a negotiation but still wants the negotiation to move forward, what can he do? I once offered this possibility:

A few years back, I went to Madrid for an extended stay. Two weeks after I arrived, I had to move from a hotel to an apartment. There were about ten boxes that I needed to move from the hotel, and I was confident that a taxi could handle the job, since one had done so for the trip from the airport to the hotel. Because there was no cab to be seen near the hotel, I had to go a few blocks to find one. When we arrived at the hotel, [the taxi driver] feigned shock at the sight of the boxes, and said that they were far too many for him to carry. I told

⁷Material lies to a seller may be similarly unfair and wrong, for example, when a buyer lies about her ability to pay on time. For discussion of the potential for wrongness in a buyer's deception, see Strudler (1997). The most famous legal case confirming that buyer nondisclosure can be wrong is *Laidlaw v. Organ* (1917).

him that I would expect to pay “extra” because of all the boxes, but he just shook his head, declaring that there were too many boxes. At that point, I gave him the money for the fare on his meter and said that I would find some other way to get the boxes to my apartment. (I ended up renting a car.) He looked incredulous and said, “Con usted no se puede razonar,” meaning, roughly, that it was hopeless reasoning with me. At this point, I finally realized that his declaration of the impossibility of moving the boxes was simple posturing, that he expected me to know this, and that I was being viewed as foolish, for wasting both our time by refusing to bargain. I felt too perplexed to continue; he left, cursing me (Strudler 1995, 817).

By saying that there were far too many boxes for him to carry, I take it that the taxi driver was suggesting that his reservation price was so high it could not be paid, that there was no amount of money for which I could get him to move the boxes because it was impossible. Apparently, his statement that there were too many boxes for him to carry was a lie, and the suggestions he made about his reservation price were false. I believe that the lie conveyed useful information dressed in drama, namely, that the cost of moving the boxes would be extraordinarily high.

As I see it, the taxi driver told an informative lie regarding his reservation price. He was playing a negotiation game, lying in a way that invited me to respond, hoping to trigger a conversational back-and-forth aimed at eventually producing some approximation of the truth regarding his reservation price. For example, I might have played along with the taxi driver, suggesting that if he could not move the boxes in one trip, perhaps he could make two trips, or he could call in a colleague for assistance. After getting a sense of how much the alternatives would cost, the taxi driver might have been willing to revisit the idea of making one trip, perhaps using the price of two trips as a new baseline for negotiation. Perhaps there were other ways that the taxi driver could have made the point that the cost of moving the boxes would be unusually high, but I see no reason to think that the taxi driver’s lying approach was wrongful. Negotiation lies form part of a fair solution to a cooperation problem: both negotiating parties come out ahead by exchanging some information about reservations prices, but each party risks predation if she unilaterally discloses reservation price information.⁸ By engaging in the cooperative practice of allowing mutual deception about reservation prices, one gets a license to lie about reservation prices and thus gets the benefit of protection against predation, while at the same time preserving a path for the mutual exchange of information. No doubt Sherwood (2021) would see things differently. He, too, recognizes the importance of cooperation in negotiation, but his focus is on the need to build trust rather than the need to get information. No doubt there are negotiations in which building trust has paramount importance, for example, when the negotiating parties contemplate the prospect of a long-term relationship (Jennejohn 2020). But the relationship between the taxi driver and me was apparently one-shot, not long term. It is reasonable to suppose that the prospects of building a trusting relationship were dim and that it made practical sense

⁸To be clear, predation may occur here because one party takes unfair advantage of her asymmetric knowledge of the other party’s reservation price.

to instead aim for the more modest objective of closing the deal on a taxi ride, which required some sort of cooperation in the exchange of relevant pricing information. But Sherwood has a larger point.

It is clearly within the spirit of Sherwood's (2021) argument to contend that the reservation price lies defended here in fact corrode negotiating relationships, undermining trust and harming the prospect of good deals. To the extent that that contention is empirical, its resolution is beyond the scope of satisfying resolution here. Still, ample reason supports doubt that reservation price lies need be corrosive.⁹ Certainly many negotiation experts disagree with Sherwood's assessment. When Robert Frank (1988, 165) declared that "the 'art of bargaining' is largely the art of sending misleading information about one's reservation price," he relayed the impression that reservation price deception is the norm, which suggests that in his experience, it was typically part of successful negotiations. Similarly, when Richard Shell (1991, 93) remarked that "in nearly all bargaining encounters, a key skill is the ability to communicate that you are relatively firm on positions when you are, in fact, flexible—in short, to bluff about your intentions," he conveyed the thought that deception about one's bottom line is typically part of successful negotiation, not anything destructive. I have already explained why I think something good comes out of a certain deception in negotiation: it helps to solve an information problem that would otherwise stand in the way of success. But why wouldn't such deception have a corrosive impact?

People expect deception about reservation prices and know how to deal with it. As Sarah Buss (2005) explains more generally, we know how to live with deception, how to flourish with it, how to build on it. Taking a cue from Søren Kierkegaard (2013), Buss (2005) suggests that deception pervades courting; it provides the elements of seduction. At the start of a romance, people put on a good face, contriving to create the appearance that they are particularly handsome, beautiful, thoughtful, considerate, cultured, and so on—more so than they really are. People do not parade their physical and emotional flaws but instead actively hide them under fancy clothes and antianxiety medications and carefully selected stories of personal glory, which is deceptive. Yet such deception, as long as it is about minor and expected things, not about one's criminal past or bigamous intentions, often (though not always) serves as the start of a good relationship, stoking romantic fires, whereas a forthright disclosure of one's flaws would smother the fire. Not all deception is the same: some of it is corrosive, some not. In our ordinary lives, whether in romance or in business, we learn to live with it and how to draw a line between acceptable and unacceptable deception. A dismissal of all deception among competent, capable adults as corrosive is overly simple. People

⁹To be clear, I am not arguing that lying about reservation prices is always not corrosive. In contexts in which trust has been established and there is a confirmed expectation of earnestness and truthfulness, any lying might be corrosive. However, in some cases, successful negotiation depends on something other than trust. As trust is not present, it does not get corroded. I do not assume that Shell and Frank agree with me, in the passages quoted earlier, about the moral justifiability of deception or that their quoted remarks show that deception is acceptable. I appeal to these authors solely for the point that deception can form part of successful negotiation.

can take deception in games lightly. I agree with Buss that we can do the same outside games. No doubt the stakes differ so substantially in parlor games and in business practice that one might think that tolerance for deception in the former should far exceed tolerance for deception in the latter (Hersch 2020). That thought gets its plausibility from reflection on the kinds of deception that seem inherently destructive, for example, deception involving materiality or exploitation of the naive. In the case of deception about reservation prices, the destructive prospects seem comparatively small and the potential benefits substantial.

5. HOW LYING CAN DELIVER TRUTH

I have maintained that the Madrid taxi driver reasonably lied as a way of conveying a truth. Following Dees and Cramton (1995), Sherwood (2021) thinks that the description makes no sense. How can one deliver truth with a lie? Here I think that Sherwood, along with Dees and Cramton, makes a simple mistake. One can tell a lie as a vehicle for delivering the truth, but the truth will not be the proposition one expresses in lying. Successful communication can occur in lying about reservation prices, not in the sense that one gets an audience to accept the face value of what one says, but because one succeeds in indirectly delivering information. The explanation of this possibility is simpler than one might think.

How can the taxi driver deliver a truth with a lie about his reservation price? The easiest way is if the process is indirect. Indirection might work in many ways. Here is one: after telling me that he could not possibly take the boxes, the taxi driver might have expected me to back off his plan to have all the boxes carried at once and propose that the driver make two trips. Perhaps on the verge of agreeing that I pay the expensive price for two trips rather than one, the taxi driver could say that because I was willing to pay so much for two trips, he can find a way to handle all the boxes in one trip. The driver's initial lie about his ability to take all the boxes need not be the end of a conversation but can be part of a process to dramatize to me that the ride would be very expensive. Even more simply, the lie could serve as a signal to me that the proposed ride would be very expensive, which is a true proposition, but not the proposition that the taxi driver expresses in his lie.

It seems, then, that one can convey a truth about one's reservation price in more than one way. Most simply, one conveys it openly and wholeheartedly. But such direct expression in business negotiation may leave one vulnerable to predation. So the cautious businessperson may lie and deceive about her reservation price, allowing the truth to emerge slowly, only after doing so seems safe.

6. CONCLUSION

Poker players lie about their cards. It is part of the game. You may not like it. Lying may bother you. But there is nothing inherently wrong with their lying. When you enter the game, you know about the practice of lying and that it is part of what makes the game enjoyable for most players. Without the possibility of deception, there might be no poker game for you to play. So you assume the risk that people will lie to you,

and you have nothing to complain about when they do. You also have nothing to complain about when a car salesperson lies to you about her reservation price. She protects herself in a way that is equally and fairly open to you. Lying about reservation price can be part of a process that fairly works to the advantage of negotiators.

In this article, I have argued that even though lying and other deception about reservation prices typically cannot be justified by appeal to considerations of consent, they may be justified by appealing to considerations of assumption of risk. I have also argued that in contexts in which the prospects of establishing trust are slim, particularly but not exclusively in one-shot negotiations with experienced negotiators, deception about reservation prices can be morally acceptable.

Acknowledgments

I thank the anonymous referees for their advice, and I thank associate editor Denis Arnold for his patience.

REFERENCES

- Alexander, Larry, Heidi Hurd, and Peter Westen. 2016. "Consent Does Not Require Communication: A Reply to Dougherty." *Law and Philosophy* 35 (6): 655–60.
- Appelbaum, Arthur. 1999. *Ethics for Adversaries: The Morality of Roles in Public and Professional Life*. Princeton, NJ: Princeton University Press.
- Buss, Sarah. 2005. "Valuing Autonomy and Respecting Persons: Manipulation, Seduction, and the Basis of Moral Constraints." *Ethics* 115 (2): 195–235.
- Carr, Albert. 1968. "Is Bluffing in Business Ethical?" *Harvard Business Review*, January–February, 143–53.
- Carson, Thomas. 1993. "Second Thoughts about Bluffing." *Business Ethics Quarterly* 3 (4): 317–41.
- Dees, Gregory, and Peter Cramton. 1991. "Shrewd Bargaining on the Moral Frontier: Toward a Theory of Morality in Practice." *Business Ethics Quarterly* 1 (2): 135–67.
- Dees, Gregory, and Peter Cramton. 1995. "Deception and Mutual Trust: A Reply to Strudler." *Business Ethics Quarterly* 5 (4): 823–32.
- Dorfman, Avihay. 2014. "Assumption of Risk, after All." *Theoretical Inquiries in Law* 15 (2): 293–327.
- Dougherty, Tom. 2021. *The Scope of Consent*. Oxford: Oxford University Press.
- Feinberg, Joel. 1988. *Harm to Self*. Oxford: Oxford University Press.
- Frank, Robert. 1988. *Passions within Reason*. New York: W. W. Norton.
- Hersch, Gil. 2020. "You Can Bluff but You Should Not Spoof." *Business and Professional Ethics Journal* 39 (2): 207–24.
- Jennejohn, Matthew. 2020. "Braided Agreements and New Frontiers for Relational Contract Theory." *Journal of Corporation Law* 45 (4): 101–30.
- Katz, Leo. 1990. "The Assumption of Risk Argument." *Social Philosophy and Policy* 7 (2): 138–69.
- Kierkegaard, Søren. 2013. *The Seducer's Diary*. Princeton, NJ: Princeton University Press.
- Koehn, Daryl. 1997. "Business and Game-Playing: The False Analogy." *Journal of Business Ethics* 16: 1447–52.

- Korsgaard, Christine. 1986. "The Right to Lie: Kant on Dealing with Evil." *Philosophy and Public Affairs* 15 (4): 325–49.
- Kukla, Quill. 2021. "A Nonideal Theory of Sexual Consent." *Ethics* 131 (2): 270–92.
- Laidlaw v. Organ, 15 U.S. (2 Wheat.) 178 (1817).
- McGinn, Colin. 2017. *Philosophical Provocations: 55 Short Essays*. Cambridge, MA: MIT Press. Kindle.
- O'Neill, Onora. 1985. "Between Consenting Adults." *Philosophy and Public Affairs* 14 (3): 252–77.
- Shell, Richard. 1991. "When Is It Legal to Lie in Negotiation?" *Sloan Management Review*, April 15, 93–101.
- Sherwood, Charles N.C. 2021. "A Lie Is a Lie: The Ethics of Lying in Business Negotiations." *Business Ethics Quarterly* 32 (4): 604–34.
- Simons, Kenneth. 2002. "Reflections on Assumption of Risk." *UCLA Law Review* 50: 481–529.
- Sinnicks, Matthew. 2022. "On the Analogy between Business and Sport: Towards an Aristotelian Response to the Market Failures Approach to Business Ethics." *Journal of Business Ethics* 177: 49–61.
- Strudler, Alan. 1995. "On the Ethics of Deception in Negotiation." *Business Ethics Quarterly* 5 (4): 805–22.
- Strudler, Alan. 1997. "Moral Complexity in the Law of Nondisclosure." *UCLA Law Review* 45: 337–84.
- Sugarman, Stephen. 1997. "Assumption of Risk." *Valparaiso University Law Review* 31 (3): 833–78.
- Vanskike v. ACF Industries, Inc., 665 F.2d 188 (1981).

. . .

ALAN STRUDLER (strudler@wharton.upenn.edu) is professor of legal studies and business ethics at the Wharton School, University of Pennsylvania. He writes about normative moral theory, deception in business, and corporate responsibility. His writings appear in various journals, including *Business Ethics Quarterly*, the *Journal of Political Philosophy*, and the *Journal of Philosophy*.