The Ethics of Negotiation

Charles N.C. Sherwood

London School of Economics and Political Science

A thesis submitted to the Department of Philosophy, Logic and Scientific Method of the London School of Economics and Political Science for the degree of Doctor of Philosophy, London.

Submitted: August 11th, 2023
Minor Amendments: November 29th, 2023

DECLARATION

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work.

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 64,058 words (excluding bibliography).

STATEMENT OF INCLUSION OF PRIOR WORK

An early version of Chapter 3 of this thesis formed my dissertation as part of my MSc degree in Philosophy & Public Policy at the London School of Economics and Political Science. A much more recent version of the chapter was published in 2022: *Business Ethics Quarterly 32*(4):604-634.

ACKNOWLEDGMENTS

I would like to thank my current supervisors, Jason Alexander and Alex Voorhoeve, and my previous supervisors, Mike Otsuka and Susanne Burri, together with David Coombs, for their encouragement and insightful feedback during the researching and writing of this thesis. Chapter 3 also benefited from the constructive contributions of *Business Ethics Quarterly*'s former editor in chief, Bruce Barry, and three anonymous reviewers. I would also like to acknowledge the tremendous support from my family.

ABSTRACT

I argue for the following:

(1) Negotiation is a *social construct* whereby the parties seek to resolve differences and promote cooperation through a process rooted in *mutual consent*.

(2) Negotiation is a feature of *imperfect* markets. Where there is a negotiated agreement, there is typically a significant cooperative surplus and, instead of a single market price, a range of possibilities as to how that cooperative surplus may be divided between the parties. The purpose of negotiation is to determine exactly where in that range agreement is reached.

(3) The fairness of a negotiated agreement is wholly determined by the fairness of the process that leads to that agreement; and the fairness of the process is itself wholly determined by the adequacy of the constraints applied to that process. This makes negotiation an instance of *pure procedural justice*. In this regard, negotiation is to be clearly distinguished from judgment or binding arbitration.

(4) Negotiated agreement itself implies some degree of consent. But beyond that there are pro tanto obligations on negotiators to (a) *behave honestly* and, in particular, not lie, (b) *seek mutually beneficial outcomes*, whereby their counterparties are left no worse off than they would be in the event of a failure to agree, and (c) *avoid the abuse of extreme asymmetric negotiating leverage*, that characterises a combination of monopoly, monopsony or cartel with essential goods and services, and discriminatory pricing.

(5) These constraints are both *necessary* and *sufficient* to deliver a fair process and, therefore, a fair outcome.

(6) Such constraints can be seen, in part, as attempts to counter *market failure* and compensate for shortfalls in *consent*.

(7) A further constraint, related to the broader demands of *distributive justice*, is *not* called for. Such a requirement is the primary responsibility of society as a whole. To the extent that it does fall on any individual, it takes the form of a separate and narrower duty to aid that can, in cases of dire need, take precedence and require that negotiations cease.

(8) The fact that these moral constraints leave the outcome of most negotiations less than fully determined is not a weakness, but a strength of this account. It properly captures the character of a negotiation as a cooperative but also partially adversarial process, during the course of which a fair outcome is *constructed* – importantly, an outcome that is *not predetermined*.

CONTENTS

Chapter 1. Introduction & Overview

Under ideal conditions there is no such thing as a negotiation: such ideal conditions imply a perfectly competitive market, a perfectly competitive market generates a single price, and a single price leaves no room for negotiation (Wertheimer 1996, 61, 139, 217-218; Elegido 2009, 30). An examination of the ethics of negotiation therefore requires an embracing of non-ideal conditions with all the implied imperfections. This is an untidy environment that does not always offer singular solutions. Sometimes instead there are multiple solutions depending on the route by which they are reached.

Game theory's so-called 'bargaining problem' is a case in point. The typical game theorist's approach to a negotiation is to look for a single pre-determined fair solution. But I suggest that no such thing exists, at least not in isolation. Instead, I propose an alternative approach: that we treat negotiation as a social construct, designed to promote cooperative behaviour based on mutual consent; and that we seek to locate its fairness in the negotiating process itself. On this view of a negotiation, a fair solution is whatever emerges from a fair process; and, in turn, the fairness of any such process is wholly determined by the constraints applied to that process. Thus described, negotiation is an instance of pure procedural justice.

Having laid out the arguments for the above (Part 1), I go on to examine what are those constraints on a negotiation that are required to render it fair (Part 2). Besides the implicit condition of consent, I conclude that there are pro tanto obligations on negotiators to (1) behave honestly, (2) seek mutually beneficial outcomes, and (3) avoid the abuse of extreme asymmetric negotiating leverage. Furthermore, I argue that these above constraints are not only necessary, but are also sufficient, to ensure a fair process and therefore a fair outcome. In particular, I reject any additional constraint aimed at delivering the broader demands of distributive justice (Part 3).

Before I attempt to construct the above argument, let me prepare the ground in a little more detail. First, what do we mean by negotiation? The term is used very broadly. I intend to employ it here in a somewhat narrower form. In what follows, I consider negotiation to be a valuable social construct, whereby the parties involved seek to agree terms of cooperation, either directly with each other, or through agents. Negotiation is thus an institutional arrangement, in the broadest sense of those terms. It is by no means the sole method of resolving differences or establishing terms of cooperation, but it is an important one that we rely upon extensively.

Essential to this understanding of negotiation is the element of consent. The goal of negotiation is an agreement, and such an agreement implies some degree of consent. That consent may, in certain circumstances, be less than complete but it must exist in some form for the procedure to qualify as a negotiation. This requirement for consent in turn entails a right to decline, to 'walk away'. Such a right to decline is thus integral to the nature of a negotiation. The moment it is surrendered, there is no longer a negotiation, but something else instead: a judgment, a binding arbitration, or an enforced order. In these latter cases, one party alone imposes an outcome upon the other, or an independent party imposes an outcome on both.

Such an imposition may be the result of a negotiation, but it is not itself a negotiation. For example, two parties might negotiate over the selection and appointment of an independent binding arbitrator. The process of establishing the terms of that appointment would very likely qualify as a negotiation. But once the arbitrator is appointed, the negotiation has ended, and a process of judgment has commenced. At this point, the parties have surrendered their unilateral rights to consent or decline. They can no longer negotiate. Instead, they must await judgment.

Another significant characteristic of a negotiation, as here conceived, is that it seeks a resolution through a reconciliation of conflicting objectives. There is an essential tension. If, to the contrary, the parties are in full agreement on everything, then there is nothing to negotiate. Equally, if one party is prepared to give up on its objectives unilaterally and comprehensively – to simply 'roll over' – then again there is no place left for negotiation. Negotiation requires some desire on the part of each party to achieve an objective that does not, at least initially, appear fully compatible with the other party's objective. There are, of course, negotiations that are constructive and cooperative, but to qualify as a negotiation there must be *some* point of contention, something to negotiate over. The outcome is thus the result of a form of conflict, however mild this might be.

In a healthy and productive negotiation, a biproduct of that conflict is an exploration of value. This may or may not be the deliberate intent of the parties involved. In both cases, though, it regularly happens that the apparently conflicting demands of opposing parties reveal ways in which additional value can be created for the benefit of all. One party concedes something of special value to the other party, in return for something to which they themselves attach special value.

This is an investigation into what makes a negotiation fair. To be clear, the focus is on what makes a particular negotiation fair. This is a matter of practical ethics. Negotiation is ubiquitous – in the words of a popular bestseller on the subject, "life is negotiation" (Voss 2022 [2016], 16). It is among the most widely practiced social constructs for resolving conflict, promoting cooperative behaviour, and securing the gains of trade. Even within the formal legal system, most cases are settled pre-trial or pre-judgment through negotiation (Hollander-Blumoff 2010, 383, 385). Negotiation is thus such a pervasive part of social coordination, of our daily lives, that to question whether it is, as a general practice, justified or not, fair or unfair, a good thing or a bad thing, would seem fruitless. It is simply there. And most of us would accept that this social construct of negotiation is vital to our peace and prosperity, and thus essential to human flourishing. We can no more easily do without it than we (or at least most of us) can do without courtship. Courtship can be conducted fairly or unfairly. But there is little purpose in debating whether the practice of courtship is itself fair or unfair. In most societies, it is not something that can be sensibly dispensed with. Negotiation is much the same. So, I will take it as understood that negotiation is a fair, that is, a morally acceptable *practice*. Instead, my focus in what follows will be on those considerations that make a *particular* negotiation fair or unfair.

Negotiation is not only a critical part of our personal lives; it also underpins most of our productive activities, especially as they relate to the world of business. The extensive and often contentious debate on the morality of the marketplace has tended to focus above all on traded markets, viewed in many cases through the lens of perfect competition: a large, impersonal collection of unrelated buyers and sellers, working in aggregate to collectively determine a single, market-clearing price. But, as James Michelman (1983, 256) comments, "for the most part, this is not the way the commercial world works. What we actually find are nominated buyers dealing with nominated sellers over prices that are not fixed, but negotiable. It is these negotiations and negotiators with which we are finally concerned." And, indeed, it is those negotiations and those negotiators that are the subject of what follows.

This is the world of largely *im*perfect markets, and it is in that very imperfection that we find many of the most pressing ethical issues. Perfect markets by definition are transparent (there is no lying), Pareto optimal (there are no 'net losers') and highly competitive with many buyers and sellers (there is no room for exploitative monopolies, monopsonies or cartels). Negotiation, in contrast, is less than perfect, and often highly uncertain in its outcome. It is a rich, if somewhat unkempt, arena in which to explore many of the wider ethical issues that face those in business or related professions.

Looking at these ethical concerns through the eyes of a negotiator offers, I believe, a new perspective on many of the familiar issues that have characterised the debate on fairness: market failure, consent, exploitation, and distributive justice. Let me offer one, perhaps overly simplistic, illustration. An egalitarian might argue that fairness requires that the benefit (the entire cooperative surplus and perhaps more) of any negotiation go in full to the least well-off party, until such a point is reached where the parties are equally endowed. But what behaviour does this require of the negotiators? Any answer would seem to have little, if anything, to do with negotiation. Is the egalitarian, then, asking the negotiators to cease negotiating? And is this really intended as practical moral guidance on negotiation? Seen from the viewpoint of a party seeking to negotiate, the challenge of establishing fairness appears rather differently. Our negotiators thought they were negotiating, but now it seems the better endowed negotiator is being asked to do something entirely different, to unilaterally surrender all the benefits from the transaction. In other words, the better endowed negotiator is asked to play a role closer to that of a benefactor; and commensurately the lesser endowed negotiator a role closer to that of a supplicant. The positions of negotiator on the one hand, and benefactor or supplicant on the other, represent markedly different perspectives from which to view this question of fairness. Seeing the world through the eyes of a negotiator thus adds a further dimension to the enduring debate as to what makes any exchange fair.

Central to my analysis is the claim that the fairness of any negotiated agreement, commercial or otherwise, is wholly determined by the fairness of the process that leads to that agreement; and that the fairness of the process is itself wholly determined by the adequacy of the constraints applied to that process. In short, that negotiation is an instance of pure procedural justice. I argue that such a claim is consistent both with the theoretical work of philosophers as opposed in their thinking as John Rawls and Robert Nozick; and with the regular modus operandi of business practitioners (Chapter 2).

In this regard, I draw a contrast between two markedly different procedures: on the one hand, the traditional concept of *judgment*, which seeks to discover a single fair solution, based on some pre-existing and independent criteria, whether that be the law or some other broad moral principle or principles (what Rawls (1971, 84-90) calls "imperfect procedural justice"); and on the other hand, *negotiation*, which, within accepted constraints, generates an agreed outcome from a range of potentially fair solutions, that agreed outcome being deemed fair to the extent that the process that led to it was fair ("pure procedural justice"). Both these procedures are valid. That is, they each have their appropriate role. But they are different. The judge's task, in applying imperfect procedural justice, is to identify and *reveal* the fair solution, a solution that in a sense pre-exists judgment. By contrast, the negotiator's task, in applying pure procedural justice, is to *construct* that fair solution by reaching agreement with another party in a fair process. To suppose, as many game theorists do, that a negotiation has a single, pre-existing fair solution is of course to confuse that negotiation with the traditional concept of judgment.

I say 'traditional concept' of judgment because there is an alternative school of thought that sees judgment as a form of legal constructivism (Nieto 2021). My objective here though is not to opine on the merits or otherwise of legal, or for that matter, moral realism. I simply wish to draw a contrast between a resolution established through the application of a procedure undertaken within accepted constraints (negotiation), and a resolution established by direct reference to an external benchmark of right and wrong (judgment). Whether that external benchmark is in turn an objective reality or merely an established law or moral norm is of no bearing in this instance.

The bulk of what follows (Part 2) is an examination of those constraints that are required to render a negotiating process fair. My methodological approach, to establishing what are these constraints, is essentially that of a (predominantly narrow) reflective equilibrium, a form of coherentism again closely associated with Rawls (1971, 48-51). In this regard, I appeal to a thoughtful and broad-minded reader's confidently held judgments, as revealed through intuitive responses to practical thought experiments. And I apply, to these responses, arguments based on consistent reasoning in order to draw more generalised conclusions. This involves significant application of what Frances Kamm, one of its advocates, terms the "case-based method of moral analysis" (in Voorhoeve, 2009 [2006], 19). I thus subscribe to "the widespread consensus that the method of reflective equilibrium, broadly understood, is our sole means of proceeding in ethics" (Street 2006, 124).

One of my aims in applying this approach is to avoid making my argument dependent on any single theory of moral reasoning, whether that be consequentialist, deontological or aretaic: Benthamite, Kantian, or Aristotelian. This is important. After all, very few of us are *pure* consequentialists, deontologists, or virtue ethicists. Few among us would feel bound by 'duty' or 'virtue' *whatever* the consequences; and equally few of us think that the *only* thing that matters is a coldly impersonal calculus of consequences. Most of us, especially when it comes to practical ethics, harbour intuitions that reflect a combination of these various approaches. Indeed, recent empirical research into the attitudes of senior business executives suggests just such a pluralistic attitude to ethical issues (Burri, Lup & Pepper 2021, 26-27; Pepper 2022, 42).

My method then is to seek the best rationale possible, consistent with a coherent set of firmly held intuitions. Underpinning this is a belief that, while competing ethical theories like those above may appear notably different to academic scholars, more often than not they *in* *practice* guide us to very similar outcomes. To borrow a metaphor from Derek Parfit (2011a, 411-419), they climb a mountain from different sides but reach the same summit. As Parfit goes on to comment, with regard to what he later calls his "Convergence Argument" (2011b, 244), "it has been widely believed that there are […] deep disagreements between Kantians, Contractualists, and Consequentialists" but in practice that "is not true" (2011a, 419). Thus, while there may be marked differences between these theories in the abstract, there is often a much closer agreement as to how our actions should be guided by them.

This argumentative approach is grounded in a series of practical thought experiments. In the main, these scenarios are intended to be realistic, that is, representative of situations that, if we have not directly experienced them ourselves, are at least close to others that we are familiar with by report. Occasionally though I employ less realistic, more extreme scenarios in order to test arguments at their limit. Armed with these tools, I seek to establish the valid moral constraints on a negotiating process that must be respected for the outcome of a negotiation to be considered fair. Thus, the search is for broad normative principles, consistent with confidently held intuitive judgments.

A key theme in this is consistency. In particular, I insist that business practitioners be held to the same ethical standards that apply elsewhere in society. This whole thesis is closely related to the field of business ethics. And yet it reflects an underlying scepticism, on my part, as to whether there is, or should be, such a field at all. I accept, of course, that ethical issues apply in business. Indeed, as a past practitioner, I feel passionately on the subject. But the concept of a *distinct* business ethics troubles me. The very term *business ethics* seems to imply a separate discipline with its own unique set of rules, at odds in some way with those of other domains of social interaction. When challenged on commonly held ethical principles, practitioners are sometimes inclined to respond along the lines of 'ah, well, this is business… things are done a little differently in commercial transactions… we have our own code'. This is often captured in the deceptively innocuous statement that 'business is business'. This seems trivially vacuous but, as Parfit observes, some tautologies can have substantive meaning. In this case, he notes that the barely hidden implication is that business is "distinctively different from other things, and must be judged on its own terms" and therefore that "ordinary moral standards do not apply" (Parfit 2011a, 71). This strikes me as a dangerous assumption, on the part of practitioners, and one I seek to refute (see, in particular, Chapter 3). I aim to reassert the continuity of morality across different domains of life.

The failure to recognise this continuity and the resultant belief in a different code of acceptable conduct unique to business has, in my view, contributed to the recent spate of financial scandals, such as the collapse of Enron (2001) and Arthur Andersen (2002), the accounting frauds at WorldCom (2002) and AIG (2005), the Madoff Ponzi scheme (2008), the US residential mortgage crisis (2008), the LIBOR fixing case (2012), the Petrobas and Odebrecht bribery cases (2014), the Volkswagen Dieselgate affair (2015), the alleged web of corruption surrounding FIFA (2015), the Wirecard insolvency (2020), the ongoing saga at 1MDB (2015-2022), and the recent alleged misuse of client funds at FTX (2022). Many of those involved in these cases would no doubt argue that their behaviour was only 'what we have always done', 'customary', and 'normal business practice', adding that 'everybody did it' and implying that there existed a 'tacit agreement or understanding about this kind of thing'. But that something is customary or common business practice does not make it right. On the contrary, as Richard DeGeorge (1992, 63) asserts, "business ethics is a subset of general ethics, and those in business are no more allowed to establish their own ethical norms than are people in any other sector of life" (see also Provis 2000a, 146; Provis 2000b, 10; Hsieh 2017, 294-297).

When academics or practitioners appeal to a code of business ethics that diverges from the ethical principles that govern normal social intercourse, the grounds for that divergence need to be explained and justified. To argue only that 'this is business' is clearly inadequate. A disanalogy must be demonstrated. That is, there must be an argument as to *why* we should treat such a business situation differently, and this must be based on aspects of the business situation that differ from the circumstances of normal social intercourse in some morally relevant way. In the absence of such factors, consistency requires that the ethical principles of normal social intercourse prevail. As R.M. Hare (1992, 19) puts it, "ethics is ethics whatever one's vocation."

This emphasis on consistency is especially important for the first of the constraints that I argue are necessary to render a negotiating process fair: that negotiators behave honestly (Chapter 3). This seems a rather sweeping stipulation, but what I intend by it is that negotiators, and in particular business negotiators, recognise they are bound by the same ethical rules that apply in society more broadly. There is no 'special exemption' that permits dishonesty simply because they are involved in a negotiation. The test case for those who argue for such an exemption is that of lying. Arguments for the defensibility of lying during negotiations draw on ideas of tacit consent, self-defence, the 'greater good', fiduciary duty and practicality. I evaluate, and seek to refute, each in turn. I conclude that, in the absence of a credible argument to the contrary, the same moral constraints must apply to lying in business negotiations as apply to lying in other contexts. Lying in a negotiation is therefore wrong and, by extension of the same argument, so is other dishonest behaviour.

The second requirement, I propose, is that negotiators pursue win/win outcomes: that is, that they seek to reach agreements that are mutually beneficial for all parties (Chapter 4). In our very non-ideal world, people at times reach agreements that are not in their own interest, whether due to immature or impaired reasoning, misinformation, gross error, conflict of interest, excessive benevolence, or extreme submissiveness. I argue that, in these circumstances, negotiators are under an obligation to seek an outcome that leaves any such counterparty no worse off than they would be in the event of a failure to agree – that is, than in the 'disagreement position'. I offer three arguments for this side constraint, based on moral imperatives (1) to share the rewards of cooperation (2) to recognise just claims in respect of property and person, and (3) to avoid doing harm to others.

This requirement of mutual benefit, however, does not, of itself, eliminate the possibility of exceptionally imbalanced agreements. There are cases where monopoly control over essential goods or services creates extreme asymmetric negotiating leverage and, as a result, leaves a counterparty vulnerable. If those made vulnerable in this way are subjected to discriminatory pricing, then we have a case of exploitation. A familiar example would be price gouging of those in urgent need of rescue or requiring life-sustaining provisions. I suggest that this concern justifies a third constraint: that negotiators must not exploit such extreme asymmetric negotiating leverage. It is intuitively unacceptable to combine (1) a monopoly, monopsony, or cartel, with (2) essential goods or services, and (3) discriminatory pricing (Chapter 5).

Some qualification is required. In all three of these cases the constraint on the negotiating process must be considered pro tanto. Clearly there may be instances when negative externalities are potentially so damaging that a party is obliged to agree to a negotiated transaction notwithstanding a breach of one of these three constraints. There are also circumstances where any duty to a counterparty is negated, because that counterparty has no right to negotiate at all. 'Negotiating' with a terrorist or hostage-taker would typically fall in this category. The terrorist has no right to threaten violence; and the hostage-taker has no right to that they seek to trade.

Furthermore, there are other occasions when a negotiation is impermissible notwithstanding compliance with these three constraints. These may be cases where the objection lies not with the negotiation per se but with the nature of the underlying goods or services being traded. Examples would include contract murder or the sale of child pornography. Or they may be instances where negotiation is simply inappropriate because of some competing and overriding obligation, such as a duty to unilaterally offer aid to those in dire need (see 6.1.2).

The three moral constraints on the negotiating process are separately derived and might appear unrelated, even ad hoc. I hope to show, though, that in all three cases these constraints move us closer to a perfectly competitive market. They each help, at least to a degree, to rectify market failure and to restore an otherwise compromised consent. In this way, they each help create something closer to an ideal world. They thus demonstrate a significant mutual coherence (see 5.7 & Chapter 7).

Importantly, I argue that these moral constraints are not only necessary, but are also *sufficient*, for a fair negotiating process (Part 3/Chapter 6). While some will advocate further constraints based on the broader demands of distributive justice, this is to fail to recognise two important distinctions. One involves a division of moral labour between the responsibility of a particular negotiator and the collective obligations of society as a whole. A single negotiation cannot be expected to right all the world's wrongs. And an individual negotiator cannot be held responsible for the collective injustice of those background conditions. Nor can they alone overcome what is a significant collective action problem (see 6.1.1). The other distinction contrasts the role of the individual *qua negotiator* with that of the individual *qua benefactor*. There are exceptional circumstances where the need of another party is so dire, so extreme and urgent, that a duty to aid falls on the individual negotiator. But when a negotiator responds to such a call for beneficence, they cease to negotiate in any full and proper sense of the word. The duty to aid takes precedence (see 6.1.2). This duty to aid is driven by need and is quite distinct from any broader considerations of distributive justice based on inequality and injustice. To attempt to impose any broader requirement for distributive justice on an individual negotiator is to commit the fallacy described at the outset, that is, to confuse the nature of a negotiation with that of a judgment.

That these moral side constraints underdetermine the outcome of a negotiation, furnishing a range of possible fair solutions rather than a single solution, is, I suggest, not a weakness but, on the contrary, a strength of the approach. It correctly reflects the way we typically look at a negotiation, recognising that the ultimate agreed outcome will, to a significant degree, reflect other non-strictly ethical considerations such as legitimate bargaining power and negotiating skill (see 6.2).

In a concluding chapter (Chapter 7), I conduct a brief review of all the substantive claims made and point again to the practical implications. If there is a thinly veiled purpose in all that follows it is to challenge “the widespread assumptions that ethical = impossible and unethical = inevitable in negotiation” (Gunia 2019, 5). Such assumptions are of course largely empirical in nature, but their plausibility or implausibility is closely linked to the normative question of what makes a negotiation ethical in the first place.

PART ONE: NEGOTIATION AS A PROCEDURE

Chapter 2. Fair Process, Fair Outcome: Negotiation as Pure Procedural Justice

ABSTRACT: Game theory's goal of a fair solution to the bargaining problem (i.e., a formula that delivers a uniquely fair outcome to any negotiation) remains elusive. Indeed, it is doubtful that any of game theory's 'solutions' is 'fair' in an ethical sense of the term. I argue that this search for a fair solution is misdirected. Rather than a fair solution, we would be better to think in terms of a fair process. Instead of seeking the appropriate axioms/constraints on the *solution*, we would then seek the appropriate axioms/constraints on the negotiating *process*. On this approach, the fairness of the negotiated outcome is wholly determined by the fairness of the process through which that outcome is reached; and the fairness of the process is itself wholly determined by the adequacy of the constraints applied to that process. This makes fair negotiation an instance of pure procedural justice.

Many philosophers have looked to game theory to provide an answer to what John Nash (1950, 155) termed "The Bargaining Problem." The focus of successive efforts has been the search for a 'solution', a formula that delivers a uniquely fair outcome to any negotiation. I suggest an alternative approach: that we seek fairness not in a single *solution*, but rather in the negotiating *process* itself, relying on that process to deliver a fair result from a range of potentially fair outcomes. On this account, then, the fairness of the chosen outcome is wholly determined by the fairness of the process by which that agreed outcome is reached; and the fairness of the process is itself wholly determined by the adequacy of the constraints applied to that process. This makes fair negotiation an instance of pure procedural justice.

2.1 A FAIR SOLUTION OR …

From its earliest post-war foundations game theory has sought to find a fair solution to the problem of cooperative bargaining. Starting with the simplest model of a two-person cooperative negotiation, Nash (1950) declared his purpose to find "a definite 'solution'" (155) to the problem of a "fair bargain" (158), while Richard Braithwaite (1955) similarly concluded his own equally prescriptive analysis by insisting that "no other behaviour would seem to me to be *sensible, prudent and fair*" (52). Both advocated a formulaic approach and a single solution to establishing a fair bargain and this has been emulated by most game theorists since, including John Harsanyi (1956), Ehud Kalai and Meir Smorodinsky (1975) among others (with the exception perhaps of Ariel Rubinstein (Heap & Varoufakis 1995, s4.4.1)).

Critique of these 'fair solutions' has focussed on their fairness, or *lack of it*. While such solutions may be seen as rational and prudent, many challenge whether they are fair in any ethical sense of the word. Ken Binmore (1994, 96) argues that the basic propositions of game theory are "necessarily neutral with regard to moral issues." Nash and Braithwaite's solutions, in particular, seem to reflect what Shaun Heap and Yanis Varoufakis (1995, s4.3.3) label "*the equilibrium of fear*."

The classic illustration of this is Braithwaite's (1955) thought experiment concerning a negotiation between Matthew, a trumpet player, and Luke, a pianist. Both wish to play their respective instruments in their shared home and must agree which of them plays when. Braithwaite's proposed solution is that Luke and Matthew play at different times in the ratio of 17 evenings for Luke to 26 evenings for Matthew (36). The justification for this weighting in favour of Matthew is that he can tolerate the cacophony of his trumpet and Luke's piano being played together, while Luke cannot (37). While both prefer to play alone, practicing together is worse for Luke than it is for Matthew. Matthew thus has the threat advantage – he can credibly assert that, if no agreement is reached, he will practice whenever he pleases. Hence, the equilibrium of fear is established in his favour. The key point here is that, while Braithwaite may have established a viable equilibrium, it is doubtful that this equilibrium bears any particular relationship to fairness. As Rawls (1971, 141) noted, "to each according to his threat advantage" is not a compelling conception of fairness (see also 134-135 n10). It seems, when this equilibrium is reached, that we can still reasonably ask whether it is fair. The question remains open, suggesting that the solution lacks real moral content (see also Sen 1970, 121-123).

Attempts have been made to address this apparent ethical deficit by proposing more obviously normative solutions. For example, David Gauthier (1986, 14, 137) argues that a fair bargain can be reached only when we equalise the proportional concession required of each party from that party's best possible outcome, expressed in terms of utility; while resourcist solutions, like that of Brian Barry (1989, 85-95), set considerations of utility aside in favour of an equitable division of the resources that are the subject of negotiation. The problems with these theories tend to mirror each other. A narrow focus on a strict division of resources can be unfair on those whose disability, special needs or lack of talents means that they require greater resources in order to enjoy a flourishing life, that is, an acceptable level of utility. Such a concern lies at the root of what is often termed the 'capability' approach (Sen 1982, 366-368; Nussbaum 2011, 57-58; Robeyns & Byskov 2020 [2011], s2.5). On the other hand, Gauthier's formulaic approach to the allocation of utility conversely can create a highly counter-intuitive division of resources, in those cases where such resources are valued very differently by the parties (for a fuller discussion of Gauthier's theory and related objections see 4.5.2).

Like their predecessors, these proposals share one thing in common: they prescribe a fair *solution* or *outcome*. In approaching a negotiation, they set out with the persona of an arbitrator who must make a judgment, revealing to and imposing upon the parties a pre-determined, formulaic solution. But it seems to me that such a judgment is a substitute for a negotiation, not a negotiation itself. It is to start at the end rather than the beginning. Indeed, their methodology is what has been coined elsewhere an "*end-state*" approach to establishing fairness or justice (Nozick 1974, 155) (discussed further in 2.2 below). They are focussed on the end position. And they presuppose that there is a discernibly fair end position to aim at. That was the very essence of Nash's achievement: from his four axioms he derived a solution, and it was a *single solution*.

I propose an alternative approach: that we look elsewhere for fairness in a negotiation; and, in particular, that we seek to locate that fairness in the negotiation process itself. The key question is then whether the *process* is fair - and what constraints might be placed on it that would render it fair. This is consistent with the view of negotiation as a social construct: a process designed to deliver cooperative behaviour based on mutual consent.

But let me go back first to the very beginnings of game theory. In their ground-breaking *Theory of Games and Economic Behavior*, John von Neumann and Oskar Morgenstern (1944) proposed just two formal requirements of their bargaining model: Pareto optimality and individual rationality, which together they believed defined the parameters of the cooperative solution. They argued that no further formal constraints could be justified. Instead, any subsequent narrowing of the field of potential solutions would need to reflect "standards of behavior" (40-42) (i.e., accepted conventions imposed by society), together with the bargaining abilities of the parties involved. The result, of course, was that there was no single solution that could be reached in a formal mathematical way. Instead, there was a range of possible outcomes, a "multiplicity of solutions" (266; see also Luce & Raiffa 1957, 205). This was viewed at the time as a problem to be solved (Luce & Raiffa 1957, 118-119). But I want to suggest that it should not be viewed as a problem – at least not when approached from an ethical perspective.

One driver of the game theorist's desire for a single solution may be an underlying consequentialist ethos: an objective of maximising something. The something may vary between advocates, but it is usually *one* thing. Some consequentialists have argued that it is possible to 'consequentialise' everything. In other words, that all moral theories can be reduced to consequentialism simply by defining the 'good' appropriately. However, Campbell Brown (2011, 749, 760-763) points to a range of theories where this is *not* possible, because they stipulate either that nothing be maximised or that multiple and irreducible things be maximised. In the latter category are agent-relative theories, i.e., theories that say that many irreducible things should be maximised, a different one for each of us.

Such agent relativity seems exactly the kind of ethical structure that we would expect to find in a negotiation. On this approach, within an agreed process and bound by certain rules of fair play, but not ones that necessarily stipulate a single outcome, negotiating parties are permitted to pursue their own ends. They do this from an unashamedly agent-relative perspective. That is, they seek to achieve their *own* objectives, albeit within a process that includes an acknowledged set of ethical constraints. Clearly, under such conditions, there is unlikely to be a single, uniquely fair solution. Rather, there will be, in all likelihood, a *range* of potentially fair outcomes. The ultimate outcome agreed will be a function in part of bargaining strength and negotiating skill. On this view of negotiation, the objective fairness of the agreed *solution* derives from the fairness of the negotiating *process* itself. The final outcome of the negotiation is clearly a specific outcome but remains only one among a number of fair outcomes that *could* have emerged, had the process unfolded differently.

This seems to me to reflect much more closely how we actually approach and police negotiations. We do not anticipate an immediate solution or necessarily believe that a single fair solution exists at all – or, at least, that any such solution pre-exists the process of negotiation. Instead, we agree, tacitly or explicitly, on a set of 'ground rules' that help ensure fairness – von Neumann and Morgenstern's "standards of behavior" – and under which we explore what matters most to each of the parties.

This portrayal of a negotiation better captures the direction of information discovery. Game theory models assume that we have all the required personal preferences (i.e., the utility functions) established at the outset. From these, we determine the one fair answer. But we know this is wholly unrealistic. Although the various parties' preferences may, to a large degree, drive the ultimate outcome of the negotiation, we rarely have full transparency over these preferences from the outset. It is largely through the process of the negotiation itself that the different utility functions are revealed and/or constructed. Indeed, this is one of the purposes of negotiation: to establish the utility that different parties attach to different elements of the transaction in order to reach an efficient Pareto-optimal outcome.

Importantly this process of establishing the utility functions is not limited to each party ascertaining the preferences of the other party. It is often only through the process of negotiation itself that each party comes to fully understand their *own* preferences. As Jason Alexander (2007, 17) puts it, "we might not have well-defined preferences until we are actually presented with a choice situation in which we have to choose among them." In the film *Sophie's Choice*, Meryl Streep's character is forced by Nazi guards to give up one of her children. Repeatedly she screams "*Ich kann nicht wählen!*" She cannot choose. The guards attempt to take her son. Only then, now *forced* to choose, can she decide which child she values the most. She gives up her daughter (Pakula 1982 based on Styron 2004 [1979], 594). Similar, if less harrowing, choices have had to be made in recent negotiations over the UK's withdrawal from the European Union. A border of sorts in the Irish Sea, which was once considered unthinkable, was – under huge pressure in negotiation – eventually deemed less unthinkable than a failure to reach agreement at all. Subsequently it appears that a need to create negotiating leverage has led to a further reappraisal of that trade-off, prompting a proposed breaking of international law (something which many would again have previously ruled unthinkable). The point is that a full understanding of these preferences only emerged in the process of negotiation – indeed as a result of the process of negotiation.

In such cases, Robert Sugden (2018, 18) argues that preferences are best seen not as revealed, but rather as "constructed" in the course of negotiation. As he writes, "individuals often come to decision problems without well-defined preferences that pre-exist the particular problem they face; instead, whatever preferences they need to deal with a problem are constructed in the course of thinking about it." This suggests that the negotiating process itself plays more than a merely epistemic role in determining preferences. It not only helps reveal those preferences, but also contributes to their very construction or constitution.

The relevant question is thus subtly changed: instead of seeking the axioms/constraints on the solution, we need to seek the axioms/constraints on the process for reaching a solution. Clearly, I have in mind here an alternative approach to establishing a fair negotiation, one that draws its ethical legitimacy from the constraints imposed on the negotiating process rather than through a direct consideration of the properties a fair solution ought to exhibit.

2.2 … A FAIR PROCESS

In closing their discussion of bargaining games, Heap and Varoufakis (1995) allude to this same possible way forward. They note that, "one of the difficulties in moral philosophy is that our moral intuitions attach both to the *patterns*, or attributes, *of outcomes* (our *ends*) and the *processes* (or *means*) which generate them" and go on to suggest that "it would be extremely helpful if we could somehow compare these otherwise contrasting intuitions by, for instance, seeing how constraints on means feed through to affect the range of possible outcomes." Under this view, "the axioms then become a way of placing constraints upon these procedures which we select because we find them morally appealing" (Ch4.5.2). Gregory Kersten, Wojtek Michalowski, Stan Matwin and Stan Szpakowicz (1988, 226) see this as particularly relevant when the issue to be resolved is "fuzzy and ill-defined," as is the case in negotiations of any real complexity. In such circumstances, they believe it is better to "not direct the model of the negotiating process towards a solution," but rather to focus on the "domain *rules*" that underpin and constrain the process.

In ethical and political philosophy this emphasis on process has found its most potent expression in two philosophers, whose work is more commonly juxtaposed: John Rawls and Robert Nozick. It is useful to rehearse their lines of thought. In *A Theory of Justice* (1999 [1971]) Rawls grounds his justice as fairness on the principle of fair process, although he prefers the term procedure: "The idea of the original position is to set up a fair procedure so that principles agreed to will be just. The aim is to use the notion of pure procedural justice as a basis of theory" (118). Rawls contrasts this pure procedural justice with perfect (or imperfect) procedural justice. In the latter case, a just process merely *reveals* a just outcome, but the fairness of that outcome stands independently of the process. For Rawls, "by contrast, pure procedural justice obtains when there is no independent criterion for the right result: instead there is a correct or fair procedure such that the outcome is likewise correct or fair, whatever it is, provided that the procedure has been properly followed" (75). Thus, in the case of pure procedural justice, the fairness of the process wholly *constitutes* the fairness of the outcome. Or to express the same in reverse, the fairness of the outcome is fully *derived* from the fairness of the process.

I submit that this reflects closely the nature of a negotiation. If all parties accept that a negotiation has been conducted fairly, then it seems strange to ask whether the outcome is fair. If the parties willingly enter a negotiation and all the elements of that negotiation are fair, then the outcome simply *is* fair. In the context of a negotiation, a fair outcome is what is agreed willingly by the parties to a fair process. A party might agree to a negotiated outcome and, at the same time, complain of the unfairness of that outcome. But, to do so consistently, the party would have to point to some element of the process that was unfair (for example, dishonesty, coercion, or the abuse of monopoly power by the counterparty). For the process to be deemed *fair* and yet the outcome to be deemed *unfair*, there would have to be some independent benchmark by which to make that judgment – perhaps a judicial ruling, a binding arbitration, or some form of self-standing moral principle. But this would represent a different procedure, one where the outcome is imposed by a third party or predetermined, rather than resulting from a negotiation between multiple parties, each of whom is free to walk away.

A key consideration here is this sense of voluntariness. It is in the nature of a negotiation that party and counterparty are not *forced* to agree to anything. As argued earlier (Chapter 1), a negotiated agreement implies at least some degree of consent. Each party has the right to decline. They are free to walk away, retaining the goods they thought to trade. This is what Benjamin Ferguson (2016b, 960) terms "The Optionality of Ordinary Transactions," pointing out that, just because a transaction is beneficial to all, or in some other manner deemed fair, this does not mean any party is obliged to consent to it.

In such circumstances, it is very hard to envisage what an independent benchmark of fairness would entail. An egalitarian might say 'split it down the middle' or 'give more of the surplus to the least well off', but if the other party declines to do so, has every right to so decline, and chooses instead to walk away, then such a judgment fails to be compelling. At least, it fails to be compelling in the context of a negotiation, because the 'imperative' that it offers is a transaction that one party is not prepared to agree to and, by the very nature of a negotiation, has *every right not to agree to*. The egalitarian is left only with a possible appeal to beneficence but, in the context of a negotiation, this can be no more than supererogatory (for a fuller discussion of the obligation of beneficence see 6.1.2). Neither party is obliged to respond to this appeal. The social construct in which they are engaged requires their consent, thus granting them each an essential right to walk away. And, as we have seen earlier, once that consent/right is abrogated, then we have moved from a negotiation to something quite different, a judgment imposed from outside. To reiterate, we are seeking here to establish what are the ethical principles governing a particular social construct, negotiation, and that social construct is rooted in this concept of consent (see Chapter 1). It is possible of course to appeal to other values, but without consent there is nothing left that can sensibly be characterised as a negotiation. In sum, then, the fairness of a negotiated outcome can reflect no more than the willingness of the parties to agree, that is, to consent, together with the fairness of the process by which the parties reach that agreement. There is nothing else that has any 'force' – that is, any of the force of a moral obligation on the parties *as negotiators*. It is always possible that one negotiating party is in such dire need, that this triggers a duty to aid, that is an obligation of beneficence. But this involves an act of charity, a gift not a trade (see 6.1.2).

It might be objected that, if new information were to arise that was known to neither party at the time of a negotiation, then such information might render what appeared to be a fair negotiation, unfair. But this is the case only in hindsight and can be true of any agreement, however reached, and indeed of any state of affairs, whether the prior subject of negotiation or not. Such a development might perhaps create a moral obligation (even if only an obligation of beneficence) to subsequently change the outcome, but it would not render unfair the *negotiation* that created that outcome. By way of illustration, let us imagine that, following a lengthy negotiation, I commit to sell you a house at an agreed price, but a month subsequent to that house prices in the area unexpectedly fall by 25%. In such circumstances, that new information, on the prospects for the housing market, might suggest that the final outcome is in some sense imbalanced, inequitable, or unfair but it would not render the prior negotiation itself unfair. The fairness of a negotiation can only be judged in the light of the information available at the time of the negotiation.

All this is not to say that there can be no other basis for fairness than the outcome of a fair negotiation. Negotiation is a particular social construct. But, as suggested above, there may be other processes available for establishing agreement. For example, the parties might choose to submit themselves to binding arbitration, accepting that they will be bound by whatever result that alternative procedure delivers.

Again, it is instructive to examine Rawls' position. In *Political Liberalism* (1993) Rawls describes his political philosophy as 'constructivist', emphasising that a "political constructivist regards a judgment as correct because it issues from the reasonable and rational procedure of construction when correctly formulated and correctly followed" (96). Thus, consistent with pure procedural justice, the fairness of his two principles is constructed upon the fairness of the original position with its veil of ignorance, that is, upon the fairness of the procedure from which those principles are derived (103). But Rawls does not claim this is the sole source of fairness. He does not seek to expand such a constructivist approach to encompass morality as a whole: "It is not proposed as an account of moral values generally" (125). Rather Rawls emphasises the "limited political purposes" (110) of justice as fairness.

Again, I argue for a parallel here in the domain of negotiation. The fairness of a negotiated outcome is a function of the fairness of the negotiating process by which it is reached and of that alone. But there are other procedures that, if adopted, might produce different outcomes that could also be considered fair. A negotiation, or bargain, is not the sole decision procedure available to those seeking "fair terms of cooperation" (4). Like Rawls, my claim is limited: negotiation is a social construct, designed to deliver a fair outcome through the application of pure procedural justice; but there are other procedures, like judgment and binding arbitration, which produce different outcomes that might also be considered fair. However, as argued earlier, these alternative procedures are not a negotiation. The parties do not retain the right to decline and walk away. They are no longer voluntary participants in a process aimed at reaching an agreement, to which they must consent. Instead, they are the subjects of an imperative imposed by a third party. The ethical rules governing negotiation are then a *part* of morality, not morality in its entirety. They govern the obligations for participants in the particular social construct we characterise as negotiation. Different principles may apply beyond that social construct, for example in determining society's collective responsibility for ensuring distributive justice (see discussion in 6.1.1).

Let me offer one further clarification. In a negotiation, a fair process ensures a fair outcome to the *negotiation*. That does not entail that the entire state of affairs that obtains following the negotiation is fair. It cannot be expected that a single commercial negotiation will right all the world's wrongs. As Alan Wertheimer (1996, 234) puts it, we can condemn the unfairness of "background conditions," while still acknowledging that, "it is unreasonable to expect the better-off party to repair those background conditions by adjusting the terms of a particular transaction." If you are an undeserving billionaire and I am, by contrast and through no fault of my own, among the less well-off, I cannot expect that differential to be significantly rectified by you selling me a car, however fair that latter transaction might be. Almost whatever the terms of the car sale, the underlying inequality, and thus perceived unfairness, will remain. The fairness of the process ensures the fairness of the outcome of the process itself, not of the welfare of all those involved in the negotiation, howsoever their lives should stand. The individual negotiator is morally responsible for neither the opening nor the closing position of their counterparty. Instead, the extent of any responsibility they bear is limited to the difference between the two, that is, the change resulting from the negotiation.

It may be argued that the state of the world as a whole determines the justice or injustice of the marketplace as an institutional arrangement (Hausman 1989, 333). Perhaps there is something fundamentally unjust about the structure of capitalism, or of neo-colonialism, or of the relationship between the northern and southern hemispheres. But the state of the world as a whole cannot determine the fairness of an individual negotiation. That would be to place an unreasonable burden on those seeking to negotiate. There is here a division of moral labour, creating a distinction between the obligations of the individual negotiator and the collective obligation of society as a whole (I discuss this further in 6.1.1).

This procedural foundation to Rawls' thinking is closely mirrored in Robert Nozick's *Anarchy, State, and Utopia* (1974), albeit with a starkly different structure built upon it. Under Nozick's entitlement theory, justice is determined not by the character of the outcome itself, whether that be based on equality, merit, need or any other pre-ordained, "patterned principle" (156), but rather by the nature of the process by which such an outcome is reached. An allocation of property, for example, can be judged fair or unfair only in accordance with the history of how it came into being, that is, the process by which the property was originally appropriated and the legitimacy of the transfers that have since occurred (153). If, as in Nozick's famous thought experiment, Wilt Chamberlain negotiates from willing fans a fee of 25 cents per game in return for watching him play basketball and acquires huge wealth as a result, then that is perfectly fair as long as the arrangement is voluntary on all sides (161-164). As Nozick puts it, "whatever arises from a just situation by just steps is itself just" (151).

The opposite, the "unhistorical principles of distributive justice," are what Nozick calls "*end-result principles* or *end-state principles*" (155). Given that, as Michael Walzer (1983, xi) points out, "money equally distributed at twelve noon of a Sunday will have been unequally redistributed before the week is out," the strict enforcement of these end-state principles would entail the constant (or at least periodic) reversal of any changes (i.e., any divergence from the previously established 'pattern') brought about through voluntary transfers between consenting adults. In Nozick's mind, to frustrate such voluntary transfers in this way is to infringe on personal autonomy and threaten individual liberty (160-164). The implications of this thesis are wide-ranging and, unsurprisingly in the views of many, this is "the most important argument" in Nozick's book (both Ryan 1981 [1977], 326; and Wolff 1991, 80).

Again, there is a clear parallel here with the character of a negotiation, which after all is very close in nature to a voluntary transfer. Indeed, they are frequently one and the same. Just as the fairness of Nozick's ownership holdings stems from the process through which they are derived, so too, I argue, the fairness of a negotiated settlement draws its legitimacy from the process by which it is reached.

To draw such a parallel is of course to hitch one's wagon to some highly controversial horses. Nozick's libertarianism is widely perceived to support a comparatively extreme positioning on the right of the political spectrum. To be clear, I do not subscribe to the political implications that Nozick draws from this structure. However, it is worth noting that even G.A.

Cohen (1995, 41), a philosopher of the left and a fierce critic of Nozick, at least to some degree accepts the structure itself: "I do not endorse a *purely* end-state theory of justice. It is evident that identical distributions of holdings can differ with respect to justice because they differ with respect to their histories."

Moreover, I would contend that the reason Nozick's entitlement theory is objectionable to so many people is not, in the main, due to its un-patterned or 'historical' character. After all, most people seem happy that substantially greater financial rewards go to those at the top of their profession, including elite sportsmen like Wilt Chamberlain. The concern is not around the freedom to transfer – or the freedom to negotiate – per se, but rather around the constraints (or *lack of them*) placed on that freedom.

The key, then, to the legitimacy of the process – whether we call it a transfer or a negotiation – would appear to be the adequacy of the constraints placed upon it. In Nozick's entitlement theory the only constraint is that the consent to a transfer be voluntary. It cannot be the result of force. There are differing views as to whether Nozick's concept of consent itself implies an absence of fraud and therefore a further constraint in that regard. Jonathan Wolff (1991, 85) argues that is does. However, Ferguson (2018, 179) questions whether a strictly libertarian position has the resources to condemn such deceit. We need not attempt to resolve this debate here (but see related discussion on dishonesty and lying in Chapter 3). The point is that, even if we allow for a prohibition on fraud, in addition to the requirement for consent, this remains a *very* modest set of constraints.

Furthermore, Nozick permits no place for a division of moral labour, by which the state takes responsibility for promoting distributive justice outside the sphere of individual negotiations. Nozick is committed to this position because he believes that any further constraints on free exchange, and/or any expansion of the role of his "minimal state" (Nozick 1974, 26-28), curtail autonomous decision-making and therefore breach individual rights. This in turn stems from Nozick's commitment to liberty based on self-ownership – the right to non-interference with one's self and, by extension, property – in precedence to all other moral values (Wolff 1991, 7-8).

To echo Bernard Williams (1981 [1975], 31), this is little more than "the anarchists' bag of moral ideas." It is clearly possible to entertain a broader set of values and, judging by the critical reaction to the implications of Nozick's libertarianism, most people do harbour such further values: values such as those entailed by mutual aid and protection against exploitation. As Thomas Nagel's (1981 [1975]) critique of Nozick puts it, "the sources of morality are not simple but multiple" (196).

My suggestion is that, when we consider what makes a negotiation fair, we adopt Nozick's structure – his emphasis on fair process – but not his values and therefore not his "minimal morality" (Williams 1981 [1975], 32), his minimal list of constraints. As I have argued above, negotiation is a process, and the fairness of its outcome is properly judged by the fairness of the process that delivers that outcome. Moreover, that fairness is itself wholly determined by the adequacy of the constraints applied to that process. But, as Wolff (1991), Williams (1981 [1975]), and Nagel (1981 [1975]), have all convincingly argued, Nozick's constraints of voluntary consent and (perhaps) absence of fraud are woefully inadequate. Indeed, they are arguably inadequate even within the confines of a strictly libertarian framework (Otsuka 2003, 3, 29-35). How might we expand them? This will be the focus of what follows. In Part 2, I will make the case for three further constraints (Chapters 3, 4 & 5) that, I claim, are necessary to render a negotiating process fair. Furthermore, in Part 3 (Chapter 6), I will argue, contrary to Nozick, for a division of moral labour, whereby society as a whole, in the form of the state, takes a collective responsibility for ensuring distributive justice.

Before leaving this discussion though, I would like to end on a practical note: this emphasis on a fair process reflects much more closely what happens in the real world of negotiation. While game theorists, with few exceptions, are focussed on the search for a fair solution or outcome, the majority of practitioners, in my experience, place much more importance on a fair process. Indeed, many practitioners are sceptical as to whether, in most negotiations, there really is such a thing as *a* fair solution; or at least, if they do believe in such a fair solution, they see it as derivative. It is the outcome of a fair *process*. A fair process is what they do feel some moral obligation to ensure. In their minds – and specifically in the context of a negotiation – if the process is fair, then the outcome is fair.

These observations are based on my own business experience, but they draw some support from independent empirical data. An example of this comes from the game theorists themselves. Ken Binmore, Joe Swierzbinski, Steven Hsu, and Chris Proulx (1993, 396 incl. Fig.5) report on a laboratory experiment involving 252 undergraduates in a simulated negotiation. They note a "highly significant relationship" (R^2=0.73) between participants' views as to what constitutes a fair outcome and the actual outcome of a *process* that those participants have recently taken part in and implicitly consented to. It is not entirely clear how to interpret this, but one explanation would be that the participants are assuming that the actual outcome of an apparently *fair process* must be a *fair outcome* – in short, that a fair process implies a fair outcome.

Further data comes from the legal domain in the form of "procedural justice research" (Hollander-Blumoff 2017, 21). These empirical studies show perceptions of legitimacy to be closely related to assessments of procedural fairness (Hollander-Blumoff 2010, 385). Moreover, they find that, in determining such legitimacy, "perhaps surprisingly, perceptions of distributive justice generally have a much more modest impact than perceptions of procedural justice" (Welsh 2004, 762-763). And importantly, this perception of procedural fairness – of a fair *process* – is highly correlated with a willingness to accept a negotiated *outcome* as fair (Hollander-Blumoff & Tyler 2008, 484, 491-492; Hollander-Blumoff 2010, 384, 417). In short, "fairness of process matters to people" and drives their assessment of what constitutes a fair outcome (Hollander-Blumoff 2010, 387; see also Welsh 2004, 759-762). The original work in this area focussed on third-party dispute resolution, but more recently there have been similar findings in models of bilateral negotiation (Welsh 2004, 764-765; Hollander-Blumoff & Tyler 2008, 478, 490-492; Hollander-Blumoff 2010, 384, 413-415). And it seems that this focus on fair process is something we learn at a very young age: a recent study reveals that "children have a robust preference for procedural justice, even when a procedurally just process leads to distributive inequality" (Dunham, Durkin & Tyler 2018, 7). In other words, children prefer a fair process even when it delivers an outcome that would, on some criteria at least, be deemed inequitable.

It may be objected that what practitioners, or for that matter children, think is of no normative relevance. The fact that someone believes that X is right or wrong (an empirical observation) does not, of itself, prove that X is right or wrong (a normative conclusion). An 'is' cannot deliver an 'ought'. However, we are seeking here to analyse a social construct – the practice of negotiation – and we need to be comfortable that our normative conception of that social construct is not too distant, too divorced, from common usage and perception. Moreover, as I made clear in Chapter 1, my approach draws on Rawls' methodology of a reflective equilibrium. A key requirement of that reflective equilibrium is a close congruence with our firmly held intuitive judgments.

Such an appeal to intuition, of course, lies at the root of most, if not nearly all, thought experiments. With that in mind, let me illustrate the argument once more, with the aid of such a thought experiment. If, during the negotiation of a corporate merger or acquisition, you suggested to an investment banker that she should be seeking a *fair solution*, I think she might look a little mystified. She probably would not see that as her role, her 'job'. She is after all more advocate than judge, striving to achieve the best outcome she can for her client, not for the other side. But if you demanded of her that she play her part in ensuring a *fair process* – if you insisted that there are certain rules of conduct in the process of a negotiation that she is duty-bound to adhere to – then I think she would understand your point entirely. This is not a function of the principal-agent relationship. Most principals, I suspect, would feel the same. The commonly perceived obligation in a negotiation is to behave properly, to 'play fair'. It is an obligation that relates to a process not an end point.

What might this investment banker understand then by a fair process, by 'playing fair'? Well, I guess that, put on the spot, she might at first be a little vague but, given a moment for reflection, I can imagine her coming up with something along the following lines: 'Be honest, that is, be straight; try to find a deal that's good for everybody; and, if you do end up in an exceptionally strong negotiating position, don't abuse that. Remember that what goes around, comes around'. And what would she not say? Well, she certainly would not see it as her job to 'maximise the product of the negotiating parties' utilities' or to 'equalise the concession of each party from their best possible result' or to 'use the transaction to rebalance the resulting endowments of the various parties' or, even more simply, to 'divide things down the middle.' In other words, when it came to assessing *fairness*, her focus would be on a fair process rather than directly on a fair outcome.

But let us further imagine that the investment banker above succeeds in her transaction. After a lengthy negotiation, involving full disclosure, an absence of any kind of coercion and a fair process in every respect, the two parties agree terms. An independent observer now declares the proposed transaction unfair. On what basis could that be? If the terms are now substantially changed, one party will likely walk away, as they have every right to do. As we have seen, such is the essential nature of a negotiation. So, what does this new independent declaration of fairness mean? I have suggested that, in the context of a negotiation, it lacks any real force. I reiterate that the only truly meaningful measure of a fair outcome to a negotiation is whether the parties reach agreement, that is mutually consent, following a fair process.

2.3 CONCLUSION

Game theory's goal of a fair solution to the bargaining problem remains elusive. I have suggested that this search for a 'solution' is misdirected. Rather than a fair solution, we would be better to think in terms of a fair process. Instead of seeking the appropriate axioms/constraints on the solution, we would then seek the appropriate axioms/constraints on the negotiating process. On this approach, the fairness of the negotiated outcome is wholly determined by the fairness of the process through which that outcome is reached; and the fairness of the process is itself wholly determined by the adequacy of the constraints applied to that process – an instance of pure procedural justice. In Part 2, I will explore exactly what are those constraints that attach to a fair process.

PART TWO: NECESSARY CONSTRAINTS

Chapter 3. A Lie is a Lie: The Ethics of Lying in Business Negotiations

ABSTRACT: I suggest that a necessary constraint on a negotiation, for it to be regarded as fair, is a requirement of basic honesty. In particular, I argue that lying in business negotiations is unacceptable. It is pro tanto wrong and no less wrong than lying in other contexts. First, I assert that lying in general is pro tanto wrong. Then, I examine and refute five arguments to the effect that lying in a business context is less wrong than lying in other contexts. The common thought behind these arguments – based on consent, self-defence, the 'greater good', fiduciary duty, and practicality – is that the particular circumstances which are characteristic of business negotiations are such that the wrongness of lying is either mitigated or eliminated completely. I argue that all these 'special exemption' arguments fail. I conclude that, in the absence of a credible argument to the contrary, the same moral constraints must apply to lying in business negotiations as apply to lying in other contexts. Furthermore, I show that for the negotiator there are real practical benefits from not lying. Finally, I suggest that the same principles that apply here to lying, apply also to other forms of basic dishonesty, like fraud and corruption. Accordingly, the first requirement that must be satisfied to ensure a fair negotiating process and resulting outcome is that the negotiating parties observe the basic demands of honest behaviour that are expected of society as a whole.

Business ethics has become an essential component of any business school curriculum or similar management training. I suggested above (see Chapter 1), that, while this is generally to be welcomed, it carries with it a risk: that we take its very subject title – *business* ethics – to imply that there is a distinct set of moral rules that apply only to business and that are significantly different from the moral rules that apply in society more generally. The usual thought is that there is some kind of exemption, whereby normal ethical standards need not be so strictly adhered to in a business context. I believe this suggestion should be firmly resisted.

Perhaps the most notable case of this has been the half-century-long debate, ignited by Albert Carr (1968), on the permissibility of lying in business negotiations. Although what follows applies, I believe, equally to other forms of basic dishonesty, my focus will be on lying as the archetypal example of such behaviour. My central claim is that lying in business negotiations is pro tanto wrong and no less wrong than lying in other social contexts. I argue, furthermore, that such lying is both unnecessary and, in most cases, ultimately to the disadvantage of the negotiator.

In 3.1, I assert that lying in general is pro tanto wrong, that is, morally impermissible unless there are exceptional countervailing moral factors that carry even greater weight. My insistence on the *consistent* application of this normative principle is the foundation of the argument that follows.

In 3.2, I lay out what I call the special exemption thesis: variously the explicit claim or implicit assumption that lying in business negotiations is not wrong, or at least not *as* wrong as lying in other contexts – that is, that business negotiations are in some way 'exempt' from the moral rules that govern lying in general. The empirical evidence suggests that this view is widely held. But those who defend this view must explain what special circumstances justify this exemption. Where is the moral disanalogy that makes permissible in this 'special' instance what is elsewhere impermissible?

In 3.3, I examine the five main arguments used by apologists to defend the special exemption thesis: the argument from consent, the argument from self-defence, the argument from the 'greater good', the argument from fiduciary duty, and the argument from practicality. The common thought behind these arguments is that the particular circumstances which are characteristic of business negotiations are such that the wrongness of lying is either mitigated or eliminated completely. I seek to demonstrate that each of these arguments is seriously flawed.

Finally, in 3.4, I reiterate that the burden of proof rests with the apologist. The apologist must show that there is a disanalogy, a moral difference, between lying in business negotiations and lying in ordinary social interaction. Because the apologist has failed to do this, *consistency* requires that the two be treated alike. I conclude that it is time for business practitioners (1) to acknowledge that lying in negotiations is pro tanto wrong and no less wrong than lying in other contexts; and (2) to recognise that such lying is both unnecessary and, in most cases, ultimately to their own disadvantage.

Let us consider the following scenario:

Zach & Abe

Zach was at a conference in Philadelphia to promote his firm's latest pharmaceutical packaging machine. As he left the conference hall, he was

> approached by a regular customer, Abe: "Zach, that's an impressive machine, you've developed. We could be a buyer. But we would be looking for a discount. Only fair to tell you." Zach was fully prepared to offer a sizeable discount. His firm had received no customer enquiries at all to date and his boss had made her feelings very clear: they must place one of these prototype machines quickly to test it in action and, if that meant all but giving the thing away, then so be it. But Zach also knew there would be huge plaudits for him personally if he could keep that discount to an absolute minimum. "Abe, you know I can't do that. This is the most advanced machine of its generation. There is considerable interest from other customers; and my boss will not permit discounting. Absolutely not. I'm sorry." Abe looked somewhat deflated, but the two of them stepped into a nearby bar and continued the negotiation over a beer. When they parted an hour later, they had agreed the deal. The discount was 20%.

In the preceding account, Zach lies to Abe. He asserts things he knows to be untrue with the intent to deceive and thus disadvantage Abe. There has been no evidence of interest from other customers and Zach's boss is not absolutely set against a discount – indeed, quite the reverse – as Zach knows. To assert the opposite is to lie. Yet, many see this kind of 'bluffing' as just normal, acceptable negotiating tactics, and therefore morally permissible. And even those who consider it morally reprehensible seem, in the main, to find it *less so* than outright lies told in other contexts (for a review of the empirical evidence, see 3.2 below). Is this view valid?

An obvious starting point would be to ask whether such lying is legal. And yet the position is not entirely clear. As Richard Shell (1991, 93-94) points out, under United States law at least, there is no duty of 'good faith', that is, there is no general obligation to consider anything other than one's own advantage. Zach's claim that there is interest from other customers could be considered a fraudulent misrepresentation, but the instances of courts finding fictitious claims of outside interest to be fraudulent are very few (96-97). As for Zach's lie concerning the availability of a discount, there is a specific exemption in Rule 4.1.a of the American Model Rules of Professional Conduct for "estimates of price or value placed on the subject of a transaction and a party's intentions as to an acceptable settlement." (Gorlin 1990, 372). In other words, in this area, even the lawyers are legally allowed to lie (Shell 1991, 96; Dees & Cramton 1991, 140; Applbaum 1999, 105-106; Burr 2001, 11-12; Peppet 2002, 92;

Hollander-Blumoff 2010, 403; Hollander-Blumoff 2017, 23). So, although there are "many gray areas," (Shell 2018 [1999], 218) Zach probably remains, for practical purposes anyway, on the right side of the law. But is his behaviour acceptable? Is it morally permissible? I shall argue that it is not.

3.1 DEFINITION & MORALITY OF LYING

There is an extensive literature covering the definition of lying (for reviews see MacIntyre 1994; Saul 2012; Mahon 2015). Much of it, though, is devoted to explaining distinctions – some of them very fine ones – that we draw between lies and a careless indifference to the truth ('bullshit') (Frankfurt 2005 [1986], 33-34, 54-55), between deceptive and so-called non-deceptive lies ('bald-faced lies') (Chisholm & Feehan 1977, 159; Carson 2006, 289-290; Sorensen 2007, 259-263), between lying to others and lying to oneself (self-deception, connivance) (Barnes 1994, 87-102), between adversarial lies and benevolent lies ('white lies', jests, prosocial lies) (Erat & Gneezy 2012, 724; Harris 2013, 12-21; Levine & Schweitzer 2014, 108; Levine & Schweitzer 2015, 89), between intentional lies and 'false' use of language (linguistic error/malapropism, metaphor, hyperbole, irony) (Saul 2012, 15-19) and between intentional statements of untruths and unintentional statements of truths (Fallis 2009, 38-39).

However, these distinctions have only a limited bearing on most business negotiations, at least as far as concerns material lies. More often than not, the latter involve lies of a narrow variety: ones that are targeted at others rather than oneself; that are serious and adversarial rather than light-hearted or benevolent; that are literal rather than literary; and that are intentional statements of known untruths. I therefore intend to employ a rather simplified definition: *you lie if, and only if, you assert something you believe to be false with the intent to deceive and disadvantage another*. In doing so, I take what I consider to be the common-sense position that light-hearted hyperbole or other language that is clearly not intended to deceive – statements offered with a smile, such as "you scared me to death" or "my children will not eat tonight unless you agree this deal" – do not constitute lies. They are not intended to be believed; nor are they intended to disadvantage another. This also eliminates from my definition so-called 'prosocial' lies that, while being misleading, are intended to *benefit* the counterparty (Levine & Schweitzer 2014, 108; Levine & Schweitzer 2015, 89; Gaspar et al. 2015, 307). My focus, then, is on what Sanjiv Erat and Uri Gneezy (2012, 724) identify in their taxonomy of lies as "selfish black lies."

Significantly, my preceding definition is limited to active lying, that is, lies of commission. Such lies are of course part of a broader category of deception that includes paltering (use of a factually truthful statement to create a false impression) and passive deception (misleading through omission). I accept that such paltering (Rogers et al. 2017, 461, 465-471) and misleading omissions (Schweitzer & Croson 1999, 243-244) are ubiquitous and can, in some instances at least, be as damaging as outright lies. But I am anxious here to maintain a focus on a specific and clearly identifiable form of behaviour: active lying. By contrast, deception covers a huge range of behaviours – to quote the title of one recent article, it comes in "Fifty Shades" of grey (Gaspar, Methasani & Schweitzer 2019, 62) – and as a result is somewhat amorphous. This of course limits the application of my findings, but it does at least establish a well-defined starting point, from which other work can follow.

As a general matter, I assert that lying is wrong: it is morally impermissible. As Sissela Bok (1999, 30) puts it, "truthful statements are preferable to lies in the absence of special considerations", that is, there is a moral "presumption against lying." But, as Bok implies above, this presumption cannot be rigid. It must allow for "special considerations," exceptional circumstances when the wrong of lying is outweighed by other moral concerns. Contrary to the well-known views of Augustine (2016 [395], 17), Wesley (1831 [1785], 276-277), and Kant (2016 [1797], 83-84), a categorical prohibition of lying is untenable. In Kant's classic test case of the murderer at the door seeking the fugitive within, almost all of us would feel that lying about the whereabouts of the potential victim was not only permissible, but also obligatory, to avoid the greater wrong of murder (Korsgaard 1996, 144; Strudler 2005, 463; Parfit 2011a, 281, 292-293). But such exonerating circumstances are rare.

My opening claim, then, is that lying is pro tanto wrong, that is, morally impermissible unless there are exceptional countervailing moral factors that carry even greater weight. I assert but do not seek to defend this position, other than to point out that it is a widely held view that seems to span many cultures. In an eight-country study of 1,583 students, Harry Triandis and colleagues (2001, 85) concluded that "in all cultures people feel bad when they lie, especially when they lie a lot… suggesting that all recognize[] that lying is wrong." This appears to be close to "a moral universal" (see also Williams 2002, 60; Tyler, Feldman & Reichert 2006, 69; Gerlach, Teodorescu & Hertwig 2019, 4; Abeler et al. 2019, 1116-1117; Sunstein 2021, 26, 50; Serota, Levine & Docan-Morgan 2022, 324).

I have no more to say on lying generally. Instead, in the argument that follows, I focus on the second, more contentious, part of my claim, *viz.* that lying in business negotiations is no less wrong than lying in other contexts.

3.2 THE SPECIAL EXEMPTION THESIS

A number of authors have claimed that lying in business negotiations is acceptable, because different rules apply from those that prohibit lying in other contexts. One of the earliest of these, Carr (1968, 143-153), argued at a wider level that "business operates with a special code of ethics," its own "rules of the game, set by law," and that therefore management is morally bound only to obey the law. Others take a different view, accepting that there are moral obligations in business that extend beyond the law. But they seek to make a particular exception for lying in business *negotiations* – what I call the special exemption thesis.

Variants of this thesis are widely held. Fritz Allhof (2003, 283) argues that "bluffing in business is morally permissible," Alan Strudler (1995, 810) that "it is permissible to bluff … by lying about my reservation price," Karl Aquino (1998, 197) that in a negotiation it is possible for a statement to be "untrue" and yet "ethically justified," and Thomas Carson (1993, 323) that, "It is not expected that one will speak truthfully about one's negotiating position." Carson, Richard Wokutch, and Kent Murrmann (1982, 13) accept that "bluffing and other deceptive practices … do constitute lying" but nonetheless consider them "typically morally permissible," while James Michelman (1983, 255) insists that "given the constraints of economic competition, one *must* negotiate by means of deception." All these statements would appear to imply that there is something special about the process of negotiation that places it beyond the normal moral constraints that attach to other social contexts. And this perspective can be detected too in the broader philosophical literature. For example, Bernard Williams (2002), who concludes that "living in the truth is just a better way" (263), insists that nonetheless "areas of commercial activity" may create "special circumstances," adding that "it is often said that no sensible person expects to hear the truth when buying a used car from a dealer." (109-110).

Empirical research suggests that this attitude extends beyond philosophers and business academics to the community at large. Using a combination of twenty-two in-depth interviews and 240 responses to questionnaires, Vincent-Wayne Mitchell and Joseph Chan (2002) created a "Consumer Ethics Index" (12) to compare UK consumer attitudes to different kinds of unethical behaviour. Of the fifty forms of behaviour compared (13-15), six involved lying and one of these lying in the context of a negotiation. This lie – "not telling the truth about your financial position when negotiating the price of a new automobile" – was rated morally *less* wrong than all but one of the other forms of lying. These others included such seemingly trivial things as "Saying one group member lost his/her ticket for an event as an excuse for getting in front of the queue and getting in sooner" and "Telling the waiter that the 'soup is too cold or it has a funny taste' when there is no problem." The only lie that was considered morally *less* wrong than the one related to negotiation was "Lying about one's age to get a pint of beer." This might suggest that the average Western consumer has relatively few qualms about lying in a financial negotiation.

One must be wary of judging a population at large on such a small sample but this picture is supported by several further studies among business school students, which reveal a similar preparedness to diverge from the truth in a negotiating setting. In one such study, Karl Aquino and Thomas Becker (2005, 666-670) simulated a negotiation between a large car and truck manufacturer ('General') and a potential supplier ('Midwest'), concerning the price of a key component: an aluminium hood. Different scenarios were tested but, in all cases, the ninety-six MBA students representing General knew the hood would become obsolete in three years, making the contract significantly less attractive for the supplier. When asked directly by the Midwest representatives about the expected length of the contract, 2% told the truth, 43% concealed the truth, and 55% lied. In another simulation, involving seventy-five graduate business students, Ann Tenbrunsel (1998, 332-334) found that a similar proportion, 55%, lied about crucial information submitted to an arbitrator, in order to achieve a negotiating advantage. And in a further simulation, involving a faulty computer hard-drive and conducted among 148 students, Schweitzer and Croson (1999, 241) found that 39% lied when asked a direct question and 25% lied unprompted. Similar cases of dishonesty have been observed in many other studies (O'Connor & Carnevale 1997, 504; Murnighan et al. 1999, 332; Boles, Croson & Murnighan 2000, 247; Brandt & Charness 2003, 125; Croson, Boles & Murnighan 2003, 154; Volkema, Fleck & Hofmeister-Toth 2004, 321, 332; Gneezy 2005, 387-388; Guidice, Alder & Phelan 2009, 543).

All the studies cited here were focussed on Western capitalist democracies with typically individualist cultures, where priority is often given to personal rather than group interests. More recent research has examined whether similar attitudes to lying and other ethically ambiguous negotiating tactics extend to other cultures that may be more collectivist. The eight-country study by Triandis et al. (2001, 83) found that in a simulation, where students were asked to declare the production capacity of their company's plant, regardless of the country or culture at least 80% factually lied, with participants from Asian countries being among those who exaggerated the most. Other such studies have also indicated that a preparedness to lie in business negotiations is prevalent across cultures and is by no means a uniquely Western characteristic (Rivers & Lytle 2007, 5-10; Banai et al. 2014, 681; Zhang, Liu

& Liu 2015, 137). Indeed, there is growing evidence that collectivist cultures, such as that of China, may disproportionately favour lying to business counterparts, because the latter are seen as part of an outgroup. That is, they are outside the relevant collective (Ma 2010, 132; Ma, Liang & Chen 2013, 651; Rivers & Volkema 2013, 26 – although, for a recent contrary view see Huynh, Rieger & Wang 2022).

These various studies are not straightforward to interpret, because the parameters differ, and the sample sizes are in most cases small. The broad conclusion though seems clear: across cultures, (1) lying generally is seen as wrong; but (2) lying in a business negotiation is seen as acceptable to many. The latter is viewed as morally less problematic than lying in many other contexts.

So, what is the basis for this distinction? Where is the disanalogy, the moral difference, that allows the apologist to justify this special exemption, to make permissible in a business negotiation what is in other contexts impermissible? Apologists offer five main arguments for such a moral difference. I will evaluate, and seek to refute, each in turn.

3.3 REFUTATION OF THE APOLOGIST ARGUMENTS FOR THE SPECIAL EXEMPTION

3.3.1 Argument from Consent

3.3.1.1 Tacit Consent

The argument from consent is often framed as an analogy between business negotiations and a game. Carr and others cite poker, while Carson has a preference for the board game Risk. The implication is that a special set of rules apply, rules that do not prize honesty as it is prized elsewhere. Carr (1968, 143-153) argues that these special rules represent a "ritualized" convention according to which truth is not required, whereas Carson (1993, 322-323) suggests that in negotiations "each party *consents* to renouncing the ordinary warranty of truth." Allhof (2003, 286-287) goes further, insisting that negotiation *is* a game and one where the moral permissibility of lying is grounded in a principle of mutual "endorsement" (see also Lewicki 1983, 87; Beach 1985, 194; Holley 1998, 633). All these theses ultimately rest on the concept of tacit consent. Whether those involved have acknowledged a ritual, have agreed that statements are not warranted to be true or have mutually endorsed falsity, the underlying premise is the same: that everyone in the negotiation – everyone 'at the table' – has accepted that they may lie and be lied to (for a similar analysis of how such a "game frame" undermines the moral character of those in the legal profession, see Cohen, Helzer & Creo 2022, 20).

I argue on two grounds that such universal consent does not exist. First, that the argument from tacit consent faces a reductio ad absurdum; and second, that it is simply inconsistent with the facts.

To take the first, suppose we grant, for the sake of argument, that there is universal tacit consent that lying is admissible. If so, then it cannot make a substantive difference if that tacit consent is made explicit. At least, if that were not the case, then the position would appear troubling: it would fail the test of publicity that has been thought morally imperative among philosophers from Immanuel Kant (1917 [1795], 184-185) to John Rawls (1971, 133). Simply to give voice to what is already understood and accepted by all should not be problematic.

But, if the *claimed* tacit understanding *is* made explicit, then the dialogue of negotiation becomes all but emptied of meaningful content. Let us take, as an example, our opening scenario. Assuming that the supposed tacit consent is made explicit, then the dialogue would read somewhat as follows:

> *Abe*: "Zach… we would be looking for a discount. Only fair to tell you."
>
> *Zach*: "Abe, you know I can't do that… There is considerable interest from other customers; and my boss will not allow a discount. Absolutely not. I'm sorry. Although I must add that this may or may not be true."

Zach's reply is only one short step from the following:

> *Zach*: "Abe, you know I can't do that, although perhaps I can… There is considerable interest from other customers, although this may not really be the case; and my boss will not allow a discount, but then perhaps she will. Absolutely not – well maybe. I'm sorry. Again, I must add that all this may or may not be true."

What Zach says here is not false. Indeed, in grammatical or logical terms, it is a tautology and therefore necessarily true. However, the dialogue has been emptied of meaning. It no longer fulfils the task of communication. It is neither truth nor lie. It draws close to Harry Frankfurt's (2005, 42-43) "bullshit" or "hot air": "It is mere vapor."

The apologist might object to the assumption in my argument above that a tacit agreement of lying applies to everything. In reality, she might feel, it applies only to certain areas of negotiation. There is evidence, for example, that people are more willing to lie when the potential harm to their counterparty is trivial than when it is more serious (Gneezy 2005, 387-388). But the same result holds. Consider the case of a used-car salesman and customer. The salesman might argue that, while he would never lie about the car's brakes, there is nonetheless a tacit understanding that he may lie about the air-conditioning. Given that, it is of course possible for the two parties to have a meaningful conversation about the brakes. But, if both salesman and customer have really tacitly agreed that whatever words the salesman utters in relation to the air-conditioning need not relate in any way to the truth, what possible purpose can a discussion of the air-conditioning fulfil? It is meaningless.

My argument, however, remains open to a further objection. Strudler (1995), in particular, claims that deceptive statements can still fulfil a purpose as a signalling device, a "trial balloon" (817) indicating intensity of feeling. He suggests that they can "express one's commitment to a price range or to the idea that one has something quite valuable to offer" (816-817). But how effective can such verbal signalling be, if there is a tacit agreement that everything said on the subject may be a lie? Any statement so qualified loses much, if not all, of its force. I would suggest that there are other more effective ways of communicating value than signalling with lies, such as describing the unique features and efficacy of the product or service on offer. But I will return to Strudler's argument when I examine the consequentialist case for lying in 3.3.3.

The point here is that communications during a negotiation are only effective if there is *some* presumption of truth. And they can only serve their purpose *to the extent* that they are believed to be true, or at least likely to be true. If statements are genuinely 'unwarranted' – either true or false with unspecified probability – and are understood as such, then rational negotiators will ignore them. And, actually, this is what often happens in practice. To return again to our scenario, if Abe believes absolutely nothing Zach says, he is as likely to reply as follows:

> *Abe*: "Zach, let's cut the bullshit. What's your best price? Just give me a number."

There is an illuminating parallel here with the world of corporate social responsibility (CSR). David Hess and Thomas Dunfee (2007, 19) have noted that the lack of assurance (regulation or standardised requirements) around corporate statements on CSR has rendered them "cheap talk", which increasingly cynical stakeholders learn to heavily discount or even ignore. They call this a "babbling equilibrium": meaningful communication has ceased. In the environmental arena, this kind of unsupported, distrusted and largely empty communication even has its own name: 'greenwash'.

In short, meaningful communication and exchange require some degree of trust. In contrast – and to borrow from Ivan Preston's (1994, 201) study of the advertising market – "When we act with utter distrust we stop acting; we paralyze ourselves. We can no longer participate in the marketplace." It seems unlikely that negotiators have really chosen that option, that is, to set aside all trust – tacitly or otherwise.

This reductio ad absurdum leaves us with two possible explanations of Zach's lying. The first is that Zach really believes there is a tacit agreement that he may lie and *intends* his lie to be understood as empty of meaning, innocent "babbling," "mere vapor." In which case, we might ask why he bothers. The second is that Zach knows there is no such tacit agreement and hopes to be believed, *at least to some degree*. His lie has purpose. He seeks to influence Abe. I believe the latter explanation is more credible.

Importantly, I do not argue that empty conversation is morally impermissible. Instead, my point is that the liar does not intend his communication to be understood as empty; he aims to be believed, at least partially. This undermines the argument from consent. After all, a genuine tacit agreement that he may lie would defeat the very purpose of him lying.

3.3.1.2 Universal Consent and the 'Closed Circle'

My second argument against the case for consent is rather more straightforward: such consent is not consistent with the facts. The apologist argument is that, while we each have a right not to be harmed by others' lies, each of us has consented to give up that right in the context of a negotiation. But, because a right can only be surrendered by the holder of that right (*I*, for example, cannot give up *your* rights), this implies universal consent among those involved in negotiations. And yet it is *empirically* false that *all* business negotiators have accepted that they may be lied to in this way.

The act of participation in an adversary game, even when the 'rules' are known to everyone, does not itself signal consent, because, as Bok (1999, 83) says, "there may be no way for [you] to 'leave the game'" or, at least, the costs of doing so may be very high. In a similar vein, *expectation* is not the same thing as *consent*: that you *know* you may be lied to does not mean that you *agree* to being lied to (Applbaum 1999, 4, 116). As Gil Hersch (2020, 213) argues, just because your daily commute requires you to live in a high-crime neighbourhood, that does not mean that you consent to being mugged.

In his reply to Allhof, Carson (2005, 401) makes this same point: "It is simply false that all participants in business negotiations endorse the practice of misstating one's reservation price." He goes on to cite the example of a naive adolescent buying a used car; and insists that he, Carson, like that naive adolescent, does not consent to being lied to (see also Sullivan 1984, 9; Koehn 1997, 1449; Provis 2000b, 8; Piker 2002, 339).

Behind this there lies an important fact of human nature: most of us *are* very trusting (Sunstein 2021, 73-76). As Roderick Kramer (2009), succinctly puts it, "To trust is human" (70). Kramer believes that this *"*presumptive trust*"* (71) derives from our evolutionary history, where it increased the survival chances of vulnerable infants. Whether that is the case or not, he is surely right that, for many of us at least, "trust is our default position" (71). And, for many of us, that trust is not set aside simply because we have entered into a business negotiation. The many financial scandals associated with the breach of that trust – for example, the fraudulent sale of payment protection insurance (PPI) in the United Kingdom and of the Madoff Ponzi scheme in the United States – are clear evidence of how trusting many of us are.

My claim, though, that there is no universal consent is open to challenge by reference to an extreme environment like the tourist bazaar. Here it can be argued that the expectation of being lied to is so marked and so universal that anyone who chooses to bargain there must surely consent to being lied to. Their very act of voluntary participation signals their consent. I accept that there is *some* substance to this objection. But my acceptance is qualified in two respects. First, for many people this is a grudging consent: their consent is, as it were, forced from them as a condition of their participation. It is a case of 'consent or go without'. Second, and perhaps more importantly, the clearly *in*essential nature of the purchases involved is vital to the claim of voluntariness. Such deemed consent cannot be extended to truly essential purchases, such as the provision of shelter, clothing, food, or basic transportation for one's family. Nor can it be applied to a broad array of typical business dealings, where business managers, by the nature of their roles, are *required* from time to time to participate in negotiation. Given the obligatory nature of this requirement, their mere participation cannot be assumed to imply their consent to being lied to.

In a commentary/reply to an earlier version of this chapter (Sherwood 2022), Strudler (2023, 768-771) argues that in many cases, where there is no consent to being lied to, there is nonetheless an "assumption of risk" (769) of being lied to. The distinction between Strudler's assumption of risk and the familiar notion of *tacit* consent seems to me a very fine one, and I am unsure what the former concept really adds. In both cases, participation is said to signal acceptance of the risk of being lied to. And in both cases, as Strudler (769-770) readily acknowledges, the question of voluntariness is vital. Strudler cites, as an example, American football, where a player by joining the game clearly communicates an acceptance of risk, in this case the risk of being tackled. Strudler then seeks to draw a similar implication for the assumption of risk of being lied to in business negotiations. But there is a crucial difference. The decision to play American football, like the decision to play poker, is clearly a *voluntary* assumption of risk. Nobody is forced to play American football. By contrast and as emphasised above, most of those in business are, by virtue of their employment, *obliged* to participate in negotiation. They have no choice. Their participation cannot, therefore, of itself signal a voluntary acceptance of being lied to, whether tacit or assumed. The same is typically true of a first-time car buyer negotiating with a used car salesman, or a young couple purchasing their first home. Many in these roles do not 'assume' the risk of being lied to. They have that risk *imposed* upon them. And, I would argue, wrongfully so.

The apologist argument from consent faces a further challenge. Most theories of consent require that such consent (tacit or explicit) be sought and received from all those affected – all those who could potentially be harmed. The apologist makes an implicit assumption that, as in a game of poker or Risk, all those affected are seated 'at the table': that they form what is sometimes called a 'closed circle'. Only thus can the apologist claim the required level of consent. This assumption is generally fair in poker. Most of the time, it is true that only those around the table are materially affected by the outcome of the game. But, while this may be largely true in poker, it is almost uniformly not the case in business negotiations. Almost never do business negotiations affect only those at the negotiating table (Badaracco 1997, 19-20; Applbaum 1999, 116, 134). As Dees and Cramton (1991, 152) point out, "Third parties, who had no choice, may be affected." They give the example of small communities impacted by labour negotiations with the community's major employer (see also Bowie 1985, 288, 290; Koehn 1997, 1452). Equally, we may return again to our opening scenario, where Abe is buying a *pharmaceutical* packaging machine. If the machine costs more, so in all likelihood will the pharmaceuticals. Ultimately this implies greater costs or less availability for customers, including public health systems, patients, and taxpayers. In addition, the outcome of Zach and Abe's negotiation will impact their various competitors. In business, there is no 'closed circle'.

The argument from consent, then, not only faces a reductio ad absurdum. It is also inconsistent with the facts. Everybody 'at the table' has not consented to being lied to and many of those affected by the negotiation are not even sat 'at the table' in the first place.

3.3.2 Argument from Self-Defence

Some have defended lying in business negotiations on the grounds of self-defence. For example, Dees and Cramton (1991, 148) condone a policy of "defensive dishonesty," whereas Strudler (2005, 462) claims that "certain deceptive acts that occur in negotiation are morally acceptable because of their defensive function" (see also Tenbrunsel 1998, 331).

Carson has been a particularly strong advocate of this approach with his "principle of self-defence" or SD, as he abbreviates it. Carson's SD states that, "It is permissible to lie or attempt to deceive others about one's negotiating position" provided that various necessary and jointly sufficient conditions are met. The real work is done by the first of these conditions, which is that "one's negotiating partner is doing the same and is likely to harm one thereby." Carson (2010a, 356-357) emphasises that it must be "*clear* that the other person is misstating his intentions." In the absence of such clarity, "the default position should be not to lie or deceive others" (2010b, 194-195).

This conditionality severely restricts the application of SD. This is because you can sometimes know with a high degree of certainty that your counterparty is lying; and you can at other times be harmed by your counterparty lying; but it is hard to see how *both* can be the case at one and the same time. After all, you cannot be harmed by deception, when you are not deceived. You may be harmed, of course, by a lack of information, but a lack of information is not itself a lie. A lie requires a false statement and, if you do not believe that false statement, it seems unlikely that it can harm you. It follows that, when self-defence would be justified, it is not actually needed. Indeed, in the absence of any harm to be averted, your counter-lies cease to be self-defence and become mere retaliation. Thus, the conditions required to justify SD produce a 'nil set'.

Carson anticipates this objection. In response, he argues that you *can* be harmed by deception even when you are *not* deceived – that is, that the set is not nil – because even failed attempts at deception can still cause one to "reveal one's own preferences" through "expressions of emotion" and "create uncertainty or fears that weaken one's resolve" (Carson 2010b, 197). Carson may be right. Perhaps there are emotional pressures that a liar can exert even when his lies are detected as such. For example, a bald-faced lie might demonstrate dominance and disrespect, placing the counterparty at a psychological disadvantage. But this feels like a marginal argument and Carson admits as much in his concluding remark: "Nonetheless, I concede that this objection supports the view that my SD principle *seldom* justifies lying or deception in negotiations" (2010b, 197; also 2010a, 358).

However, there is a rejoinder available to an advocate of SD. It is possible that you might know your counterparty is lying or will lie to you, without knowing for sure *which* statements of hers are lies. A clear case is where the counterparty makes two statements that are inconsistent: both cannot be true, but you may not know which is false. In theory at least, this *is* an instance where you can know the counterparty is lying *and* may be harmed by that lie. Again though, this feels very much at the margin. After all, if you know for certain that your counterparty is lying regarding at least *some* of what she states, are you really likely (absent other supporting evidence) to attach great credence to *anything* that she states? And sufficient credence to expose you to harm? It seems unlikely.

All the foregoing casts doubt on whether you can *knowingly* be harmed by a lie; but, even if you *do* know you are being harmed, it is far from obvious that lying in response would remove that harm. In the case of physical attack, most justifications of self-defence sanction actions that are designed to parry the attack, that is, to fend off the blows and prevent further blows. It is difficult to see how retaliatory lying would achieve anything equivalent.

The apologist could perhaps argue that retaliatory lying, while not removing the original harm, might nonetheless neutralise it in some way, through a balancing of the scales. In this vein, she might appeal to a distinction sometimes drawn between justifiable conduct and excusable conduct (Cohen 2000, 158). Joseph Heath (2018a) characterises the latter as morally undesirable behaviour that has become habitual throughout a business domain and which all participants must emulate in order to remain competitive. Such behaviour is seen as excusable in that, although it is wrong, it is not considered blameworthy (523-526). The implication would be that, if everybody else is lying, then you cannot be blamed for lying yourself, in order to 'level the playing field'. But Heath sets two conditions for such unjustifiable yet excusable conduct. First, the cost of not behaving in this way – in this case, of not lying – must be "unreasonably high." (529). This is rarely the case with lying, as I seek to demonstrate in 3.3.5. Second, the party must not condone the practice as a whole, but instead must take "positive action" (529) to curtail it. This would certainly rule out pre-emptive lying. And, as Heath goes on to warn, under this "'everyone else is doing it' argument" (520), there is a real danger that "'we had no choice, our competitors were doing it'" can all too easily slip "between 'were doing it' and 'were *going* to do it', 'were *probably* going to do it', or '*might* have done it'" (526). The twin dangers of pre-emption and escalation are all too obvious.

What is clear is that Carson's 'principle of self-defence' at best justifies the apologist position in a very narrow group of cases. For example, it certainly could not justify Zach lying to Abe, as he does in our opening scenario. There is no evidence that Abe has lied first, or indeed ever lies, and therefore no case for self-defence, as described by Carson.

The self-defence argument above might, though, be reframed by the apologist in the language of reciprocity: if your counterparty lies to you and thus fails to show you what Stephen Darwall (1977, 40) calls "moral recognition respect," perhaps you might then be relieved of any obligation not to lie to her. But this line of argument sits most easily with acts that are supererogatory, where my uncalled-for gesture of help, kindness, or courtesy places you under an implicit obligation to reciprocate; and, to the contrary, where my failure or refusal to make such a gesture leaves you without any such obligation. It sits much less easily with circumstances involving harm. To take the extreme case, you would not typically view the obligation to refrain from murder as one imposed by the principle of reciprocity. If someone commits murder, you may have a right of self-defence as described above, but not a right to reciprocate. In like fashion, the principle of reciprocity alone seems insufficient to justify the inflicting of harm through lying. Such a response would best be described not as one of reciprocation, but rather as one of retaliation. Stripped of its defensive aspect, reciprocity of this kind becomes a case of 'a wrong for a wrong', one harm justified by another – the morality of the gangland and the blood feud.

This concept of negative reciprocity is similar in some respects to a more permissive version of the argument from self-defence, advocated by Dees and Cramton (1991, 144). They argue for a "notion of defensive fair-play," based on their "mutual trust principle." This formulation has gone through revision over time, and it remains unclear whether it should be seen first-and-foremost as a moral or a prudential maxim (Dees & Cramton 1995, 829-830). But it essentially states that the obligation to negotiate honestly is diminished, or indeed eliminated, when there is an absence of mutual trust between the parties. In effect, this inverts the burden of proof stipulated by Carson. To justify lying, you need no longer prove that your counterparty is lying. Instead, you need only satisfy yourself that you have no reason to trust that he is not lying.

While Carson's SD seems too restrictive, Dees and Cramton's mutual trust principle would appear too permissive. In many instances, it would seem to raise again Heath's concern above and justify a pre-emptive lie, with all the escalation effects that entails. It is, as the authors acknowledge, a principle suited to "a Hobbesian state of nature" and counter-intuitive to most of us (Dees & Cramton 1995, 828). Furthermore, it is vulnerable to extension ad absurdum. As Strudler (1995, 807-808) points out, "the Principle seems to allow one to do anything to one's opponent, no matter how horrible." If I am permitted to lie to you in the absence of positive evidence that you will not lie to me, then presumably the same principle must apply to other forms of pre-emptive injury, such as unprovoked physical assault. In other words, in the absence of evidence that you are not going to attack me, it would appear morally permissible for me to attack you. Such a pre-emptive right of injury, whether it be through physical assault or lying, is untenable.

Advocates of these self-defence arguments face then a two-horned dilemma. On the one hand, if Zach is required to be nearly certain that he is being lied to by Abe (in line with Carson's SD), lying in self-defence is almost never justified, because Abe's claimed deception cannot harm Zach when Zach is not deceived; and anyway, it is unclear how a counter-lie renders the original lie harmless. On the other hand, to the extent that this requirement is weakened (in line with Dees and Cramton's mutual trust principle), then to that extent Zach is permitted to lie pre-emptively but also, by extension of the same principle, to commit all sorts of other unprovoked injuries – injuries that cannot possibly be justified. It appears that any apologist case based on self-defence is thus, at best, extremely limited.

It seems the apologist must still look elsewhere for a cogent defence of the special exemption. The foregoing debates rest on foundations that are essentially deontological in character. They assume implicitly that people have rights and, in particular, a right not to be lied to; but that this latter right can be waived through consent or forfeited if, by threatening others, you give them a right of self-defence. I will now turn to an apologist argument from a different ethical tradition, consequentialism, that is, an argument that appeals to the consequences of lying to assess its moral rightness or wrongness.

3.3.3 Argument from the 'Greater Good'

Applbaum (1999, 6, 187-197), Jonathan Cohen (2002, 118), and Nathaniel Davis (2016, para.2) all see the 'greater good' argument lurking behind many of the special exemption claims of the 'adversary professions' – in business, as well as in law and government. Each cites Adam Smith (2012 [1776]) and his theorem of the 'invisible hand'. They do not suggest that Smith himself advocated lying. But they do claim that behind many of the arguments, marshalled in support of the permissibility of lying in negotiations, there rests an assumption that, through mechanisms similar to those Smith describes, society's greater interest is served – specifically, that the self-interested pursuit of individual gain, in an adversarial negotiation, generates a clearing price for the market and ensures an efficient allocation of resources to the benefit of society as a whole; and crucially that, should part of that amoral, self-interested pursuit of individual gain involve the use of deceit, then this is permissible as long as it is an essential element of the above mechanism that serves that morally laudable objective, the greater good. Or such anyway is the apologist case. It is, for example, an argument that has been employed to defend 'spoofing' (deceptive orders and order cancellations) in high-frequency trading (Cooper, Davis & Van Vliet 2016, 3, 10-13; Khorasanee 2022, 11). But, as Applbaum (1999, 180) warns, "the appeal to the good ends of adversary institutions in equilibrium is notoriously underdemonstrated." A closer analysis is needed. Let us break down the key claims into two components.

First, is business negotiation, in particular price negotiation, vital to the 'greater good'? This is contestable, but I will happily concede to the apologist that, in many instances, the answer is *'yes'* – at least in a liberal, capitalist, market economy.

But the apologist must also prove a second contention: that lying is an essential element of that business negotiation. This contention is much more doubtful, and I argue false. It is important to note that the question here is not about protecting the interests of individual protagonists (I address this in 3.3.5); rather, it is about the successful maintenance of the system for the benefit of society as a whole. In this context, it is very hard to see how lying is beneficial, let alone essential. Seen from the perspective of society as a whole, negotiation is about communicating relative availability and need, so that these can be better matched. Lies are false communications. They introduce an *information asymmetry*, where one party has more information than the other, that is, knows the truth, while the other is deceived. To return to the example of 'spoofing', it has been argued that, because other traders are able to see only half the trade, they are left "misinformed about the real supply and demand for financial contracts" (Hersch 2020, 218). In economists' terms, such misinformation, or asymmetry of information, represents a market *failure*. An efficient market requires properly matched information (Applbaum 1999, 192; Heath 2014, 37) and lying undermines this. This would suggest that, although a suitable process of negotiation may be crucial to delivering the greater good, that process should *not* include lies.

This conclusion, however, faces a potential objection. While the consequentialist case for lying in negotiations remains largely implicit, one particularly strident apologist, Strudler (1995), does make the case explicitly with his "mutual advantage view" (818). Posing the question "Is deceit wrong?" he insists that "no plausible case [can] be made even for the *prima facie* wrongness of deception" (812). Strudler argues that, on the contrary, deception is "a mutually advantageous tool" (813) that creates "a safe device for indirect communication" (818). He bases this, as we saw earlier, on the role of lying as a form of signalling. In other words, Strudler believes that lying actually improves the flow of information and, presumably therefore, enhances market efficiency and the general welfare (816-817).

Dees and Cramton (1995, 824) have pointed out that there is one sense at least in which Strudler's argument is self-defeating. If the purpose of deception is really to improve information transparency and to enhance the knowledge of the counterparty, then it is not deception. And in practice we do not treat it so. Sometimes a comment is so very obviously figurative, light-hearted, or clearly intended to mean something beyond its face value, that it cannot be interpreted literally. It is not intended to deceive, nor to disadvantage, and, consistent with our opening definition, does not therefore constitute a lie. But this does not characterise most lying in negotiations, especially when such lying is between parties who do not know each other well or relates to matters other than reservation prices. Most 'deceit' is just that: intended to deceive (even if it actually fails to do so). It does not aim to improve information transparency or market efficiency.

It is possible though to put a different interpretation on Strudler's mutual advantage view. One might accept that individual deception is self-interested and intended to deceive, while arguing that somehow these various deceptions interact with each other to create a mutually advantageous outcome. This would appear consistent with Strudler's (1995, 818; 2023, 771) signalling thesis; and with his unconventional claim that lying is the solution to the lack of trust between the parties, rather than its cause. But again, it is hard to see how Strudler's logic would apply beyond disclosure of the reservation price and, even then, it would seem to fly in the face of empirical evidence that supports the more conventional direction of causality, that is, that it is lying that creates distrust, not the other way around (Boles, Croson & Murnighan 2000, 255; Tyler, Feldman & Reichert 2006, 75; Schweitzer, Hershey & Bradlow 2006, 15, 17; Gaspar & Schweitzer 2013, 168). Indeed, Raymond Friedman (1993), studying success and failure in labour negotiations, has concluded that "the one factor that is consistently most important is trust" (435) and that *the* crucial contributor to that trust during the bargaining process itself is "believing what the other side says" (451). It is by *not* lying, by instead developing a reputation for speaking honestly, by building confidence, that you earn the right to be believed, to have what you say taken at face value. It might seem naively simple, but the most effective communication is when someone who is trusted tells the truth.

The worst instances of the deceit that Strudler advocates can lead to a failure to transact at all. Strudler (1995, 816; 2023, 771-774) describes a personal experience, which seems to illustrate just this point. He recounts an incident in Madrid, when a taxi driver "feigned shock" at the sight of Strudler's extensive luggage, declaring that there was simply too much to be accommodated in the car (a lie). Strudler took this at face value and chose to rent a car instead, leaving the taxi driver outraged that Strudler had not sought to negotiate further. Strudler seems to feel that this demonstrates the importance of correct signalling, but it appears an equally potent illustration of the adverse consequences of lying in a negotiation.

As Leslie John (2016, 117) warns, far from being an essential element of negotiation, “Lying … can be a real impediment to the creation of value in negotiation.” There is a danger that, as in Strudler’s case, false statements about the parties’ bargaining positions lead to a breakdown in negotiations, where actually a mutually beneficial transaction is possible (Dees & Cramton 1995, 825; Wertheimer 1996, 237-238; Hollander-Blumoff & Tyler 2008, 475-476; Rogers et al. 2017, 468-469). This is particularly the case with reservation prices. For example, let us suppose that Abe is actually prepared to pay 90% of list price, while Zach and his boss, pushed to the limit, would accept as little as 75%. This is clearly a transaction that ought to be easily achievable as in the interest of all. After all, there is an overlap in acceptable outcomes (i.e., a cooperative surplus) equivalent to 15% of list price. But, if Zach insists (falsely) that his boss will not countenance *any* discount and Abe (equally falsely) claims he can pay no more than 70% of list price and they both believe each other’s lies, then there is a real risk that they will fail to reach agreement at all and ‘walk away’. A 15% cooperative surplus has been transformed, through mutual lying, into an apparent 30% gap and a ‘deal that should happen’ from a consequentialist perspective instead collapses.

Furthermore, there is a wider concern. As I have argued already, repeated lying undermines trust (Boles, Croson & Murnighan 2000, 255; Tyler, Feldman & Reichert 2006, 75; Schweitzer, Hershey & Bradlow 2006, 15, 17; Gaspar & Schweitzer 2013, 168; Sunstein 2021, 27). This in turn threatens the cooperative efforts that underpin social life (Lewicki 1983, 74). For example, Dan Ariely (2009[2008], 215) reports a conversation with an Iranian student at MIT, who complained that the lack of a platform of trust in his country meant that no one would pay in advance or offer credit, and all hiring had to be done within the family. Normal commercial relations were no longer possible.

As Asha Rao and Stuart Schmidt (1998, 669) argue, trust “limits transaction costs,” because it encourages more informal exchange of information, without the need for exhaustive validation (see also Arrow 1973, 314; Burr 2001, 14; Bowie 2017, 31-32); and conversely, lying increases those costs and can even lead to a collapse in the means of exchange. George Akerlof (1970) vividly illustrates this latter point with his seminal analysis of “lemons” (over-priced, poor quality, used cars) in the automotive industry. Because of a lack of honest information, customers are only prepared to pay the price consistent with the average quality of all cars available in the market. Sellers of high-quality cars know their cars are worth more than this and therefore withdraw from the market. As a result, the average quality and consequent price of cars in the market declines and more sellers withdraw. Before long the only cars available for sale are the ‘lemons’, customer trust is undermined, and the market is threatened with collapse, thus highlighting “the major costs of dishonesty” (Akerlof 1970, 495; see also Arrow 1973, 307-308). As Roger Sullivan (1984, 2) warns, this “leads to the destruction of the system of value traded for fair value on which business as an enterprise is based” (see also Hare 1992, 11-12; Bowie 2017, 32).

Communication and cooperation within a society rely on the presumption that most people are telling the truth most of the time. The survival of this presumption cannot be taken for granted, as much empirical work has attested. Research reveals not only a tendency for lying to lead to more lying by the individual concerned (Barnes 1994, 44; Shell 2018 [1999], 177; Gaspar, Methasani & Schweitzer 2019, 72); it also has an escalatory effect, where early use by one party encourages subsequent use by the other, a form of “behavioral contagion” (Reid et al. 2023, 262-264; see also Volkema, Fleck & Hofmeister-Toth 2004, 333; Tyler, Feldman & Reichert 2006, 75; Sobral & Islam 2013, 290-291; Olekalns, Kulik & Chew 2014, 24). This “ethical fading,” as it has been called, weakens those very bonds that hold societies together and facilitate communal life (Tenbrunsel & Messick 2004, 224; Rees, Tenbrunsel & Bazerman 2019, 28).

But this may still understate the importance of trust. After all, for most of us, it has more than just an instrumental value: “trust is a basic human aspiration” (Banai et al. 2014, 673), something that meets one of our fundamental needs as social beings. It has a status that borders on a ‘good’ in its own right – “a social good to be protected just as much as the air we breathe or the water we drink” (Bok 1999, 26).

In short, the consequentialist case for negotiation is strong; but the consequentialist case for the permissibility of *lying* in negotiation is not. There is no necessary conflict between the search for truth and the search for agreement. Indeed, quite the reverse is true. And the costs of lying for society as a whole are considerable.

3.3.4 Argument from Fiduciary Duty

The fourth apologist argument is somewhat different in character, or at least in its target. It does not need to contest that lying in business negotiations is pro tanto wrong. Instead, it argues that whether it is pro tanto wrong or not is largely irrelevant in practice, because any pro tanto wrongness is consistently outweighed by another consideration: that of management’s fiduciary obligation to its stakeholders. The argument runs that managers may be obliged to lie in order to optimise the outcome for the other stakeholders in the enterprise. This is their fiduciary duty, and it stands before any 'queasiness' concerning falsehood.

There has of course been a longstanding debate as to which stakeholders are owed such an obligation. The traditional emphasis on owners is associated with Milton Friedman (1962, 1970); whereas others, led by R. Edward Freeman (1984), have advocated a more inclusive approach, extending to customers, suppliers, employees, and the community at large. A popular solution among practitioners has been to refer to a duty to protect and further *the interests of the company*, thus avoiding any specific or preferential allegiance, while retaining the notion of a fiduciary obligation to promote the greater interest of a wider constituency that still falls short of society as a whole.

This particular language is very much tied to that of business, but it has its parallels elsewhere. It is part of what the lawyer Charles Curtis (1951, 16) calls "the special moral code which governs a man who is acting for another." Applbaum has made a close study of the legal profession in this regard, examining the moral status of lawyers, making statements in or out of court that they themselves believe to be false, based on an avowed obligation to promote the interests of their clients. This "argument of constituted description," as Applbaum (1999, 89) characterises it, claims that a lawyer fulfils different roles as a lawyer and as an individual. In the former role he has a duty to his client, which places him beyond many of the strictures of any broader moral code. As Curtis (1951, 9) puts it rather baldly and contentiously, "one of the functions of a lawyer is to lie for his client."

A similar distinction is evident too in the British Civil Service. Sir Michael Quinlan (1993 [1990], 542), a past permanent secretary of the Ministry of Defence, describes the need occasionally to bow to ministerial wishes and defend "a square circle." He does not countenance direct lying, but does believe that unelected officials, once they have argued the case in confidence with ministers, are duty bound to promote it in public, whatever their private reservations (Applbaum, 1999, 61-74).

These different appeals to a fiduciary obligation – the business executive's based on a duty to protect his company's interest, the lawyer's based on a similar duty to protect that of her client, and the civil servant's based on the principle of democratic accountability – vary in degree, but they all point to a common distinction between principal and agent. The argument is conceptually similar in all three cases: that, in the event of conflict, the agent has a moral obligation – to protect and further the interests of the principal – that overrides other moral considerations. Accused of deceit or other wrongdoing, the agent may respond that they are merely fulfilling the duty to their principal that their role requires. If, in the case of a negotiation, protecting the interest of the principal entails lying or distortion, and that is the will of the principal, then such also is the obligation of the agent.

I believe this argument fails for four reasons. First, it is questionable whether lying really is in the interest – at least the long-term interest – of the principal. In the corporate context, we should ask whether it truly benefits the company to have its executives lie on its behalf. Most progressive companies put considerable store by their reputation for honest dealings, a reputation often promoted and defended at great expense. But I will leave this point aside for now and take it up again later when I examine the argument from practicality in 3.3.5.

Second, the weakness of the principal-agent argument comes into much clearer focus when we view the situation from the perspective of the victim, *viz.* the recipient of the lie. Through the action of others and through no fault of their own, the victim's interests have been set back; the victim has been wronged. The agent and principal may seek to shuffle between them the responsibility for this wrong, but such shuffling does not make it go away. It does not negate the wrong. The moral impermissibility of the lie is not extinguished (Applbaum 1999, 11, 109).

Furthermore, the principal-agent argument can be extended ad absurdum, in much the same way as was the case with the permissive version of the self-defence argument. If the agent can justify lying in the interests of her principal, then why can she not similarly justify any manner of other wrongs, all the way from fraud to murder? This is no merely abstract concern. Many appalling crimes have been excused in just this manner. In particular, a common defence of war crimes has rested on a duty to carry out orders from higher authority – in other words, the duty of agents to fulfil the wishes of their principals. To take the very obvious example of the Nuremberg Trials, as the Holocaust Encyclopedia reminds us, "the defendants generally acknowledged that the crimes they were accused of occurred but denied that they were responsible, as they were following orders from a higher authority" (US Holocaust Memorial Museum 2021). I am not of course suggesting that lying can be equated with genocide, but the argument from fiduciary duty is essentially the same in both cases, revealing its structural weakness.

Finally, the fiduciary duty defence of lying fundamentally misinterprets the relationship between principal and agent. If the negotiator is, to employ Friedman's (1970, 211) language, "the agent of the individuals who own the corporation" or for that matter of the stakeholders more broadly, then it follows that the manager's obligation must be to stand in the owners', or stakeholders', shoes. Crucially though, this implies that *what would be morally impermissible* *for the owners must also be morally impermissible for the agent*. The obligation of truthfulness is replicated, not exempted. As most commentators agree, "what is wrong for principals to do directly does not become right when delegated to their agents" (Cohen 2002, 119) and "I cannot (ethically) *hire* done on my behalf what I would not (ethically) *do* myself" (Goodpaster 1991, 68; see also DeGeorge 1992, 65-66; Bowie & Freeman 1992, 9; Quinn & Jones 1995, 35; Heath 2014, 276, 315-316; Bowie 2017, 194-195).

The apologist might rejoin that, whereas the owners are free to choose their own moral position, the manager as agent is bound by undertakings and, specifically, by an implicit promise to maximise profits whatever that takes. But such a promise would itself be morally impermissible and void. Applbaum (1999, 8) captures the point well: "The answer cannot simply be that the professional has promised. Whether the promissor is a contract killer or a contract liar, a promise to wrong another has no moral force" (see also Quinn & Jones 1995, 35-36; Carson 2010b, 173; Heath 2014, 275).

The fiduciary duty, or principal-agent, defence of lying fails then on multiple grounds. Its basic premise that the interests of the principal are best served by lying is highly questionable (see further in section 3.3.5); its shuffling of responsibility between offending parties cannot erase the wrong suffered by the victim; and its line of argument by extension permits any number of other wrongs that are clearly impermissible. Above all though, it essentially misinterprets the relationship between principal and agent, a relationship that cannot turn wrong into right.

3.3.5 Argument from Practicality

The final apologist argument is different in character again. It rests not on a question of principle, but rather on one of practicality. It insists that lying in a negotiation must be permissible, because a prohibition on such lying is simply infeasible to abide by, placing an unworkable burden on the negotiator. The typical test case is the reservation price. Surely, the argument goes, a negotiator cannot be expected to reveal her reservation price. That would be to unreasonably weaken her 'hand'. And yet, the argument continues, if she is asked "Is that your best offer?" and does not reply in the affirmative – *even if that is a lie* – then she has done just that. She has undermined her own negotiating position. This pragmatic defence of lying is, as far as I can tell, absent from the philosophical literature, but it is regularly cited by practitioners actually involved in negotiation. When challenged as to why they have claimed dishonestly that their offer was "best" or "final", they are prone to reply along the lines of "Well, what else was I supposed to say? If I'd said anything different, the other side would have known I had more in my pocket!" I will offer two responses to this argument from practicality: a principled response and a pragmatic response.

The principled response might be briefly characterised as 'so what?': Doing the right thing often is difficult and disadvantageous, but that doesn't mean it's not the right thing. After all, there are plenty of parallel instances, where we would not consider potential disadvantage an adequate justification for morally impermissible behaviour. For example, few would accept that, because not participating in fraud, insider trading, or corporate espionage disadvantages your position, that such disadvantage justifies, or even excuses, you committing such acts. Furthermore, it is quite normal for companies to incur significant costs pursuing corporate citizenship/CSR agendas that go well beyond the requirements of the law (for an extensive list, see Norman 2011, 49). Adopting a similar logic, perhaps we should conclude that, however disadvantageous a prohibition on lying may be, that is the necessary price demanded by morally permissible conduct.

But I accept that there is something unsatisfactory about this principled response – at least when unsupported by a practical course of action. As Badaracco (2002, 39) says, "No one manages to climb even part way up the greasy pole without paying close attention to his or her interests." For a business manager to act consistently against their self-interest would be a recipe for Darwinian extinction or, as Badaracco puts it, "martyrdom and self-immolation" (52). A moral requirement that they do so would be ultimately self-defeating. There needs to be a more practical alternative (Provis 2000b, 10-11).

This is a challenging area, but experienced negotiators have found a practical way to meet this challenge – at least to a significant degree. They tend to avoid statements about their reservation price, including references to 'best' or 'final' offers. Instead, they seek to establish a consistent pattern of behaviour that means that a refusal to engage in such statements does not signal one thing or the other. They focus on the offer they are prepared to put forward at any given moment, reflecting the other terms and conditions and the knowledge then in their possession. They tend to say less rather than more, to focus on the facts rather than the rhetoric. They certainly do not volunteer the 'whole truth'. Instead, they withhold any comment on their reservation price and simply engage in a series of offers and counteroffers, none of which constitutes lying (Boatright 1992, 72; Carson 2005, 402). This approach is not dissimilar to that Shell (2018 [1999], 185-186, 189) advocates in his book on negotiating strategy, *Bargaining for Advantage*.

How might this play out in our opening scenario? How might Zach negotiate effectively without lying? Well, one could imagine a conversation that goes something like this:

Abe: “Zach… we would be looking for a discount. Only fair to tell you.”

Zach: “Abe, I know where you’re coming from, but this is the most advanced machine of its generation, and we are expecting a lot of interest. You know our price.”

Abe: “Yes, but is that your best price?”

Zach: “Well, you’re negotiating. That’s fair enough. But you know me, by now, and I never discuss ‘best’ or ‘final’ offers. That just makes for nonsense and bluster. You know the price we’re asking. And this is an excellent machine. I believe it can make a real difference on your production line. It’s worth paying up for.”

Abe: “Yes, but how much do I need to pay up?”

Zach: “Well, how much are you prepared to pay up? [Zach slaps Abe on the arm.] Come on, let’s discuss this over a beer.”

Zach has not lied; nor has he undermined his negotiating position. Yes, he has signalled *some* preparedness to compromise. But Abe almost certainly knew he would get *some* discount and Zach has really given nothing else away. Zach has also managed to put Abe into the position of being obliged to make the next move, a key strategy favoured by honed negotiators. Moreover, Zach has reinforced his personal standing and integrity with Abe, which may be significant in future dealings. By contrast, in the original scenario, Zach makes a categorical statement that his boss will not permit a discount, only an hour later to reveal by his own action in granting a discount, that he was previously lying to Abe. Such behaviour, repeated over time, will lead Abe to discount much of what Zach says, even when he is telling the truth.

The preceding case illustrates an important point: lying might confer a short-term advantage, but it also incurs long-term *costs* for the liar that include loss of trust, retribution, and diminished reputation (Shapiro & Bies 1994, 18, 29-31; Friedman & Shapiro 1995, 251; O’Connor & Carnevale 1997, 513-514; Boles, Croson & Murnighan 2000, 254-255; Croson, Boles & Murnighan 2003, 157-158; Tyler, Feldman & Reichert 2006, 75; Boero et al. 2009, 876; Rogers et al. 2017, 465-466; Lewicki, Barry & Saunders 2020, 190). Maurice Schweitzer, John Hershey, and Eric Bradlow (2006) have researched the impact of lying on trust in particular. They find that while trust can generally be restored following a period of untrustworthy behaviour, "trust harmed by deception never fully recovers" (15).

More broadly, Joseph Reitz, James Wall, and Mary Sue Love (1998, 13) enumerate four significant costs in lying: (1) "rigidity in future negotiations" – the inability to alter one's offer, having initially adopted a dishonest position; (2) "a damaged relationship with the opponent" once falsity is uncovered; (3) "a sullied reputation" in the wider community; and, in their view most important of all, (4) "lost opportunities" – the failure to identify mutually beneficial trade-offs, because of false signalling around the value attached to different elements being negotiated.

This fourth component, the ability through open and honest debate, to find mutually beneficial outcomes, has received particular attention. Such mutual gains bargaining (MGB), as it has been termed, was first popularised by Roger Fisher and William Ury in 1981 (Fisher, Ury& Patton 2011) and has since become the subject of an extensive literature, with a notable focus on labour negotiations. Beyond simply increasing the resource to be divided, Dean Pruitt (1983) identifies four strategies for delivering MGB – or "integrative agreements" as he prefers to call them. The detail of these strategies is unimportant here. The crucial thing common to them all is that they rely on an open exchange of information, for example "information about the nature of the two parties' interests and their priorities among these interests" (41). It is exactly this sort of information exchange, and the consequent value created, that the liar risks forgoing.

The apologist's appeal to practical necessity rests on an erroneous assumption that those who fare best in negotiations are what Andrea Schneider (2000, 24-28) labels "'Rambo' negotiators." In her study of that perhaps most adversarial of all professions, the law, Schneider found the exact opposite was the case. In her survey of 690 lawyers in Milwaukee and Chicago, 53% of "adversarial" (i.e., 'Rambo') lawyers were rated by their peers as "ineffective," while the equivalent for "problem solving" (i.e., non-'Rambo') lawyers was under 4%. Rebecca Hollander-Blumoff and Tom Tyler (2008, 493) draw a similar conclusion from their own empirical studies of "the 'lawyer as shark'." Aggressive negotiating tactics, including lying, are not necessarily outcome-maximizing, even for the aggressor (Burr 2001, 10; Cohen 2002, 118; Cohen, Helzer & Creo 2022, 21).

A similar conclusion emerges from a simulated market entry exercise conducted by Rebecca Guidice, G. Stoney Alder, and Steven Phelan (2009) among 112 undergraduate and MBA students. In successive rounds, students declared their intent (or not) to enter the market and then acted accordingly (or not). Points were awarded for entering an under-supplied market and deducted for entering an over-supplied one. In effect, the students were negotiating with each other for the right to participate in the market. A declaration of intent that was different from the subsequent action was rated a "bluff" (in this case, a knowingly false statement intended to deceive, i.e., a lie). Guidice, Alder and Phelan found that such "bluffing had a significant negative effect on performance" and concluded that "while it is often assumed that bluffing is advantageous, it actually is not" (547).

It is widely believed that it is the more experienced, 'street-wise' negotiators who are most inclined to lie, but again this is the inverse of the truth. Attitudinal studies, across several cultures, have reported a greater *reluctance* to lie among older negotiators (Volkema 2004, 75; Morse & Cohen 2019, 15). Deepak Malhotra and Max Bazerman (2007, 213) too have found that their more seasoned executive students are much *less* likely to condone lying in negotiations than their younger MBA students. They attribute this to the fact that their executive students "have many more years of negotiation experience" and "have witnessed and experienced the consequences of negotiating in bad faith." Asked about lying, Malhotra and Bazerman's advice is crystal clear: "Don't do it. Don't ever lie" (Malhotra & Bazerman 2007, 218; see also Nyerges 1987, 21, 24; Shell 2018 [1999], 194).

However, to conclude that lying is invariably to the disadvantage of the liar – that 'honesty always pays' – would be unduly sanguine. There is evidence that lying can confer at least short-term advantage (O'Connor & Carnevale 1997, 504; Aquino 1998, 208; Croson, Boles & Murnighan 2003, 156-157; Gaspar & Schweitzer 2013, 161; Rogers et al. 2017, 464-465; Gunia & Levine 2019, 134; Gaspar, Methasani & Schweitzer 2019, 64, 74; Meyer & Choo 2023, 8). It often goes undetected (Ekman & O'Sullivan 1991, 916; Bond & DePaulo 2006, 222, 230; Adler 2007, 70). And there are clear cases of negotiation that are inherently 'single-shot', such as the purchase or sale of a used car or a transaction in a tourist bazaar. In these cases, we must accept that lying may give some partial benefit to the liar. But this does not characterise most business negotiations. In the main, business negotiations are held between parties in the same or related industries or communities. Such negotiations are not by their nature 'single-shot', but instead form part of an extended relationship or at least involve parties who know of each other through reputation (Gunia 2019, 9). Here the long-term costs of lying are much greater.

In summary, there are methods available to protect one's negotiating position, while refraining from lying. And, lying does not come 'free'; it carries costs of its own, especially in the longer-term. Experienced negotiators, like the economist and merchant banker John Rutledge (1996, 78), know that, far from being systematically disadvantaged, "Nice guys often finish first.” And the inverse of that holds too. As another financier, the Goldman Sachs ‘whistle-blower’, Greg Smith (2012), declared in the *New York Times*, “If clients don’t trust you they will eventually stop doing business with you. It doesn’t matter how smart you are.”

But the apologist may raise a further objection. The practical strategy, outlined above, could lead to the accusation of inconsistency. The apologist might argue that it relies on an untenable moral distinction between active lying and passive withholding of information. Why, she may ask, is passively withholding information from the other side morally any different from actively lying to them? Both can prove seriously misleading. And, clearly, the consequences for the counterparty in each case can be equally damaging.

This is a complex area. Some radical consequentialists, like Peter Singer (1972, 229-243; 2009, 15-16), are reluctant to accept that there is any morally significant difference between doing (acts of commission) and merely allowing (acts of omission). And yet, if we examine the fundamental objections to lying, whether they be Kantian, consequentialist, or aretaic, it does seem reasonable to draw such a distinction. From a Kantian perspective, lying uses others as a mere means, while withholding information only represents a refusal to be useful to others, that is, a refusal to subordinate our ends to their ends – something that is arguably supererogatory. From a consequentialist perspective, there is a significant difference in the outcome for the party misled. Withholding information leaves that party not knowing the truth; but at least knowing that they do not know. By contrast lying may leave them not knowing the truth; and not knowing that they do not know. Furthermore, withholding information, except in the most egregious cases, does not engender the same broad, consequential mistrust as does lying. It does not fundamentally threaten social cooperation and cohesion. Indeed, it seems that some degree of withholding of information is essential to that social cohesion (Sullivan 1984, 11; Sunstein 2021, 33). For example, it is hard to imagine a functioning society where people volunteered *all* their thoughts about each other *all* the time – from sneering contempt to unbridled lust. The film *Liar, Liar* provides a humorous illustration of just that (Shadyac 1997). Virtue ethicists in their turn might argue that, whereas a refusal to lie is part of the virtue of honesty, willingly divulging one’s reservation price would betray an excessively open character, bordering on naivety. In Aristotelian terms, a simple refusal to lie, but to say no more, might reasonably be seen to sit at the ‘golden mean’ between opposing vices of manipulative lying on the one hand (a deficiency of truthfulness), and a mindless and naive refusal to withhold information, which one is not obliged to disclose, on the other (an excess of truthfulness).

Finally, this ethical distinction between lying and simply withholding information draws some support from the law. This takes the form of the right to remain silent. As Chris Provis (2000a, 155) points out, "If questioned by the police, I may not lie to them, but I have the right to remain silent, and fundamentally the same distinction is available to negotiators."

In summary, passively withholding information is morally different from actively lying. Not lying, but at the same time refusing to reveal, for example, one's reservation price, *is* therefore a morally viable option. The apologist's argument from practicality is neutralised because there are means available to the ethical negotiator that allow her to protect her position. And there are many other practical benefits that accrue from refusing to lie.

3.4 CONCLUSION

My claim is that lying in business negotiations is pro tanto wrong and no less wrong than lying in other contexts; and, furthermore, that such lying is both unnecessary and, in most cases, ultimately to the disadvantage of the negotiator. I have focussed on business negotiations but suggested, at various points, that this may extend to other professions as well, including the law, trade unionism and even government.

My case is simply that lying in business or professional negotiations is no different from any other form of lying and is therefore equally wrong. This is a reasonable presumption in the absence of any sustainable argument that there is a moral difference. My own argument has been largely one of refutation. The burden of proof rests with the apologist, who seeks to create a special dispensation for lying in business negotiations. The apologist must show there is a disanalogy, a moral difference that distinguishes business negotiations from other social contexts. Failing that, consistency requires that they be treated alike.

I have analysed the five main arguments put forward by apologists in defence of this moral difference: arguments based on consent, self-defence, the 'greater good', fiduciary duty, and practicality. I have sought to demonstrate that each is seriously flawed. The argument from consent looks unsustainable, if that consent is made explicit, and anyway fails empirically, because such consent is insufficiently universal. The more restrictive argument from self-defence is seldom, if ever, applicable, as even its most vocal proponent accepts; while the more permissive argument from self-defence would appear to justify all sorts of pre-emptive, unprovoked attacks that are clearly morally impermissible. And the argument based on the promotion of the 'greater good', while it does support the importance of the negotiating process, cannot justify lying, because lying, far from being an essential element of that negotiating process, actually threatens to undermine it and society at large. Meanwhile, the argument from fiduciary duty fails too, because it cannot erase the wrong inflicted on the victim; because it by extension permits any number of other wrongs that are clearly impermissible; and because it fundamentally misinterprets the relationship between principal and agent. Finally, the argument from practicality fails to address the moral principle involved; and anyway, is circumvented, because there *are* practical means available for protecting a party's negotiating position while avoiding lying. Furthermore, and contrary to the apologist position, the weight of evidence suggests that not lying is, in most cases, the more advantageous strategy. Far from being a burdensome constraint, the imperative not to lie is generally in the self-interest – especially the long-term self-interest – of the negotiator.

I conclude that my claim stands. It is time for those involved in business negotiations to recognise that a lie is a lie. The moral rules that apply to business negotiators are the same moral rules that apply in society at large. It was Carr's article of 1968, above all, that catalysed this whole debate. So, it is fitting to end with a reply to that article from the wife of a businessman, one Mrs Philip D. Ryan of Wyckoff, New Jersey: "Plainly the true meaning of a man's job escapes Mr. Carr. A man's work is not a card game; it is the sum of his self-expression, his life's effort, his mark upon the world, his pride. Men who… buy their business at the cost of their integrity… had best stick to card games and send their wives out to work" (Blodgett 1968, 166).

The focus of this chapter has been on lying. Space does not allow for a similar analysis of all the other forms of dishonesty, like fraud and corruption. However, I do believe that the same fundamental principles apply. In particular, there is no special exemption that prohibits such dishonesty in normal social exchange and yet permits it in a business or negotiating context. Accordingly, the first requirement that must be satisfied to ensure a fair negotiating process and resulting outcome is that the negotiating parties observe the basic demands of honest behaviour that are expected of society as a whole.

Chapter 4. Win/Win Negotiating: The Search for Mutual Benefit

ABSTRACT: A second constraint, required for a fair process, addresses our attitude to the division of value in a negotiated transaction. People at times reach agreements that are not in their own interest, whether due to immature or impaired reasoning, misinformation, blameless or culpable error, conflict of interest, excessive benevolence, or extreme submissiveness. This raises a serious ethical question for negotiators: where these non-ideal conditions occur – that is, where there are such departures from ideally rational and self-interested negotiation – what, if any, moral obligation is there to consider the interests of a potentially disadvantaged counterparty? I argue for an obligation of mutual benefit. I offer three related arguments for this based on moral imperatives (1) to share the rewards of cooperation (2) to recognise just claims in respect of property and person, and (3) to avoid doing harm to others. This requires a sharing among the parties of the expected cooperative surplus. Hence, each party to the negotiation is morally obligated to seek a win/win – that is, a mutually beneficial – agreement, whereby all parties enjoy an outcome that is equal to or better than that of the disagreement position. The disagreement position reflects not merely the status quo, but also any fully credible threats faced by, or opportunities open to, the parties. These include the reasonable expectation of being able to conclude an alternative transaction with another party. However, there can be no specific obligation to 'match the market' because negotiation is a tool for reaching agreement in exactly those circumstances where a clear market price does not exist. Instead, the obligation of mutual benefit is best seen, in broader terms, as requiring a certain attitude. Such a win/win attitude demands a sensitivity to the interests of the other party, a concern for the outcome viewed from their perspective and a commitment to creating value that can be shared by all participants in a transaction. Thus, a moral principle, usually considered supererogatory, proves instead to be obligatory, and represents the second constraint required to ensure a fair negotiating process and resulting outcome.

Having examined in the previous chapter a requirement that the negotiating parties observe the basic demands of honest behaviour expected of society as a whole, I turn in this chapter to a second constraint on a fair negotiating process: that the parties seek their *mutual benefit*.

Contrary to ideal theory, people at times reach agreements that are not in their own interest. This can be for many reasons. Some have only partially developed or impaired rational faculties. Others are uninformed or misinformed or simply get things wrong – whether due to negligence, inexperience, or blameless error. And still others are unduly influenced by conflicts of interest or the opposite, an excessive sense of benevolence or submissiveness that subjugates their own interests to those of others. This raises a serious ethical question for negotiators: where such non-ideal conditions occur – that is, where there are such departures from ideally rational and self-interested negotiation – what, if any, moral obligation is there to consider the interests of a potentially disadvantaged counterparty?

I propose an obligation to seek a win/win outcome, but what exactly does that entail? Having examined and rejected demands for impartiality and reciprocity, I suggest that the ethics of negotiation are governed by an obligation of mutual benefit. Negotiators are morally bound to seek an outcome whereby all parties are left equally well or better off as a result of agreement – that is, no worse off than in the case of a failure to agree (the 'disagreement position'). I argue for such an obligation on three grounds: moral imperatives (1) to share the rewards of cooperation (2) to recognise just claims in respect of property and person, and (3) to avoid doing harm to others. I also highlight parallels with recent work done on the implicit morality of markets.

I go on to explore the extent of this obligation of mutual benefit. On first appraisal it appears to impose only a modest burden on the negotiator. The negotiator is at liberty to optimise their own position with just one side constraint, that they seek an outcome whereby their counterparty is left no worse off than in the disagreement position. However, examination of the disagreement position and its implications suggests a more onerous obligation: to 'match the market'. Ultimately though I show such a tight prescription to be inappropriate. Negotiation and a market, of the kind presupposed by a 'mark to market' prescription, are mutually exclusive: where there is such a market to match, there is no room for negotiation.

I conclude that the obligation of mutual benefit is best seen, in broader terms, as requiring a certain attitude. Such a win/win attitude demands a sensitivity to the interests of the other party, a concern for the outcome viewed from their perspective and a commitment to creating value that can be shared by all participants in a transaction. A moral principle, usually considered supererogatory, proves instead to be obligatory, and represents the second constraint required to ensure a fair negotiating process and resulting outcome.

It is perhaps a sad comment on the state of the world at that time that the negotiating tactics prevalent in the 1970s and early 1980s are most associated with terrorist hostage crises, shuttle diplomacy between warring nations, and often violent disputes over labour relations. With regard to the last of these, the period was epitomised by what Frederick Post (1990, 495) calls variously the "adversarial collective bargaining process" or "a combat model based on power." Underlying this model was an assumption that such negotiations constituted a zero-sum game, with a fixed and finite pool of value, meaning that one side's gain would always be balanced by the other side's loss. Management (as the agents of the owners - in many instances, governments) were pitted against trade unions (as the agents of the employees) in a win/lose contest.

As early as 1960, Thomas Schelling (1980 [1960], 4) pointed out that such "pure conflict, in which the interests of two antagonists are completely opposed, is a special case," that is, not representative of most negotiations; but it was two books in the early 1980s that were particularly influential in challenging the adversarial stereotype. R. Edward Freeman's *Strategic Management* (2010 [1984]) offered a new model he called "the stakeholder approach" (27). This confronted the conventional logic of those who, like Milton Friedman (1962 & 1970), advocated the primacy of shareholder interests in the management of the corporation. As noted in the previous chapter (3.3.4), Freeman argued that a wider set of interests should be taken into account, including those of customers, suppliers, employees and the community at large. This thesis has not gone unchallenged. For example, a concern about conflicting objectives – to generate profits *and* safeguard employment *and* protect the environment – has led Joseph Heath and Wayne Norman (Heath 2014 [2004], 61) to object that lack of performance on one dimension can always be blamed on the need to meet objectives along another dimension. As they put it, "having three bottom lines is equivalent to having no bottom line." Nonetheless, the wide acknowledgement of the stakeholder approach has spurred the belief that good business management is not a zero-sum game, but one where all parties can and should gain from a transaction.

In the other notable book of the period, the popular bestseller *Getting to Yes* (2011 [1981]), Roger Fisher and William Ury, advocated what they then called "*principled negotiation"* (xxvi) and what has subsequently become better known as mutual gains bargaining (see also 3.3.5). They note, in the Preface to the Third Edition, that the past generation commonly asked of a negotiation, who is the winner and who the loser. The authors add that, "the idea, that both sides could benefit, that both sides could 'win', was foreign to many of us" (Fisher, Ury & Patton 2011 [1981], x). Their homely illustration of the unrealised potential for mutual gains bargaining concerns two sisters negotiating over the division of an orange. In the end they agree to cut it in half. One then eats her share of the fruit, discarding the peel, while the other uses her share of the peel to bake a cake, discarding the fruit. Clearly a missed opportunity to optimise mutual benefit (58-59, 75; see also Pruitt 1983, 35).

Each of these books has spawned a literature of its own. Much of it though is aimed at persuading business academics and practitioners that seeking positive outcomes for counterparties, whether they be customers, suppliers, employees, or the wider community, is actually in the interest – the rational self-interest – of the party themselves. In the continuing debate over the obligation of management, this has been termed the *strategic* or *instrumental* stakeholder approach (Goodpaster 1991, 57-58). Management is encouraged to promote, where possible, the interests of customers, suppliers, and employees because this is the most effective way of maximising the return to shareholders. Similarly, advocates of mutual gains bargaining pursue cost-effective ways of improving the outcome for counterparties in the expectation that this will, more often than not, prove an equally effective route to maximising gains for the parties themselves, the respective parties being in this case typically management and employees (Bowie 1985, 284; Post 1990, 499; Hunter & McKersie 1992, 319-320). According, then, to both these theories – instrumental stakeholder theory and mutual gains bargaining – a win/win outcome is to be pursued by each and all simply because it delivers expected value maximisation to each and all. A case of "good ethics is good business" (Bowie 2017, 171).

The distinction between prudence and morality remains nevertheless a significant one. It is of course reassuring to know that there is often such a congruence between what might commonly be thought 'good behaviour' and rational self-interest. However, I want to look beyond this prudential motivation and to ask whether there is also a *moral* obligation to seek outcomes that are beneficial for all, independent of whether doing so is prudentially wise as well. This is important, because there are instances where it is prudentially optimal *not* to seek a win/win outcome, and where, as a result, we must rely on moral reasoning to deliver a result that is fair to all.

Game theorists, in particular, seem to take it as an axiomatic condition of any admissible solution to a negotiation that each party's expected utility must equal or exceed that which would accrue to them in the absence of agreement. As we saw earlier (2.1), in their pioneering analysis of two-person cooperative games, John von Neumann and Oskar Morgenstern (1944) argued that the set of acceptable cooperative solutions was a function of just this condition, together with that of Pareto optimality (Luce & Raiffa 1957, 118-119). Six years later John Nash (1950) similarly introduced his solution to a "fair bargain" by stipulating that, "we shall consider only those cases in which there is a possibility that both individuals could gain from the situation" (158). Most game theorists since, from Richard Braithwaite (1954, 21, 26) John Harsanyi (1956, 144-145), Ehud Kalai and Meir Smorodinsky (1975, 514) on, have again adopted this same condition for a fair solution. And this is equally evident in the literature on fair division (Heilmann & Wintein 2017, 3490). Of course, this largely reflects these writers' assumption of universal rationality. They start from an assumed precondition of mutual benefit and only then ask what other conditions are needed to deliver fairness. As Johanna Thoma (2015, 3339) puts it, "we take for granted that agents will never agree on a set of utilities that affords either of them a utility smaller than that of the disagreement position."

This game theoretical position is reflected too in economic theory. As Luigino Bruni and Robert Sugden (2008, 48) observe, "it is fundamental to market exchange, as represented in economic theory, that each transaction, considered in isolation, provides benefits to everyone." This follows from the assumed voluntariness, together with "the presumption that each individual acts in his own interests."

Of course, this assumption of universal rationality rests on an ideal view of the world. In this ideal world, actors have perfect and symmetric information, use that information to make error-free calculations of their own best interest, and pursue that self-interest with the singular focus of the economist's *homo economicus*. This holds true too for more recent models of rational man, like Michael Moehler's (2018, 95-107) more risk averse, Hobbesian *homo prudens*. In the colourful imagery of Richard Thaler and Cass Sunstein (2008, 6-7), these various actors "think like Albert Einstein, store as much memory as IBM's Big Blue, and exercise the willpower of Mahatma Gandhi." In other words, they never make mistakes, they are never misled, and they never fail to pursue their own self-interest.

This creates something of a lacuna when we seek to address our real, very non-ideal world, peopled not by *homo economicus* or *homo prudens*, but, as Thaler and Sunstein suggest, by error-prone *homo sapiens*. It is not that the game theorists' assumptions are never justified. It is just that, by making the assumptions they do, game theorists fail to address the situation where their ideal conditions do not hold. From a game-theoretic point of view, it makes little sense to ask whether a party to a negotiation need be concerned with the outcome for their counterparty, and in particular with whether that outcome offers a positive expected utility compared to that of the disagreement position. It must do so, a game theorist will reason, otherwise the counterparty would never agree to transact (Sugden 2018, 44). From such a perspective, any moral requirement to look out for the interests of the counterparty appears superfluous.

And yet in our rather less tidy, non-ideal world, people do act against their own interest, that being the result of any number of things, ranging from simple error, inaccurate and asymmetric information, negligence, and internal conflict of interest to an excessively self-sacrificial sense of benevolence or submissiveness towards others. There *are* negotiations that leave one or more parties in a worse expected position than they would have been without a deal. This raises an obvious question: to what extent does a fair negotiation require each party to consider the expected utility/value, not only for themselves, but also for their counterparty? In particular, is there a moral obligation to seek a win/win solution for all parties to an agreement and, if there is, exactly what requirements does that impose?

4.3 THE NON-IDEAL WORLD

Let us explore further this non-ideal world, where parties to a negotiation do at times act against their own interest. To explain exactly what it means to act against one's own interest is not straightforward, and I will defer attempting that task until a little later. But there clearly are instances where negotiators make "ditzy blunders" (Thaler & Sunstein 2008, 79) or otherwise compromise what would seem to be in their best interest. The potential causes of this are various.

First, there are cases where parties are less than fully rational, either because their rational faculties are not yet fully developed (children) or because those faculties have been constrained in some way (impaired reasoning). Typically, we do not hold such parties liable for the agreements they make. We believe that they cannot be relied upon to have properly determined what is in their own interest. Instead, they exhibit what Alistair MacLeod (1984, 711) terms "blameless stupidity".

Second, there are situations where parties lack the correct information necessary to negotiate fully in their own interest; or alternatively have that information but nonetheless act in error. People make mistakes, either because they are uninformed or misinformed, or because they simply get things wrong (Sunstein & Thaler 2005, 177). This may be through their own carelessness or negligence – what MacLeod (1984, 711) calls "culpable stupidity" – or it may be because of unavoidable circumstances that create opacity and/or limit the time for analysis and reflection. As Thaler and Sunstein (2008, 256) put it, "when things get complicated, Humans can start to flounder."

In some cases, such ignorance may be a perfectly rational choice. A familiar example would be many contracts of adherence, where the fine print laying out the terms and conditions is so extensive, and the transaction value sufficiently small, that it makes no sense for the purchaser to do anything other than 'tick the box'. Colleen McCullough (2016, 801-802) calls this "the problem of rational ignorance," pointing to empirical evidence that "only one or two in every thousand consumers even glance at adhesive contracts in online transactions." Ignorance and error come then in many forms, both rational and irrational.

A third, less obvious, case arises where there is a conflict of interest within the party themselves, usually between an agent and a principal. Here the agent agrees a transaction that is in *their* own interest but to the detriment of their principal. The interest of the principal (the ultimate party to the agreement) is compromised because consent to the agreement does not come directly from the principal themselves but rather from the conflicted agent. Or, if it does come directly from the principal, it does so on the guidance of the conflicted agent. This is a common situation in many business transactions, where middle and senior management, chief executives and board directors are typically negotiating *on behalf of* their employers and shareholders. The directors' and management's interests are rarely, if ever, identical to those of their shareholders and other stakeholders.

Examples of this abound, especially where ownership is highly fragmented. Such fragmentation leaves shareholders with limited leverage and accentuates what Alexander Pepper (2022, 27) calls "the investors' collective action problem." Research by Lucian Bebchuk and Jesse Fried (2003) highlights the resulting agency problem: the negotiation over salaries and bonus arrangements in public companies is often perceived as a form of self-dealing, whereby value is transferred from shareholders to rent-seeking senior management. Board directors, who are in theory there (as agents) to defend the interests of the company and its shareholders (the principals), are often felt to lack independence; and to be overly deferential towards the chief executive, frequently relying on that same chief executive for their own subsequent re-election. In such circumstances, board directors might well feel reticent about negotiating forcefully with an influential, incumbent chief executive (73-75). Inevitably there will be instances when "directors do not adequately represent shareholders' interests in bargaining with managers over their pay" (88). This "fundamental agency problem" (72) is particularly acute in the case of negotiations over mergers and acquisitions, where there are regular accusations of managerial empire building to the detriment of shareholder value; and where all too often the key 'sticking point' in such negotiations is the future roles and remuneration of senior management, especially those managers leading the negotiations (89).

Public shareholders rarely vote down these arrangements. As in the case of contracts of adherence, this 'hands off' behaviour on their part may be entirely rational. As Pepper (2022, 27) points out, it may make no sense for an investment firm with £50 billion in assets to concern itself with the governance arrangements in a public company, where it holds only a 1 per cent shareholding.

But such conflicts are not limited to public companies. In the early days of the leveraged buyout industry, it was standard practice for an acquiring private equity firm to negotiate a fee for themselves on any new transaction – a fee that was determined by the private equity firm (the agent) but, in effect, paid for by the investors the firm represented (the principals). Many investors viewed this as yet another form of self-dealing. It was many years before collective pressure from investors grew strong enough to eliminate the practice and force the private equity firms into a closer alignment of interest. These are contentious issues, but in all these instances the concern is the same: that the principals – whether public shareholders or private equity investors – are, via their agents, party to a negotiated outcome that is not necessarily in their interest.

Fourth, there is the complex area of benevolent or 'other regarding' preferences, where one party places such importance on the welfare of their counterparty that they are happy to agree a deal that, while highly beneficial to that counterparty, on the face of it, is disadvantageous for themselves. There are difficulties here. First, there is of course a question as to whether such a party is really behaving against their own interest. It might reasonably be argued that, if their utility function is such that the party ascribes great value to the wellbeing of their counterparty, then they too are a beneficiary of any transaction that benefits that counterparty. How comfortable we should feel with this rationale depends, I think, on an important distinction between two different forms of benevolence.

On the one hand, there is the benevolence based on family ties and/or deep affection – what Amartya Sen calls "sympathy" (1977, 326-328). In this case, the benevolent party may well be making a self-denying offer in favour of a close relative or lifelong friend. But such exchanges are rarely characterised as negotiations. They are perhaps best described as accommodations, collaborations, or shared endeavours. They fall in a different category of exchange, closer to the spirit of a gift or to the ethos of the camping trip, described by G.A. Cohen (2009, 3-11), where a group of friends work together in pursuit of a shared goal.

The other type of benevolence is more relevant to the current analysis and more perplexing. This is the attitude of general selflessness, public spiritedness or just conformity to an expected norm – what Sen (1977, 326-330) calls "commitment," and what Philip Pettit (1995, 310), following Philip Wickstead, contrasts with "strong non-tuism." Cristina Bicchieri (2010, 8-28) ascribes such benevolence to "social norms," a desire to conform to established norms and meet the expectations of others. Parties of this second kind are inclined to place unusual importance on the welfare, preferences, and expectations of others, for whom they may carry no particular affection, but to whom they feel a generic obligation. In some cases, this perceived obligation may reflect a desire to conform to a social convention, or to observe a perceived moral or religious norm (see Elsa Kugelberg (2021, 376) on how widespread "harmful gender norms" in particular constrain women's choices). In other cases, it may reflect an extreme form of normative altruism that subordinates self-interest to a point bordering on submissiveness (see below). It is hard to say that these sentiments are in all cases wrong. But we might feel, along with J.R Lucas (1959, 10), that this "selflessness can be carried too far," and sometimes to such a degree that the party's own true interests are compromised.

This focus on expected behaviour and adherence to norms raises a further related concern: what Martha Nussbaum (2011), following Sen and John Elster, calls "adaptive preferences" (54, 83). This is where people for whom certain goods and benefits have never been available, cease or never learn to aspire to having those goods, or to enjoying those benefits. Nussbaum cites the case of a Gujarati woman, Vasanti, who reports no dissatisfaction with her illiteracy or lack of access to political life, because she has never been led to expect anything better (55). Women like Vasanti have "deeply internalized the idea that a proper woman does not go in for schooling" (83-84). When negotiating over family roles, such women are very likely to accept that there will be no educational provision either for themselves or for their daughters. And they are very likely to accept this *willingly* because they have no expectation to the contrary. But equally we may feel that such an outcome is not in their interest, only serving to perpetuate their disempowerment (153-154).

The above cases of other regarding and adaptive preferences point towards a final, and particularly troubling, cause of negotiating parties striking deals that are not in their own best interest: submissiveness. This may reflect adherence to a norm of hierarchical subordination, or it may express a simple feeling of powerlessness in the face of domination, or indeed both. Submissiveness of this kind is usually associated with ethnic minorities, the economically disadvantaged, the socially stigmatised, and women in highly patriarchal societies. But Kate Manne in her book, *Down Girl* (2019) [2018], suggests a still wider problem. She argues that even educated women, in apparently liberal and progressive democracies, tend to negotiate ineffectually with men, because the *"Hers to give*" and the "*His for the taking*" (130) norms of gender relations are so embedded. Women are socially conditioned to cater to men's desires without demanding fair reward. In this way, "misogyny makes people so irrational" (290) that "bad gendered bargains" (xiv) are inevitable.

In all the cases above it seems to me clear that the parties are acting against their own best interest, but it is not straightforward to explain exactly what is meant by that. What are the criteria by which we measure such an interest? Two broad possibilities present themselves, but neither is entirely satisfactory (Fletcher 2016, 151). First, we could use an objective benchmark of value – a measure perhaps of well-being that stands independently of the party's own immediate preferences. However, such an objective benchmark is notoriously difficult to establish. Even in transactions which seem to be straightforwardly limited to a single measure of value, such as money, there are often other factors present, for example, the certainty that a deal will be concluded.

Alternatively, we could use what might be termed a subjective counterfactual measure of value. This would be the subjective preference of the party – to use Thaler and Sunstein's (2008, 5) much quoted phrase, "*as judged by themselves*" – were they to act rationally under ideal conditions, that is, as we would expect them to act with fully developed rational faculties, with symmetric information to us, in the absence of any conflict of interest between agent and principal, and with the typical intent of a *homo economicus* or *homo prudens*. This requirement creates clear epistemic and potentially even existential challenges: Robert Sugden (2018, 13), for one, doubts whether such "'true', 'underlying' or 'latent' preferences" exist at all. Most though would accept that there are instances where it is reasonable to assume that a declared preference does not fully reflect a deeper set of preferences that would reveal themselves in conditions closer to the ideal, that is in the absence of such factors as misinformation and error.

The epistemic challenge, involved in revealing such latent preferences, is obvious though. William Jevons (1879 [1871], 15) famously insisted that "every mind is […] inscrutable to every other mind" and many others since have grappled with the same concern (Arrow 1951, 9; Harsanyi 1955, 317-318; Yaari & Bar-Hillel 1984, 8; Hausman 1995, 489).

And yet, this scepticism can be pushed too far. Humans are pretty good at being able to gauge the contents of other people's minds. And, enough so, that at times we can and do judge confidently that someone is making a mistake. Indeed, I would argue, along with Thomas Scanlon (1991, 44), for "the impossibility of doing without such judgments." People often do overlook things of great importance to them, threatening to badly set back their own interests. And other people do frequently recognise this and warn them of it. Without such an ability, the world would be sadly diminished.

But what are we seeking to achieve in correcting such a mistake? Are we correcting solely for cognitive error – that is, conducting what Douglas Bernheim (2021, 387) calls "implementation critiques" – or are we also challenging the validity of the other party's declared preferences? Let us take the case of a submissive wife, who desires to satisfy her husband's demands, however unreasonable they may be, not out of affection but out of a perceived obligation to respect an established norm of female subordination. This seems to be neither a cognitive error nor a conflict of interest. Instead, it is an, in our view, ill-advised preference. She wants to please/obey her husband. We just feel that the reason she wants to please/obey him – in this case, conformity to a norm – is not a good one. We are suggesting that her preference is morally flawed. To do so is clearly to open ourselves to accusations of paternalism and arrogant epistemic superiority. And yet, equally, for her husband to make unreasonable demands of his wife, relying on his wife's loyalty to this patriarchal norm, seems exploitative and unacceptable.

We face here a dilemma that is not easy to resolve. There are circumstances when to discourage a counterparty from entering a transaction is to frustrate their clearly stated wishes; and yet to complete the transaction is to be complicit in setting back their interests, in a manner that seems morally unacceptable. Many of the cases discussed above appear to call for an element of either explicit paternalism (for example, children and the mentally impaired) or implicit paternalism (for example, the excessively benevolent or submissive). Whether this paternalism is acceptable may to some extent depend on an 'all things considered' determination of the degree of potential harm done, if ignoring such a call.

Sugden (2018) adopts an alternative approach, but it ultimately faces this same dilemma. He argues that, when judging a party's best interests, the relevant criterion should be not their preferences – at least, not directly – but rather the set of opportunities open to them and, in particular, the extent to which that set of opportunities is expanded (84), since "the more opportunity a person has, the more effectively his preferences will be satisfied" (85). His claim is that you can only benefit, and never harm, someone by expanding the range of opportunities open to them – that is, for example, by offering them an opportunity to transact that they may choose to either accept or decline. There are, though, challenges to this approach. First, it is not clear that additional opportunities are always in a party's best interest. Such a "Just Maximise Choices mantra," as Thaler & Sunstein (2008, 183) characterise it, has been strongly contested in areas like financial planning (Sethi-Iyengar, Huberman & Jiang 2004, 88-92), prenuptial agreements (Wertheimer 1996, 302), sweatshop regulation (Kates 2015, 195-199), euthanasia (Velleman 2015, 5-20), organ donation (Rippon 2014, 146-148; Walsh 2014, 153-154; Koplin 2018, 310-313), and self-enslavement (Nussbaum 2011, 26). Moreover, expanding a counterparty's opportunity set does not eliminate the possibility of misinformation, miscalculation, conflict of interest, excessive benevolence, or extreme submissiveness. We are still left with the question of what obligation, if any, we owe in the event that our counterparty places great weight on an option – for example, a transaction – that we are convinced is not to their benefit. Having an additional option (e.g., the option to transact) may generally be of positive value, but actually exercising that option (i.e., transacting) is often not so.

This debate remains highly contentious but, as a practical matter, it seems to me that, when during a negotiation we do take the time to consider the interests of another party, we employ both of the broad approaches outlined above, imperfect though each is. On the one hand, we take into account what we consider to be objective measures of value; while on the other hand, we try to place ourselves in the shoes of the other party and reflect on what might be best for them given *their* expressed values, preferences and priorities. Of course, this is an empirical observation, and yet such a position appears also normatively reasonable. To my mind, a morally tenable combination of the above would adopt an objective basis of value, but with one of the inputs to that objective value being the way the counterparty feels about things. For example, a measure of objective value might reflect the importance of adequate educational provision, while recognising that particular perspectives on what type of education is appropriate may vary markedly across different counterparties.

It is worth noting that, while this task of establishing another party's true interest might seem daunting, it does not strictly require an interpersonal comparison of welfare. Consideration of whether a transaction is against a counterparty's true interest, requires only an assessment of *their* well-being with a transaction, compared to *their* well-being without a transaction. To make that judgment does require some understanding of their values but it does not require a strict indexing of their values to our own. As we will see later (4.5.3), we are not seeking to trade off the one (our well-being) against the other (their well-being). Nor are we looking to balance the two in any way. We need to know only whether the counterparty's well-being would be increased or decreased. And, as argued above, there are many instances when we can with confidence judge a counterparty to be so ill-equipped, so ill-informed, so compromised by internal conflict, or so mistaken in their value judgments, that they are agreeing to something that will materially set back their interests. In such circumstances, an element of paternalistic concern may be called for. Unsettling though that might be, the alternative may be yet more unacceptable: to exploit youthfulness, inexperience, ignorance, significant misunderstanding, conflict of interest, excessive benevolence, or extreme submissiveness.

None of this is perfect. But I'm not sure what more there is that can reasonably be asked of a negotiator. Hopefully what we have here is at least a healthy *attitude* of concern for our counterparty's best interests, however these are precisely defined. In the sections that follow (4.4 & 4.5) we will consider whether such a normative obligation exists and, if so, what is its nature.

4.4 THE EXISTENCE OF A MORAL OBLIGATION

Given this admittedly imprecise measure of value, do we then have a moral obligation to consider the interests of our counterparty in cases like those described above? I think it is clear that we do – at least in some instances. Let us take each of the cases in turn.

Where our counterparty is a child or suffers from impaired reasoning, we have an uncontentious obligation to look out for their interest or to ensure that another competent party is representing that interest (Brenkert 1998, 11-13). We are not at liberty to contract with them on whatever terms they are prepared to accept, without considering carefully whether such an agreement is truly protective of their interest. In the case of children, we will often say that they are below the age of consent. By this we mean that they are not yet capable alone of reaching agreements that sufficiently safeguard their own interests. Those who engage with them must recognise a special obligation to consider those interests, an obligation over and beyond what might be reasonably expected in other negotiations.

Cases of misinformation, lack of information, misjudgment and simple error are more complicated to assess, partly because the circumstances are so varied. It is rarely the case in any negotiation that all parties have access to symmetric, let alone perfect, information; or that they have equal capabilities in processing that information. Negotiators make mistakes. And often it is only reasonable that they bear the consequences of those mistakes. But there are instances where we know with a high degree of certainty that our counterparty has received misinformation or has significantly misinterpreted that information *and* that the resulting misjudgment is sufficiently prejudicial for the counterparty that the transaction would no longer be in their interest to pursue. Under such circumstances, knowingly taking advantage of the other party's ignorance or error seems wrong. It offends our moral instincts, especially those in the Kantian tradition. The counterparty, on eventually realising their mistake, will legitimately feel used. Their ends will clearly have been subordinated to our own, as "a mere means, a thing, a tool," to fulfilling our purpose, notwithstanding the loss that they suffer as a result (Korsgaard 1996, 141).

Furthermore, there is again a question of consent: as Derek Parfit (2011a, 195) says, "people who don't understand the most relevant facts […] cannot *validly* give or refuse consent." Tom Dougherty (2013, 179) has argued in the context of sex that not only coercion, but also deception, where known 'deal breakers' are concerned, invalidates consent. There appears a parallel in this case. There are instances where a failure to point out a clearly evident error to a counterparty, an error that is quite obviously a 'deal breaker', amounts to deception and invalidates consent. The counterparty is no longer a fully voluntary participant in the transaction.

MacLeod (1984, 713) offers a similar argument in support of this same conclusion. He suggests that, where this "adequate-information condition" is not met there can be "no assurance that the contracting parties are parties *to one and the same agreement*." In other words, the advantaged party cannot reasonably claim that the counterparty has in any voluntary sense signed up to the same deal. And Joel Feinberg (1984, 116) echoes this concern, insisting that there cannot be full consent where one party has "been 'taken advantage of' by a person who had information not remotely accessible to [them]" (see also Wertheimer 1996, 26). There is no real consent and thus any negotiation conducted on this basis cannot be seen as morally binding.

Contracts of adherence are again a case in point. If the terms and conditions, contained in the fine print, are egregiously onerous, we feel there is an obligation to point this out and, in many markets, for example insurance, this is a regulatory requirement. Typically, such terms must be communicated in bold print and/or stated verbally. Without such clarification, we may well doubt whether the counterparty has fully consented to the relevant terms of the transaction.

Similarly, where we are aware of a conflict of interest on our counterparty's side, we may not unreasonably exploit that for our own advantage. I say 'unreasonably' because there is inevitably a matter of degree involved. If we offer our counterparty's agent a pleasant meal over which to negotiate, that in most cases is acceptable; if our agreement with them includes a future consultancy contract for the agent, that is more troubling; and if it involves a cash transfer into a foreign bank account, that is outright bribery. In many cases, this is not easily remedied by disclosure alone. If the agent is a senior manager with authority to make an agreement on the principal's (i.e., her company's) behalf, there may be nobody to make a disclosure to beyond the agent herself. Her interest as agent is not identical with that of her principal. And yet she is not a detached adviser either. After all, she has the clear authority to commit the principal. In such a situation it seems misplaced to say there is a lack of consent, but it does feel as though that consent has been in some way compromised. There must be a moral limit to the extent that a negotiating party may take advantage of this disconnect. And, indeed, in most jurisdictions that moral limit is reflected too in a legal limit (e.g., the US Foreign Corrupt Practices Act 1977 and the UK Bribery Act 2010). But, leaving the law to one side, the key point is that, where a party is negotiating with a counterparty that is evidently subject to a material conflict of interest, the party is not morally entitled to exploit that in whatever manner it chooses. To the contrary, the party is morally obliged to consider the true interest of the counterparty's principal.

To return to an earlier example (see 4.3), such a conflict of interest is evident in the negotiation of chief executive compensation by board directors. In many cases, such negotiations are truly 'arm's length'. But, in others, there is such a level of mutual dependence that this cannot be the case: the chief executive's remuneration is determined by the directors, while the directors' re-election and remuneration are in turn heavily influenced by the chief executive. As Bebchuk and Fried (2003, 73-74) observe, frequently the director slate proposed by management is the only one offered to shareholders. Directors are thus very much beholden to management. As a result, it has been suggested that "many CEOs de facto set their own pay" (Betrand & Mullainathan 2001, 902). In such circumstances, Jeffrey Moriarty (2009, 235) argues that chief executives are not at liberty to exercise to the full their negotiating leverage over the board's directors. He is surely right that there is "a moral limit on how much compensation they can accept, and hence seek in negotiation, from their firms." To relate this back to the broader point, the agent (in the form of the directors) is conflicted and, as a result, the potentially advantaged party (in this case the chief executive) has a moral obligation to consider not only their own interest but *also that of their principal counterparty* (being the company as a whole).

Finally, there is the complex and deeply interwoven area comprising benevolence, adaptive preferences, and submissiveness. We might look on these sentiments as an expression of positive virtues: self-sacrifice, a willingness to support norms that are valued and respected by the community as a whole, a selfless preparedness to accept one's 'lot', and a sense of modesty regarding one's own position in the social structure. Or we might view them instead, with a sense of despair, as a source of disadvantage and domination. Whichever is the case, though, we can probably agree that exploiting them for our own advantage is wrong. Indeed, Feinberg (1990, 183), in his taxonomy of exploitation, considers the abuse of such "*moral virtues*" and "*conscientiousness*" to be "the clearest example of *unfair* advantage-taking."

Again, let us return to a previous example (see 4.3). In a patriarchal society, it is not morally permissible for a husband to ask of his wife whatever he seeks, however unreasonable, knowing that she will oblige out of an ingrained submissiveness, or because she has no expectation of anything better, or because of her desire to conform to a social, moral, or religious norm. Such behaviour is abusive. A sadly very real instance of this is husbands in parts of the developing world asking their wives to give up for sale their 'redundant' kidney. In many cases the wives are said to have grave doubts (Satz 2010, 196; Malmqvist 2014, 116; Malmqvist 2015, 116; Koplin 2018, 311). It is hard to be certain why they do agree. In some cases, it may be a simple act of submission; in others, a reluctant acknowledgment that this is expected of them; and in still others, a genuine desire to conform to an accepted norm. It seems wrong to say that they have not consented, but there is a sense in which they appear nonetheless to have been coerced, if only by convention. Whichever of these explanations holds true, taking negotiating advantage of such circumstances is not morally acceptable. As in the other cases, we cannot avoid some obligation to further consider the true interests of our counterparty.

It is difficult to say categorically, in any of the cases above, that the parties were forced against their will. And yet, in all these cases, there is a sense in which their consent remains incomplete, whether that be through lack of a fully developed rational faculty, epistemic failure, conflicting interests, pressure of convention or exceptional submissiveness. In most instances, such pressures are very slight, and we will feel that any agreement reached is nonetheless in the interest of everyone. But, in more extreme circumstances, it may be that the element of distortion created by these pressures is so acute as to generate an outcome that is no longer in the interest of those concerned. Such cases raise important ethical concerns.

We have reached then two, I think quite modest and yet significant, conclusions. First, when negotiating, parties do at times reach agreements that are not in their own interest. And second, in these cases, we are under a moral obligation to at least take such a counterparty's interests into consideration. But what exactly is the form of that obligation?

4.5 THE NATURE OF THE MORAL OBLIGATION

There are a number of different forms that such an obligation could take:

4.5.1 Impartiality

The most demanding form would presumably be a duty of impartiality, that is, the adoption of a disinterested perspective that treats our counterparty's interests on a par with our own. For Brian Barry (1995, 12) such a perspective reflects "a certain kind of neutrality" between competing viewpoints, that is central to his conception of morality. While for John Rawls (1971, 190) an impartial judgment is one that would be rendered in his original position from behind a veil of ignorance. It is what Thomas Nagel (1986) calls "The View from Nowhere": blind justice, without fear or favour. Such impartiality entails that whatever division we seek to agree with our counterparty is independent of who receives which share.

This perspective has its place of course. It recalls for us the image of Adam Smith's "impartial spectator" (2009 [1759], 133). And, as Smith argues, such impartiality represents the ideal of the "equitable judge," whose judgments must reflect the law without favouring the interest of one party over that of another. It is as though we were adjudicating between strangers (Parfit 2011a, 329).

But this has no place in a negotiation. This was discussed at length in Part 1 but is important enough to reiterate specifically in relation to this concept of impartiality. A suitably impartial perspective might be valued in a facilitator helping to organise a neutral process, but it is not realistic or reasonable to expect such an attitude on the part of those negotiating. For an impartial judgment we look to an adjudication or a binding arbitration, not a negotiation. If a judicial or quasi-judicial determination is available, cost-effective, and acceptable to all parties, then there is no role for negotiators. The essence of a negotiation is that it pits *directly* against each other *partial,* and to some extent conflicting, viewpoints as a route to finding a cooperative solution. A binding arbitration, or other kind of impartial judgment, may achieve the same ultimate goal but it is an entirely different kind of process. In short, a negotiation without partiality is not a negotiation at all.

It might seem that this dismissal of impartiality rests on a play of words, a matter of definition. But there is a substantive point at issue here. We have a procedure, whether we call it negotiating or bargaining, that is clearly distinct from a parallel and different procedure we typically refer to as judgment or binding arbitration (see Chapter 1). To conflate the two is a mistake because, while the product of each may be similar (an agreed end position i.e., a solution), the means of getting there (i.e., the process) is fundamentally different. To talk of an 'impartial negotiator' is to combine two incompatible concepts. The process of judgment or binding arbitration is morally grounded on the principle of impartiality. I suggest that the process of negotiating or bargaining is morally grounded on a different principle, that of mutual benefit.

4.5.2 Reciprocity

Before I turn to mutual benefit though, let me first address another form that, it might be argued, the obligation to consider the interests of our counterparty should take, what we might loosely call reciprocity.

One very specific version of this is central to David Gauthier's (1986) work. Gauthier argues in the contractarian tradition (9-10) for a morality closely tied to rational self-interest. For Gauthier, "to choose rationally, one must choose morally" (4) and he locates that rational morality in a form of proportional reciprocity that he terms "*minimax relative concession"* (137). Drawing on a model first sketched by Howard Raiffa (1953, 361-387) and subsequently axiomatized by Kalai and Smorodinsky (1975, 513-518), Gauthier (1986) argues that a fair bargain can be reached only when we equalise the concession required of each party from that party's best possible outcome, as a proportion of the delta between that best possible outcome and the disagreement position (14, 137). In other words, it requires a proportionately equal concession from each of the parties (140). It is broadly consistent with Lawrence Becker's (1980, 420) definition of reciprocity as "a fitting and proportional return of good for good;" and Barry has succinctly coined this "Splitting the Difference" (1989, 388-392).

Gauthier's theory has been much criticised due to its apparent exclusion of the unendowed, that is, those such as the severely disabled and children who might have nothing to contribute to the creation of the cooperative surplus (Buchanan 1990, 230-233; Gibbard 1991, 271-273; Barry 1995, 42-46). But this critique, while quite to the point with respect to much of Gauthier's broader project, fails to hit its mark when we limit our analysis strictly to what makes for a fair negotiation. A party with nothing at all to offer cannot expect to negotiate. They might rightly feel that justice requires that their interests should be respected; that they should, at the very least, not be harmed. Furthermore, they might look to the charitable nature of their counterparties, in the hope of generosity. But this cannot constitute a *negotiation*.

However, reciprocity, in this very prescriptive form, runs into trouble on other grounds. First, there is an issue of transparency. The importance Gauthier's theory places on each party's best possible result makes the 'fair' outcome vulnerable to manipulation, to each party claiming improbable, 'unreal' upsides in the calculation of their own position. To borrow from Duncan Luce and Howard Raiffa (1957, 133), "certain aspirations are merely empty dreams" and should not be allowed to determine a fair bargain (see also Moehler 2018, 55).

But, even allowing for perfect information, there is a still more serious challenge – a variant of what Cohen (1995, 184) calls the "*value allocation problem*": how to allocate value when the creation of any value at all is utterly dependent on the combined contribution of multiple parties? The difficulty is that, while a proportionately equal sharing of the cooperative surplus may appear reasonable in many cases, we rarely feel that it is the *only* outcome that can be considered fair. To take a somewhat extreme, yet oft rehearsed thought experiment (e.g., variants in Meyers 2004, 324-325; Elegido 2009, 33; Zwolinski, Ferguson & Wertheimer 2022, sec2.1.3, sec2.2.1):

> *Desert Hiker & Water Vendor*
>
> Let us imagine someone who sells a hiker, lost in the desert and near death with thirst, a bottle of water for $1,000. Intuitively this might feel wrong – an unfair transaction, based on discriminatory pricing and exploitative in the pejorative sense of the word. And yet, if the hiker values her life at say $1,000,000, then almost any application of 'Splitting the Difference', when applied to the cooperative surplus, will generate a fair value for the transaction much *higher* than this. In the case of Gauthier's minimax relative concession, if the most that the hiker has available to offer to the seller in cash and future earnings is $250,000, then an equal 20% concession by each side, from their maximum possible gain, would entail a price for the bottle of water of $200,000. Perhaps on some putative criteria this is a fair bargain, but such a prescriptive outcome seems highly counter-intuitive. It feels as though some other consideration must come into play beyond such a mechanical division of value. Moreover, it is important to note that this same problem will apply to almost any prescriptive principle of division when applied to the cooperative surplus. To take, for example, a simple equal sharing of the cooperative surplus: if the bottle of water is worth near nothing to the seller, but a lifesaving $1,000,000 to the hiker, then the 'fair' price would be $500,000. This cannot be a fair solution, let alone the *only* fair solution (for a parallel case, see Wertheimer's physician scenario (1996, 107)).

Allan Gibbard, drawing on the work of Schelling, has proposed a different, less specific, less formulaic, and less prescriptive version of reciprocity. He also recognises that reciprocity underpins many of our "moral sentiments" (Gibbard 1990, 262) but believes this reciprocity can be based on different divisions of the cooperative surplus that depend in turn on varying criteria – criteria the force of which is derived largely from their contextual salience, their "subjective prominence" (Gibbard 1991, 273). Schelling (1980 [1960]) calls this "the intrinsic magnetism of particular outcomes" (72) in "a lumpy, discrete world" (77). These salient features include various forms of symmetry, as well as such factors as "precedent", "custom" and "familiarity" (Gibbard 1990, 262-263). For Gibbard, there is "no unique standard of fairness" (262). Instead, "reciprocity needs terms of trade" (261, 264) and in many cases there is "slack" (264) in what those terms might be. In other words, the principle of reciprocity has a major influence over our moral judgment but does not of itself prescribe a specific outcome.

Is there, then, any interpretation of reciprocity that we might consider a valid obligation on negotiating parties in respect of their counterparties? I believe there may be if we take reciprocity in its most minimalist guise of 'you can't expect something for nothing'. This states that it is not reasonable for one party to expect benefits from a negotiated agreement without offering *something* of positive value in return to the other party, some minimal reciprocation. In this, its barest form, a necessary condition for reciprocity is mutual benefit.

4.5.3 Mutual Benefit

The obligation of mutual benefit requires that the various parties have a shared intent that each of them draw some anticipated advantage (or at least suffer no anticipated loss) from any mutually agreed transaction. That is, that each party achieve an equal or greater expected value than they would in the event of a failure to agree. In common parlance, that they seek a win/win outcome. This is arguably the bare minimum that could be asked of any negotiator. It permits the negotiator to optimise their own position to the full, with just a single side constraint: that their counterparty come away with no less value than they would get in the absence of an agreement. Importantly this side constraint only has teeth in non-ideal conditions. By contrast, in ideal conditions, the side constraint becomes redundant, because the counterparty will not agree to anything that is not in their own interest, that is, to any deal that would bring the side constraint into play. It is only in non-ideal conditions, when the counterparty behaves in a manner that compromises their own interests, that the side constraint does bite – and then, again, only to the extent that the counterparty might be left worse off than in the event of a failure to agree.

Before I attempt to defend this obligation of mutual benefit, let me offer three points of clarification. First, it is important to note that this obligation is a necessary, but not sufficient, condition for a fair negotiation. There are other obligations too such as basic honesty (not lying) (see Chapter 3) and avoiding the abuse of monopoly power (see Chapter 5). These requirements are independent of the obligation to seek a win/win outcome. For example, many of the more egregious outcomes in the Desert Hiker case above would breach a further constraint on the abuse of monopoly power.

Second, as emphasised previously (Chapter 1), like many moral strictures, the obligation of mutual benefit, while being substantive, must nonetheless be considered pro tanto. Clearly there may be circumstances, when potential externalities are so damaging that a party is required to agree to a negotiated transaction notwithstanding that it is expected to inflict a net loss on the counterparty. There may also be instances when the counterparty has committed sufficient wrong to deprive itself of any claim to mutual benefit. An obvious example would be negotiations with a terrorist or hostage-taker.

Finally, this can only be an obligation to *seek* a mutually beneficial outcome. Nobody can ensure that everyone benefits (or at least suffers no loss) from a negotiation. The best they can do is to seek that outcome. The normative component thus attaches to the intention. It is in this sense, in part, attitudinal. In this respect, the obligation of mutual benefit is little different from many other familiar moral imperatives. For example, most utilitarians are what Parfit (2011a, 375) calls "*Motive Consequentialists*," attaching individual moral praiseworthiness and blameworthiness to the intentions of those involved. Consequently, and consistent again with other moral doctrines like utilitarianism, the relevant measure of value here is each party's expected value (in Parfit's terminology, "*Expectabilism"* (2011a, 160)). The actual value may be uncertain and subject to the vicissitudes of fortune. To take a trivial example, if you sell me a lottery ticket at a sufficiently discounted price, then my expected value will be positive but it remains possible, indeed likely, that this will not prove a winning ticket and that, as a result, my actual value will in due course turn out to be negative. In sum then, the most that can be asked of any party to a negotiation is that they *intend* that all parties enjoy an *expected* mutual benefit – that is, together they *seek* a mutual benefit.

Is there then an obligation on negotiating parties to seek a mutual benefit, or at least, if possible, to avoid any counterparty suffering a net loss? Alan Wertheimer (1996, 72, also 210) appears to feel this is so clearly the case that no further argument is required. He simply states that, "it is unconscionable for sellers to make *any* profits on deals that are likely to prove harmful to the buyer." However, he may have a narrower class of transactions in mind. I do believe that an argument is required, and I offer it in three related forms based on moral imperatives (1) to share the rewards of cooperation (2) to recognise just claims in respect of property and person, and (3) to avoid doing harm to others.

The first form of my argument focusses on the sharing of the additional value created by the cooperative exchange that in turn results from a successful negotiation. This is what I have called above the ‘cooperative surplus’ and is what many economists term the ‘social surplus’ (Hausman 1989, 318; Wertheimer 1996, 21). It is the newly created value, resulting from the negotiated agreement or exchange, that is available to be divided between the parties. To seek a mutually beneficial outcome is in essence to accept that this expected value must be shared in some manner. The intention must be that each side comes away with something or, at a bare minimum, suffers no loss. In contrast, for me to seek otherwise is to demand that I take more than 100 per cent of the available benefit, leaving you with less than none. That is, worse than ‘empty handed’. There is something intuitively unacceptable about this latter approach. It seems to fly in the face of the very idea of social cooperation and exchange that many see as a central part of morality. Most of us intuitively recognise that a negotiated agreement, as a basis for cooperative action or exchange, should offer *something* (or, at least, no less than nothing) to each party that contributes to that cooperative action or exchange.

The second form of my argument focusses on just claims in respect of property and person. It is integral to the meaning of a fair negotiation that the parties are trading goods or services to which they have a just claim. By way of contrast, there can be no obligation to respect the interests of a hostage-taker, because in that instance the counterparty has no rightful claim to that they seek to trade. Such a rightful claim to property is then implicit in any fair negotiation and legitimate exchange. When we bring these collective goods and services of ours together in a cooperative action or exchange, we aim to create additional value (otherwise there would be no purpose in such a transaction). If I now insist that all that additional value should accrue to me and none, or less than none, to you, then I fail to acknowledge the legitimate claim that you have to participate in the additional value that you have helped create. I deprive you of any reward for your contribution to our joint endeavour. Moreover, to the extent that you are left worse off than in the absence of an agreement, I induce you not only to forgo any share of the surplus value you have helped create but also to surrender a part of the property that you originally brought to our joint endeavour. Such an outcome seems fundamentally inequitable.

Furthermore, in pursuing such an outcome, I treat you as an instrument for my purpose; as purely a means to the satisfaction of my ends and not as an end in yourself. I subordinate your interests entirely to my own – indeed, to such a degree, that you are left worse off from the transaction than had you not participated. I am thus disrespectful of you as an autonomous individual, failing to acknowledge your worth as a fellow rational being, your claim to "moral recognition respect" (Darwall 1977, 40). As Stephen Darwall explains, "To have recognition respect for someone as a person is to give appropriate weight to the fact that he or she is a person by being willing to constrain one's behavior in ways required by that fact" (45). By totally disregarding your interests, to the extent of leaving you worse off as a result of the transaction, I fail to show such constraint. In short, then, I wrong you in two ways, by failing to acknowledge both your just claim to property and your right to respect as a person.

The final form of my argument focusses on the imperative to avoid doing harm to others. If I knowingly pursue an outcome where you are left with an expected value less than you would enjoy in the absence of agreement, then I accept that you become a 'net loser', that your interests are set back, that you are harmed (Feinberg 1984, 33-34). Now, it might be argued that you bring that harm upon yourself by agreeing a disadvantageous bargain. That may be the case, but whatever the cause, it remains the fact that I seek a particular outcome and it is one where I know that you will be harmed. But it may be further objected that being harmed is different from being wronged. There are what Feinberg calls "common non-wrongful harms" (36). For example, an entrepreneur who develops a new vastly superior product thereby harms her competitors, but she does not wrong them (Brennan 2011, 144-145; Kurki 2018, 443). The difference, though, in the case of our negotiation, is that I seek your agreement and support. I require that you agree to my terms, in order that I may secure the gains of trade. But, in doing so, I ask you to contribute to an outcome which, I know, will result in your harm. You are a participant in our joint enterprise, our joint endeavour, whatever that should be, and yet my intention is that you should not share in its rewards – indeed, that you should instead pay a penalty for your participation. In short, I invite you to join an arrangement where (1) you benefit me; and (2) I harm you. There is, at a minimum, something morally unsettling about this.

All attempts to explain such a patently unfair outcome will fail what Cohen (2008, 41-46) calls the "*interpersonal test*" (42). I cannot coherently articulate to you why my share of the surplus should exceed all that is available; why your contribution to our joint endeavour should be valued at less than nothing; why I should be permitted to use you in this manner; or why your willingness to help me should be rewarded only by my willingness to harm you. I cannot justify these actions "in explicit I-thou form" (44) and to that extent I stand outside the "*justificatory community*" (43). My insensitivity to the norms of fairness in the community mean that I am behaving, to use Cohen's memorable analogy, like a Martian (44).

So, what are the implications of this obligation of mutual benefit? As discussed already, it appears, on first appraisal, to impose a very modest burden. It requires only that our counterparty be left no worse off than in the case of a failure to agree. Thus, we are permitted to optimise our own position to the full, with just a single side constraint: that we seek an outcome where our counterparty comes away with no less value than they would enjoy in the disagreement position. In ideal conditions this side constraint in effect falls away because our counterparty will not agree to anything that is not in their own interest, that is, to any deal that would bring the side constraint into play. As explained, it is only in non-ideal conditions, where our counterparty behaves in a manner that we believe compromises their own interests, that the side constraint does bite; and then, again, only to the extent that the counterparty might be left worse off than in the disagreement position. But, as we will shortly see, this apparently undemanding constraint has potentially far-reaching, and sometimes perplexing, implications which need to be considered carefully. To understand why that is, we must look at what is meant by the disagreement position.

4.6 THE DISAGREEMENT POSITION & ITS IMPLICATIONS

A commitment to leave each party no worse off than in the disagreement position is only meaningful if there is some way of establishing what constitutes that disagreement position. As Sugden (2018) writes, "mutual benefit has to be defined relative to some baseline" (240) and "the baseline must be non-cooperation" (241) – what he also terms the "baseline of non-interaction" (241, 243) or "non-participation baseline[]" (253). Broadly then, a party's disagreement position is the expected outcome of their selected course of action (or inaction) in the event of a failure to agree – in other words, their *next best alternative* among the remaining options open to them. This needs to reflect *all* the circumstances arising in the event of such a failure to agree. After all, the moral obligation we established above is to seek an outcome whereby all parties are left equally well or better off as a result of agreement – that is, no worse off than in the case of a failure to agree. The 'case of a failure to agree' is just that, namely the position our counterparty will find themselves in should the current negotiation collapse.

While this sounds straightforward, the issue is made more complex by the presence of threats and opportunities; and the uncertainty that goes with them. This is especially the case where the opportunities and uncertainty relate to *alternative transactions.* For example, it is often the case, in a transaction between seller and buyer, that the seller's next best alternative is to sell to another buyer; and the buyer's, to seek another seller. In this case, the expected value of the seller's and the buyer's disagreement positions must reflect, to some degree at least, the price available in the market once the current negotiation has failed.

Now it might be objected that this is too literal an interpretation of what is meant by the disagreement position. It might be argued that the disagreement position should instead reflect merely the status quo, assuming that the current transaction has failed, but ignoring the potential for alternative transactions. While this might be a convenient simplifying assumption, I don't believe that such alternative transactions can be wholly ignored. To do so seems inconsistent with the concept of the disagreement position. The disagreement position is not how things are here and now or in some sense prior to the negotiation. Rather it is the position that will be faced by each party in the event of a failure to agree. And this must reflect, to some degree, the possibility of alternative transactions. This can be illustrated in the following thought experiment:

> *Business (1) Competitive Threat*
> Let us imagine that you are a foreign buyer negotiating to purchase my business. I know that, if I do not sell, you will establish your own competing business in my marketplace.

Do we need to reflect that (perfectly legitimate) threat when calculating my disagreement position? The answer must be 'yes'. It is a simple fact of business like any other. And, if the sale were instead to a third party, they would certainly take this competitive threat of yours into account, when valuing the business they were about to acquire.

> *Business (2) Expansion Opportunity*
> Alternatively, what if, in the event of our negotiation breaking down, my next best alternative is an attractive opportunity to take my business into your market?

Should we consider that opportunity in calculating my disagreement position? Again, yes, all legitimate threats and opportunities seem relevant to the disagreement position. And, again, if the sale of my business were to a third party, they would in all likelihood place a similar additional value on this expansion opportunity and reflect that in any offer.

Business (3) Sale Opportunity

But let us suppose the nature of my opportunity was to realise value, not through business expansion, but instead through an immediate sale to that third-party buyer.

This is merely another kind of opportunity and there is no obvious reason why it should be treated any differently. It too should be included when calculating the disagreement position. Thus, it seems that any alternative transaction must be a valid consideration in assessing the disagreement position. It is integral to the concept of the disagreement position that it should reflect *all* the circumstances predictably arising in the event of disagreement. To include some and not others would be arbitrary.

But, with that in mind, let us consider the nature of the resulting obligation with the aid of a further thought experiment:

Car (1) Faulty Exhaust

Imagine that I am negotiating to buy a car from you. Let us suppose that, known only to you, there is a major recurring fault with the exhaust system, the replacement of which will cost substantially more than the proposed sale price of the car. You do not disclose this, and I buy the car in ignorance of this problem.

In this case, it seems highly probable that you have not met your obligation of mutual benefit. If known to me, this would most likely have been a 'deal breaker'. In other words, the outcome of the sale is very likely, for me, inferior to that of the disagreement position. I am left a 'net loser' and you know that I am a 'net loser'. However, …

Car (2) Scrap Metal

What if, instead, I am an experienced scrap metal dealer? The running condition of the car, including the exhaust system, is of no relevance to me. You reason correctly that I must have a potential buyer for the car, a metal fabricator of some sort, who will in due course pay me a scrap value well in excess of the price currently under negotiation.

In this case, the obligation of mutual benefit is respected, because I am better off under this agreement than I would be in the case of disagreement. The opportunity I have to sell the car on to the metal fabricator meets and exceeds the disagreement position.

But, if we take into account such transactional opportunities on my (i.e., the buyer's) side of the negotiation, presumably we should also take into account similar transactional opportunities on your (i.e., the seller's) side. After all, the obligation to seek mutual benefit falls on both of us. Bearing this in mind, …

> *Car (3) Imminent Offer*
>
> Now let us imagine that I have near certain information that my potential buyer, the metal fabricator, instead of transacting with me, would rather transact directly with you. I alone know that the metal fabricator is, at this very moment, picking up the phone to offer you, the present owner of the car, without any need for further inspection, a cash price well in excess of my own offer.

If I conclude our transaction now and buy the car from you at the lower price, benefiting from your ignorance, do I fulfil my obligation to seek our mutual benefit? It seems not. I know with near certainty that, if you agree the deal with me at the lower price, then you will be worse off than you would be in the disagreement position, given that you will imminently have as your next best alternative a significantly higher offer. It seems that, in order to meet the obligation of mutual benefit, I must either withdraw from the deal, alert you to the coming alternative proposal, or match the higher price that I know you will very shortly be offered.

Can this be correct? In these exceptional circumstances, I believe, it possibly can – at least to some degree. I agree with Parfit (2011a, 207) that it is at least plausible that "just as I ought to warn you if the product that I am selling is in some way defective, I ought to tell you that you could buy this product much more cheaply elsewhere." While this appears very demanding, it is rare in practice that such totally asymmetric information, regarding other parties' intentions or preparedness to trade, can be known by just one party with such certainty. In most transactions there is considerable *un*certainty not only regarding competing parties' intentions but also in respect of the value of their potential offers. Part of the problem is that what we mean by value so often encompasses a great deal more than just price: for example, product quality, availability, service, nature of payment (cash, deferred consideration, payment in kind etc.), the importance of a long-term strategic relationship, and crucially confidence that the transaction will actually be consummated. Even in the case above, which appears so clear cut, there remains uncertainty that the imminent proposal from the metal fabricator will turn into a sealed deal. All of this makes a definitive comparison of competing offers much more challenging. In practice, it may be very difficult to establish what is the true value of the next best alternative – very difficult with regard to our own next best alternative, let alone our counterparty's.

Furthermore, it is possible to conceive of circumstances where such an obligation to inform a counterparty of a more attractive competing opportunity might be outweighed by other factors. Consistent with other such constraints, this obligation can only ever be pro tanto. For example, the creation of a new business or the very survival of an older one, upon which many employees depend, might represent an argument, 'all things considered', for the exercise of some discretion in this regard.

But, where we do have near certain knowledge that our counterparty is acting in ignorance of a significantly more attractive alternative, then the analysis above would suggest that we are indeed under an obligation – in most instances – to at least explore this with our counterparty. Thus, a moral imperative, which appeared initially to place only the lightest of burdens on the negotiator, in the form of a seemingly modest side constraint, looks now to be significantly more demanding than on first appraisal. To borrow a past strapline of a popular British retailer, there would appear to exist a duty to be "Never Knowingly Undersold."

However, we must be careful. There is a danger here of sinking into absurdity. What if, as the strap line suggests, we look similarly at the seller's side of this negotiation:

> *Car (4) Impasse*
>
> Let us imagine that just as I, the buyer, am aware of a possible higher buyer, unknown to you; so also you, the seller, are aware of a possible lower seller, unknown to me. Let us further assume that the former threshold is well above the latter.

Here, the obligation of mutual benefit would appear to rule that I am not permitted to offer a price *below* that of the higher buyer, while you are not permitted to accept a price *above* that of the lower seller. Given that the former threshold is well above the latter, then we have an impasse. Thus, potentially, we have a buyer (me) prepared to offer a price that is acceptable to a seller (you), but with neither of us permitted morally to transact!

Actually, the result is a little less absurd than it seems. One logical outturn, in the case above, would be for the higher buyer to first transact with the lower seller, this being the transaction with the largest cooperative surplus. With these parties thus satisfied and eliminated from the marketplace, I, the lower buyer, could then continue to seek a (more marginal) transaction with you, the higher seller.

Nonetheless, this suggests that the obligation of mutual benefit must be applied with care. And this can be further brought out with yet another variation on the thought experiment above:

> *Car (5) Reservation Price*
> Let us suppose that there are no other obvious buyers; and that I offer you a price for the car that you find acceptable. We seem to have the makings of a deal. However, I know perfectly well that, if pushed, I would be prepared to pay a higher price. Indeed, so determined am I to own the car, that I feel sure that, were our negotiation based on the current price to fail, I would certainly revert to you with an offer at this higher price.

In this case, I appear to have sure knowledge that your disagreement position, your next best alternative, in the event of a temporary breakdown in the current negotiation, would be a further offer from me at the higher price. So, if I reach agreement with you on any price below that, it cannot be in your interest. Thus, in order to fulfil my obligation of mutual benefit, it seems I need to offer you the higher price now. I need not only, then, to match the market. I must also match my own best offer. The implication is that a negotiator is obligated from the outset to offer their reservation price, which of course would negate the whole negotiation process.

How do we resolve this apparent absurdity? There are reasons for doubting the force of the objection, implicit in this latest scenario. To begin with, it creates a false distinction between the current and a future negotiation. If I really am prepared to pay a higher price, and it really is required to secure a deal, then why would I allow the current negotiation to fail without at least moving towards that higher price and, if necessary, offering it in full?

Moreover, while it may indeed be the case that I, the buyer, am prepared to pay more than I am currently offering, it may equally be the case that you, the seller, are prepared to take less than you are currently asking. We each have a reservation price. If there is a viable deal to be done, then there is presumably a common area of overlap between the two. But to assert that the result of a future negotiation will be to settle at my reservation price, rather than at your reservation price, or at a point between the two, is to prejudge the result of that subsequent negotiation. It is to beg the whole question. Thus, the parties' *reservation prices* cannot be treated as *predictable outcomes*, any more than can those of other parties yet to join the negotiation. They cannot constitute 'facts' of the situation that determine a next best alternative.

While there may be an obligation to consider the broad parameters of market value involved in any negotiation, these will always be somewhat vague. This vagueness is not a matter of chance. It is not ad hoc. It is *necessarily* the case. Negotiators cannot 'match the market' in any precise way because there is no such market to match. It is in the very nature and circumstances of a negotiation that the perfectly competitive market of economic textbooks, or anything close to that, does not exist (see Chapter 1). Where there is such a perfect market, with an unlimited number of participants etc., the price and other terms are dictated by the market. There is no room for negotiation. As we have seen, negotiation is a tool for reaching agreement in exactly those circumstances where a market, of this kind, does *not* exist. These are circumstances where there is typically an overlap between the seller and the buyer's reservation prices, the task of negotiation being to determine where in that overlap final agreement is reached. It is *possible* in a negotiation that a third party's bidding intentions are known with near certainty, but this is the exception and, even then, as we have noted, a firm bid still remains some way from a sealed deal.

This is not like a traded market, where execution is near instantaneous. In most negotiations, the buyer is at best aware of another potential buyer that *might* pay more, just as the seller is at best aware of another potential seller that *might* accept less. This epistemic uncertainty, together with their typically provisional character, means that such potentially competing proposals cannot impose any precise moral obligation on others. Such an obligation would presuppose a clearly defined, efficient market with near instant execution and, as argued above, such a market leaves no role for negotiation. Negotiation and the kind of market, that would permit of such an obligation, are mutually exclusive. To express the same point in different terms, an efficient market is a feature of the *ideal* world, whereas negotiation is very much a creature of the *non-ideal* world.

These contrasting worlds can be seen as a form of division of labour. A perfect market dictates a single price and related terms; a negotiation determines an outcome between two reservation positions, where no such market or price clarity exists. There is a continuity, of course, between the two. But they do not overlap. Rather one gives way to the other. The more efficient the market, the narrower the spread between the 'offer' and the 'bid', and the more limited the room to negotiate. In the extreme case of a perfect market, the 'offer' and the 'bid' are one and the same, that is, there is a single price. And, in that case, there is no room at all for negotiation.

To suggest that a negotiator must 'match the market' is to confuse these two different roles. If there is a definitive market price – precise, instantly deliverable, and transparent to all

– then the market can generally be relied upon to 'do its job' and deliver a well-defined outcome. In contrast, the role of the negotiator is to agree an outcome where those very characteristics are absent. And to be clear, reservation prices and bidding intentions are not market prices. Instead, they are by their very nature imprecise, subject to various conditions, and far from transparent to all. They are typically starting points, waymarks, or parameters in a negotiation. They are much less commonly end points. In sum, there can be no moral requirement that a negotiator take account of the market, in anything more than the broadest of senses.

What then, in practical terms, is required by the obligation of mutual benefit? It is perhaps best described as the adoption of a win/win attitude. Such a win/win attitude requires a sensitivity to the interests of the other party, a concern for the outcome viewed from their perspective and a commitment to creating value that can be shared by all participants in a transaction. The objective is clear: to ensure that all parties to a negotiation are left as well or better off, as a result of agreement – that all parties are, to some extent, winners. This cannot entail matching a market, that after all does not exist, but it does require an awareness of and concern for the interests of any counterparty, together with a sincere desire to see those interests furthered as a result of any agreement reached. This will require consideration of the counterparty's next best alternative and that in turn will reflect, but only *in broad terms*, the parameters of any wider market for the good or service being negotiated.

The obligation of mutual benefit is characterised, as we have seen, by significant elements of uncertainty regarding the determination both of a counterparty's true interests and of their next best alternative. But again, this is no more the case than with many other moral theories, including utilitarianism. The objective of utilitarianism is clear enough: maximisation of the good. But how one precisely defines the good, or indeed establishes and evaluates the consequences from now until eternity of any given action, is another thing altogether. It is possible to adopt a utilitarian ethos – that is attitude – but beyond that is probably more than can be asked of anyone. The same is equally true of the obligation of mutual benefit.

4.7 PARETO IMROVEMENT & CONSENT

Although negotiation and the market are mutually exclusive, functioning in different domains, there is an illuminating parallel between the two. Several commentators have argued that there is an "*implicit morality*" (McMahon 1981, 254) or aretaic "*telos*" (Bruni & Sugden 2013, 151) to the market, based on the market's ability to deliver "mutual benefit" (153) and/or Pareto optimality (see also Norman 2011; Heath 2014; Sugden 2018; Heath 2019). Every market trade is deemed to leave all participants equally well or better off, thus over time promoting the welfare of all at the expense of none. Furthermore, this ensures that, over the course of many trades, the market is correctly priced to properly allocate resources where they are most needed (McMahon 1981, 255). It is Pareto optimal. In a similar fashion to the game theorists, this reflects an assumption of ideal conditions: a perfectly efficient market, with no barriers to competition, no externalities, symmetric information, and rational behaviour.

Of course, these conditions rarely, if ever, pertain. Heath (2014, 5; 2018a, 529; 2019, 25-26) argues that, where these conditions break down, management is under a moral obligation not to exploit such market failure, because to do so would undermine the Pareto efficiency that underpins the normative justification of the market itself. Christopher McMahon (1981, 255) similarly argues that such conditions impose what he calls "efficiency imperatives" on "economic agents." These leave management duty-bound to protect the Pareto optimality that justifies the system under which they operate. And Bruni and Sugden (2013, 156), in their turn, lay down a closely related requirement that a "virtuous trader" not seek to raise barriers to entry, not participate in cartels and not exploit other forms of market failure.

In more recent work, Heath (2020) expands this efficiency imperative from the private into the public sector. He sees it as the "normative commitment" that underpins the civil service's application of cost-benefit analysis. In his view, the bureaucracy's use of cost-benefit analysis largely employs a "Pareto rationale" (346) to resolve society's collective action problems by pursuing those outcomes that would have come about in the absence of market failure (note that this is an 'actual' Pareto rationale rather than just a 'potential' Pareto or Kaldor-Hicks rationale (191-192)).

The argument for such an imperative can be naturally extended from the macrocosm of the market, and its public sector equivalent, to the microcosm of an individual negotiation. Under ideal conditions, the outcome of any individual negotiation will also represent a Pareto improvement, that is, each party will be left equally well or better off as a result. But, as with the wider market, conditions are not always ideal. Information may be asymmetric and, as we have seen, participants may not behave in their own best interests. As with the wider market, when ideal conditions break down in this way, this can be seen to create a parallel moral obligation on the part of individual negotiators to restore Pareto efficiency. In other words, there is an implied requirement to seek an outcome that is a Pareto improvement, offering something to all parties involved or, at least, leaving none worse off – in short, an implied obligation of *mutual benefit*.

This obligation of mutual benefit, in effect then, directs negotiators to narrow the gap between non-ideal and ideal market conditions. As we saw earlier, it entails a side constraint: that no party be left worse off than they would be in the case of disagreement. Again, as we saw earlier, that side constraint has teeth only in non-ideal conditions. But, as we can now see, its effect, in these non-ideal conditions, is to drive the outcome closer to that which could be expected from ideal conditions. Negotiators are obligated to correct anomalies, where these are sufficient to leave their counterparty worse off than in the case of a failure to agree. In such circumstances, they must seek to protect from loss children and the rationally impaired, to correct cases of egregiously asymmetric information or gross error, to safeguard against conflicts of interest wherever they should arise, and to ensure that those acting under soft forms of coercion really are genuinely at ease with any agreements to which they are party. And, to the contrary, they may not exploit these anomalies in such a way as to leave their counterparty a 'net loser' from the transaction. After all, such a 'net loss' would represent neither a Pareto improvement, nor a mutual benefit.

This emphasis on Pareto improvement, and the implied mutual benefit, has a bearing on the issue of consent. We noted earlier (see 4.3) that many of the cases where people fail to act in their own interest might be characterised, in very broad terms, as a lack of full consent. Children and the mentally impaired are often considered unable to give consent on their own behalf; the grossly misinformed may be deemed not to have consented to the same agreement as their counterparties; conflicted agents may be judged less than fully authorised to deliver the required consent on behalf of their principals; and the overly benevolent or submissive may in some instances be felt coerced against their 'true will', if only by convention. Such a characterisation may appear stretched, placing perhaps too heavy a burden on the concept of consent. However, if one does adopt such a characterisation, then the obligation of mutual benefit might be seen as partially correcting these non-ideal factors and thus rectifying, to a degree at least, any deficit in consent. By insisting that every negotiated agreement is an expected Pareto improvement, the obligation of mutual benefit produces a result that is in the interest of all parties and therefore could rationally have been the subject of a full and uncompromised consent.

4.8 NEGLIGENCE & MORAL HAZARD

Before closing, I want to address a potential objection: the possibility of moral hazard. The claim that the negotiator has a moral obligation to seek an outcome that leaves their counterparty no worse off than in the disagreement position comes very close to what Benjamin Ferguson (2016a, 486) calls a "'come-what may' dut[y]." This presents the counterparty with an, at least theoretical, opportunity to game the situation. The counterparty might choose to risk a very poor outturn in the expectation that any downside, that is worse than the disagreement position, would be protected by this obligation:

> *Moral Hazard*
>
> For example, as a would-be buyer of a business, I might choose to save the expense of a normal 'due diligence' investigation of the company's books, title to land, compliance with environmental regulations etc., in the expectation that, if the putative transaction were to be an exceptionally bad one for me, as the buyer, then you, as the seller, would feel morally obligated to bring this to my attention. I might, in effect, adopt an attitude of 'someone else is obliged to look out for me; hence I don't have to do it for myself'.

I do not believe this objection can be wholly dismissed, but I will attempt to address it briefly, as best I can, both in practical terms and in principle.

In practice it seems to me very unlikely that a buyer would really be prepared to bear such a risk. To rely on what, after all, is only a *moral* duty of *intent* on the part of the seller, and that only in respect of an outturn for the buyer worse than the disagreement position, would surely be precarious in the extreme. It would be to expose oneself to great risk.

Furthermore, and addressing now the principle, while there is something potentially unfair about such 'free insurance', we all do, as Ferguson himself notes, regularly face obligations that place on us burdens that we have done nothing to warrant, for example, "a duty to aid all forms of vulnerability" (Ferguson 2016a, 508). That the seller in the case above faces such a duty is hardly a unique situation. Indeed, it is a situation that is common to many regulated industries, like investment management, where the service provider is duty-bound to assess the client's level of knowledge and bring potential concerns to their attention.

A closely related objection might challenge the implied approach to negligence as a whole: can it really be that a negotiating party is obliged to protect a counterparty from the consequences of the counterparty's own wilful negligence? I suspect that people's intuitions may vary on this, but my response would be 'yes'. At least, 'yes' to the extent of protecting a counterparty from a net loss, that is, an outcome worse than that in the absence of agreement. In this respect, the behaviour of the negligent counterparty is not relevant. The important point is that each of the parties is under a separate and independent obligation to seek a transaction that is beneficial for all. For all the reasons rehearsed above, that moral imperative is integral to the very nature of social cooperation and mutual exchange. And it is an imperative that stands, whether or not one party acts negligently.

My argument here potentially faces twin objections. First, there is the concern that a failure to exploit such negligence, on the part of a counterparty, may frustrate the natural course of competitive attrition. Poor performers will no longer be forced to exit the market. I will concede that, where such negligence is systematic rather than one-off, perhaps a different approach is called for. For example, one might at least factor in the systematic savings enjoyed by the counterparty, over time, through their failure to devote suitable resources, intellectual and otherwise, to evaluating those transactions that they are party to. This could allow a conscientious negotiator to reason that an otherwise disadvantageous transaction for their counterparty might after all be in that counterparty's interest 'all things considered', and therefore morally permissible.

Second, it might be felt that such an imperative places a conscientious negotiator at a disadvantage in a competitive marketplace, thus threatening that party's continued participation in the market. However, over many decades of commercial negotiation this has not been my experience. To the contrary, I am convinced that long-term success in transactionally oriented businesses is built, in significant part, on a track record of value-creating transactions, from which all parties have come away as winners.

4.9 CONCLUSION

Let us pull together the arguments above. Contrary to ideal theory, people at times reach agreements that are not in their own interest. As we have seen, this can be for many reasons: some have only partially developed or impaired rational faculties; others are uninformed, misinformed or simply in error; and still others are excessively influenced by conflicts of interest or the opposite, an excessive sense of benevolence or submissiveness that subjugates their own interests to those of others. When negotiating in such cases, we intuitively sense an obligation to take such counterparties' interests into consideration. We do not feel at liberty to press our advantage to the limit against children, the mentally impaired, the grossly misinformed, the evidently corruptible, or the overly benevolent or submissive.

We have examined what forms this obligation might take. It cannot entail impartiality. A negotiation without partiality is not a negotiation at all. Instead, it is a judgment. Nor does a fair negotiation necessarily demand a proportionately equal division of the cooperative surplus, ‘splitting the difference’. But a fair negotiation does require that each party seek an agreement that leaves not only themselves, but also their counterparty, with some benefit or, at a minimum, no loss: in other words, an expected value equal to or greater than that available in the disagreement position. I have argued for this obligation of mutual benefit on three grounds, being moral imperatives (1) to share the rewards of cooperation (2) to recognise just claims in respect of property and person, and (3) to avoid doing harm to others. I have also drawn on a parallel with recent work done on the implicit morality of fully efficient markets.

The *extent* of this obligation of mutual benefit is more challenging to establish. On first appraisal, it seems to impose only a modest burden on the negotiator. The negotiator is at liberty to optimise their own position with just one side constraint, that they seek an outcome whereby their counterparty is left no worse off than in the disagreement position. However, the disagreement position reflects not merely the status quo, but also any fully credible threats faced by, or opportunities open to, the parties. These must include, to some degree, the reasonable expectation of being able to conclude an alternative transaction with another party. This would appear, then, to impose on negotiators a much more onerous obligation: to ‘match the market’.

However, there is a limit to this further requirement. After all, the role of negotiation is to facilitate agreement in exactly those circumstances where a market, in this sense, does *not* exist. To ‘match the market’ implies a market that is sufficiently developed to dictate price and terms and that, consequently, leaves no room for negotiation. Negotiation and a market, of this kind, are mutually exclusive. It is in the nature of negotiation that it operates in the grey space left between the reservation prices of buyer and seller; and in an undeveloped market, that is often a wide and indeterminate range. Thus, any requirement that the obligation of mutual benefit imposes on a negotiator cannot, by its very nature, be precise. It cannot go beyond taking account of alternative parties to the transaction in the most general terms.

The obligation of mutual benefit is then best seen, in a much broader manner, as the adoption of a win/win attitude. This requires a sensitivity to, and sincere concern for, the interests of any counterparty. The objective is clear: to ensure that all parties to a negotiation are left as well or better off, as a result of agreement – that all parties are, to some extent, winners. My central claim is that this principle, widely considered as supererogatory, is in fact obligatory.

What does this mean in practice? Negotiators are morally bound to enter a transaction seeking a win/win result for all. They must pursue an outcome that, to the extent of their reasonable judgment, is beneficial to their counterparty, as well as themselves. If the counterparty offers to transact on a basis that appears clearly damaging to the counterparty's own interest, then the potentially advantaged party, if they are committed to a fair process and a fair outcome, must ask themselves why that is. If the counterparty appears a 'net loser' for whatever reason – whether they are not suitably equipped to negotiate (e.g., are mentally incapable or insufficiently experienced), have made some evident gross error that would otherwise be a 'deal breaker', are compromised by an internal conflict of interest, or are seriously disadvantaged by an excessively benevolent or unduly submissive disposition which makes them vulnerable to exploitation – then a fair negotiation cannot be concluded on that basis. It is not acceptable to reach such an agreement, knowing that it will leave a counterparty in a worse expected position than would be the case in the absence of agreement. While this might seem an unexpected obligation, we should recall that in those more efficient and transparent markets, where buyers and sellers rationally seek to maximise their own interests, this outcome is likely to be delivered anyway by the natural forces of the market. The moral imperative only becomes relevant as a result of non-ideal conditions, that is, market imperfection or failure.

In sum, then, we have a second requirement of a fair negotiating process: that the parties seek an outcome that is to their mutual benefit.

Chapter 5. The Abuse of Monopoly: Extreme Asymmetric Negotiating Leverage

ABSTRACT: A negotiation can be truthful and mutually beneficial (i.e., it can satisfy the prior two constraints), while nonetheless being monopolistically exploitative and, therefore, unfair. I argue that a negotiation is monopolistically exploitative if, and only if, it involves all of the following: (1) a monopoly, monopsony, or cartel, (2) an essential good or service, and (3) discriminatory pricing. If one of these conditions is unfulfilled, then either there is no vulnerability (in negotiating terms), or that vulnerability is not exploited. This supports the necessity claim. Meanwhile, the argument for the sufficiency claim is as follows: the combination of a monopoly, monopsony, or cartel with an essential good or service creates extreme asymmetric negotiating leverage and therefore vulnerability. The further combination of that vulnerability with discriminatory pricing entails exploitation – taking unfair advantage. I argue that such monopolistic exploitation renders a negotiation pro tanto unfair, that is, morally impermissible. Accordingly, the absence of such monopolistic exploitation represents the third requirement of a fair negotiating process and resulting outcome.

Thus far we have considered two constraints on negotiators: that they behave honestly and that they seek mutually beneficial outcomes. But are these sufficient? The answer is clearly 'no' because it is easy to come up with scenarios where these constraints are respected and yet the resulting negotiating terms are morally offensive. The philosophical literature is rife with such examples, typical cases being our previously described desert hiker, dying of thirst and obliged to pay an exorbitant sum for water (see 4.5.2); or that of a drowning swimmer, forced to give up all her future earnings and work in servitude, as the price of rescue. The common feature of these cases is that they reflect extreme asymmetric negotiating leverage, created by monopoly control of essential goods or services, together with discriminatory pricing. Our moral reactions in such cases suggest the need for a third constraint: a prohibition on the abuse of monopoly.

Let us consider a not wholly unrealistic scenario of our own:

Libyan Captain & Refugee Family

The sun was close to setting over this small port, not far from Tripoli, its fading light competing with the illuminating bursts of shell fire in the distance. It seemed certain that a dreadful battle would engulf the port that night. Somewhere out to sea, in international waters, lay foreign vessels ready to rescue fleeing refugees. But first those refugees needed to escape Libyan waters. The captain briefly surveyed his converted fishing boat, the last vessel left in the port. It was half filled with those desperate for a passage to anywhere but where they were. Each had paid him $1,000. The voyage out to sea was not without some risk and he had found this sum a little disappointing, but it reflected the considerable competition among the various boats in the port earlier that day. Now only his boat remained. He turned his eyes to a family of five that had just arrived on the quay. He approached the father and got straight to the point: "What money and valuables do you have?" The father hesitated. The captain turned to leave. The father looked panicked and gestured to a bag his wife was carrying. "We have cash, property deeds, and jewellery totalling $40,000, enough to get us started in a new life." The captain smiled. "OK, I'll take all that. In return I'll get you and your family out of this hell hole. The boat is seaworthy, but beyond that I can give no assurances. If we are stopped by the authorities, I will simply hand you over. If you're not prepared to deal on my terms, there are other families who are. Do we have a deal?" The father looked at the tracer fire on the horizon, turned momentarily towards his young family, and then nodded.

Is the captain's behaviour morally permissible? Most, if not all, of us will feel that it just cannot be so. And yet, by the two criteria we have discussed thus far, the captain has behaved properly. First, he has not lied. Indeed, he has not hidden in any way the risks entailed by his proposition. Nor has he attempted to disguise the nature and severity of his terms. It would be hard, therefore, to say that he has been in any way dishonest. Second, this is a win/win transaction. There is mutual benefit, assuming, as seems reasonable, that almost any outcome is better for the family than remaining on that quayside. Is this then, after all, a fair negotiation? Again, most of us will feel not. But why is that?

The reason, I believe, is that this is a case of *extreme* asymmetric negotiating leverage, resulting from monopoly control over an essential good or service; and such extreme bargaining power, as this, carries with it a moral obligation not to exercise discriminatory pricing in a manner that exploits the vulnerable. By discriminatory pricing I mean, specifically and exclusively, terms that are egregiously and adversely out of line with those offered to others not acting under compulsion; and by the vulnerable, I mean, specifically and exclusively, those who have no real choice other than to accept the terms offered. In essence, my claim is that for such a negotiation to be monopolistically exploitative there are three necessary and sufficient conditions: (1) a monopoly, monopsony, or cartel (2) an essential good or service, and (3) discriminatory pricing. My argument for this will be based on the concept of vulnerability through the absence of real choice, and our intuitive moral judgments in that regard – as applied to a diverse range of scenarios, including those above. If one of the preceding three conditions is unfulfilled, then I shall argue that either there is no vulnerability (in negotiating terms) or that vulnerability is not exploited. This supports the necessity claim.

In parallel, I will argue for the sufficiency claim as follows: the combination of a monopoly, monopsony, or cartel with an essential good or service creates extreme asymmetric negotiating leverage and therefore vulnerability. The further combination of that vulnerability with discriminatory pricing entails exploitation – taking unfair advantage. Intuitively, we judge that such monopolistic exploitation renders a negotiation pro tanto unfair, that is, morally impermissible. The absence of monopolistic exploitation thus represents the third requirement of a fair negotiating process and resulting outcome.

5.1 PRELIMINARY CONSIDERATIONS

Three preliminary clarifications are required. First, we must recognise that any obligation to avoid monopolistically exploitative negotiation, as with obligations not to lie and to seek mutual benefit, can only ever be pro tanto. There will always be circumstances, where urgent need and the exceptional demands of, for example, distributive justice may call for an element of discriminatory pricing. For example, it might be justifiable to charge a premium for an essential medical procedure in richer, developed countries in order to permit making that same medical procedure freely available in poorer, developing countries. Arguably this would be a case where the pro tanto obligation to avoid monopolistic exploitation in negotiation was outweighed by a separate demand for distributive justice, itself unrelated to negotiation (see further discussion in 5.5.1).

Second, I acknowledge that these terms – monopoly, essential goods (together, implying vulnerability) and discriminatory pricing – are not precise. At least, not precise in the real world of practical ethics. I will seek to explain further what I intend by these terms below. But, for now, let us recognise that these terms are not absolute and there may therefore be *degrees* of exploitation, dependent on the extent to which a negotiation is characterised by these three features.

Finally, before turning to my central claim, I want to address a potential objection of a rather fundamental kind: it might be argued that what occurs in the opening scenario is not a negotiation at all. There is a clear sense in which a deal struck between a monopolistic seller, on the one hand, and a desperate buyer with no alternative, on the other hand, cannot be seen as truly voluntary on the latter's part. To the extent that the buyer cannot walk away, her participation in the transaction is forced upon her. Terms are dictated – even if benevolently so – not negotiated. In an enquiry into the ethics of *negotiation*, it is therefore tempting to simply pass over such monopolies, dismissing them as irrelevant to the discussion. But I prefer not to do that for a number of reasons.

To start with, there is no clear division between a monopolistic dictating of terms and a negotiation. They sit on a continuum. As the number of potential parties to a transaction increases, we move from a monopoly to a meaningful negotiation and, indeed beyond that, to an auction and a market, including a traded market (Michelman 1983, 256; Satz 2010, 30). These are not perfectly defined, or perfectly delineable, concepts. The distinction between the first two is particularly unclear. A pure monopoly is almost unheard of (Walzer 1983, 11). There is nearly always an alternative of some sort, however unpalatable. For example, in the scenario above, the family could always sit out the night on the quay, hoping beyond hope to find another boat the following day. Or the father might attempt to partially negotiate by initially bluffing and walking away. Given this continuum, an analysis of the ethical obligations entailed in a negotiation needs to at least address the situation where negotiating power rests very firmly on one side.

Furthermore, many of the most difficult and challenging ethical issues arise when the balance of negotiating leverage is at its most asymmetric. And the extreme case of this is monopoly control of essential goods and services. Conversely, where a market is fully efficient and leverage at its most balanced, there is little space left for negotiation or related concerns around pricing. As Juan Elegido (2015, 562) observes, "the more a market approaches a situation of perfect competition, the more all participants in that market are price-takers and then there is no room for making decisions about pricing, and no need for ethical guidance in making these decisions." Monopoly power is therefore of particular interest to business ethicists.

Moreover, such situations where negotiating power is notably, even egregiously, asymmetric are too common to be ignored. Most of us, in our lifetimes, will face an exceptionally demanding negotiation – let us call it a 'take-it-or-leave-it offer' – from which we feel we simply cannot walk away. Equally, at some stage, we will likely be a party to a negotiation where our counterparty cannot walk away. When that moment comes, we need a clear perspective on what obligations we owe.

Thus, while accepting that a monopoly over essential goods or services, in its *purest* form, with its implied offering of strictly non-negotiable terms, cannot be seen as a true negotiation, nonetheless for practical purposes I would like to examine monopolies, and the ethical obligations that go with them, as an extreme case of asymmetric negotiating leverage. In talking of monopolies, much of what I have to say has direct relevance also to monopsonies and indeed cartels. And this might be further extended to certain oligopolies, where a small number of suppliers or buyers are, in practical terms, working in concert. All represent similarly extreme cases of this imbalance in negotiating power.

5.2 MONOPOLY

Monopolies have long been the target of moral opprobrium. In England, Max Weber (1992 [1904-5], 259n110) traced this negative sentiment all the way back to the Long Parliament's political struggle with the Crown and the 17th Century competition between smaller bourgeois Puritan capitalists and overbearing financial magnates. Meanwhile, on the other side of the Atlantic, Homer Blosser Reed (1916) was prompted to attack the immorality of monopolies by the newly emerging industrial combinations of the late 19th and early 20th Century. As Joseph Schumpeter (1976 [1942], 100) put it, "monopoly became the father of almost all abuses."

But, as Schumpeter emphatically countered, monopolies can be a necessary and efficient solution to a market need. There are economic domains where small, fragmented competition cannot deliver the required scale efficiencies. A typical example is utilities. And there are other areas where monopolistic constraints on trade – patents and copyright – appear justified in the interests of innovation and creativity. This is recognised by most modern economists: while monopoly and other anti-competitive features, such as tariff protection and barriers to market entry, are widely seen as forms of market failure (Heath 2014, 37; Heath 2020, 161), there is an equal recognition that, in a world of imperfect markets, they may be indispensable, or at least that "there are reasonable grounds for supposing they are not economically odious" (DeMarco 2001, 152; see also Michel 1999, 189). The General Theory of the Second Best suggests that sometimes it is better to 'double up' on market failure – for example, through regulation of monopolies – rather than try to eliminate that market failure itself (Lipsey & Lancaster 1956-57).

In sum, while monopolies are open to potential abuse and in some circumstances can threaten market efficiency, they can also play an important role in creatively and cost-effectively delivering certain goods and services that the market demands. We cannot simply castigate the whole class as morally impermissible.

5.3 DISCRIMINATORY PRICING OF ESSENTIAL GOODS & SERVICES

What is it about certain monopolies then that makes them intuitively offensive? An obvious concern is what is sometimes called 'price gouging'. This can seem particularly objectionable when markedly different prices are demanded for the same product of different would-be buyers, depending on their ability to pay.

On the other hand, we accept such differential pricing willingly in many instances. After all, discounts for students and senior citizens imply a price premium for those in their middle years, better able to pay; while the reduced cost of travel for season ticket holders and others that book far in advance equates to a surcharge on those who, often travelling on business, buy their tickets as they board their transport. Such differentials can become quite extreme in areas like airline ticket pricing. And yet this kind of market segmentation – a form of what economists call third-degree price discrimination (Nicholson & Snyder 2017 [2012], 507) – is widely accepted (Elegido 2011, 637-638; Heath 2014, 191) and indeed, under certain conditions, can plausibly contribute to market efficiency (Aguirre, Cowan & Vickers 2010, 1611). We do not typically accuse those, who offer such student discounts and season tickets, of wrongful exploitation. Indeed, Alexei Marcoux (2006, 61) argues that, if equality of welfare is the correct measure of fairness, then price discrimination is morally necessary since it alone can deliver equality of "*utility enhancement.*" In contrast, unitary pricing offers a greater share of the cooperative surplus to some than others, dependent on the parties' differing reservation prices. Furthermore, Elegido (2011, 643-644) points out that, since price discrimination often entails charging higher prices to those who can more easily afford them, it more often than not has "positive distributional effects." Leveraging a monopoly position to charge the highest price that a counterparty is willing to pay does not then, *of itself*, strike most of us as morally impermissible.

This point is easily illustrated. To take an example of Friedrich Hayek's (1960), let us imagine someone who would like their portrait painted by a famous artist, but is frustrated because the artist will only fulfil the commission at a very high fee. Here, the price charged is not immoral because the good or service is inessential. Although the seller is in a monopoly position, the buyer is not forced by necessity into the transaction. As Hayek himself puts it, "so long as the services of a particular person are not crucial to my existence or the preservation of what I most value, the conditions he exacts for rendering these services cannot properly be called 'coercion'" (119) and are therefore not morally impermissible. Jeremy Snyder makes much the same argument with respect to "diamond-encrusted tiaras" (2013, 351); and widescreen, high-definition televisions: "because this product is inessential, television vendors can ethically charge whatever price the market will bear" (2009a, 286). These transactions are not exploitative. It seems, then, that neither monopoly control nor discriminatory pricing is, of itself, morally impermissible.

Hayek and Snyder's examples point though to a third element that, when combined with monopoly control and discriminatory pricing, *is* capable of rendering such a transaction exploitative and objectionable: the *essential nature of the good or service* to the counterparty. The risk of exploitation is always there when one party has no real choice other than to participate in a transaction – when they cannot, for one reason or another, walk away from the deal. Because it would threaten their core interests to an unacceptable degree were they to walk away, they are left *vulnerable.* This gives over-wielding power to the monopolist. As Mikhail Valdman (2009, 14) argues, it creates a "*prima facie* moral obligation not to extract excessive benefits from people who cannot, or cannot reasonably, refuse our offers." Clearly this does not apply to the art commission, diamond tiaras or widescreen, high-definition televisions described above, because these putative transactions are not essential for either party. But it does explain our intuitive objection to the behaviour of the Libyan captain towards the refugee family. He is exploiting their vulnerability. He exploits them by charging a *discriminatory price* well in excess of what he has charged others – a price limited only by what they are able to pay. And he is in a position to do this because of their vulnerability: with the last boat in the port, he has *monopoly* control of an *essential* service; and, as a result, they have no real choice but to agree. They face what Valdman (2009, 9) calls "unacceptable non-transaction costs." Importantly, the exploitation is the consequence of these three factors *in combination*: monopoly, an essential good or service, and discriminatory pricing.

There is an extensive debate as to what represents exploitation, in its pejorative sense. There is broad acceptance, though, of the starting point: it implies, in common parlance, "tak[ing] unfair advantage" (Feinberg 1990, 179; Wertheimer 1992, 212-213; Wertheimer 1996, 10, 207; Wilkinson 2003, 131; Snyder 2010, 188; Lawlor 2011, 253; Powell & Zwolinski

2012, 468; Snyder 2013, 351; Arneson 2013, 392; Lawlor 2014, 200; Ferguson 2016a, 485; Ferguson 2016b, 955; Brennan & Jaworski 2016, 149; Steiner 2017, 377; Malmqvist 2017, 476, 480; Ferguson & Steiner 2018, 533; Miklós 2019, 60; Richard 2021, 3; Kates 2022, 3; Zwolinski, Ferguson & Wertheimer 2022, sec2). This suggests two necessary and sufficient elements. First, unfair benefit, or self-enrichment – the idea of enjoying a gain over and above what one might fairly have expected. And second, taking advantage of another's vulnerability – reflective of an imbalance of power. Addressing exploitation very broadly, Allen Wood (1995, 142) labels these two factors "*advantage-exploitation and benefit-exploitation*." The exploiter takes advantage of some vulnerability on the part of the exploited in order to derive an unfair benefit there from.

There is a dissenting Marxist view, defended for example by Nicholas Vrousalis (2013; 2021), that takes a somewhat different position. Vrousalis agrees that exploitation requires "an unequal exchange of labour" (2021, 105-106) and entails "the instrumentalisation of vulnerability, the using and playing of another's weakness" (2021, 110). But Vrousalis resists the suggestion that this requires some unfairness, instead insisting that exploitation is present in *any* unequal exchange where a condition of domination exists, even if this is in the context of a fully competitive market. He thus countenances "fair exploitings" (2013, 149n), seeing these as characteristic of a "systematic relationship" (2013, 132, 138) between exploiter and exploited, between capitalist and employee. This, in my view, is too broad an application, going well beyond the more customary use of the word 'exploitation'. In what follows, I employ the narrower usage that requires unfairness, rather than this more structurally driven, class-oriented interpretation.

In particular – and in the context of a negotiation – I argue that the benefit from exploitation must be an *excess* benefit. It cannot be simply the benefit that derives from the normal relationship between buyer and seller – for example, between capitalist and employee – that occurs in a more competitive setting. This excess benefit is an essential element of the exploitation, which in turn is made possible by the vulnerability of one, and only one, of the parties (Valdman 2009, 11-13).

To take our scenario above, if the captain sought to charge only the *same* terms for the use of his boat as he would when his counterparties were not forced, by personal circumstance and lack of competitive alternatives, to accept these terms (in this case the $1,000 per head established earlier in the day when there were many boats in the port), then most of us would *not* view his behaviour as exploitative or morally impermissible – even if those terms were beyond the means of the refugee family. At least, we would not feel that his *negotiating* behaviour was exploitative or morally impermissible. We might hope that he would feel an obligation of charity, a duty to aid in time of crisis. But this would be, in a sense, supererogatory or, at least, it would be distinct from his negotiating position. We would not see his negotiating position itself as unreasonable. After all, he would only be demanding the rate for his boat that he has asked of and received from others in a more competitive market. To offer a discount to, and thus make an exception for, those unable to meet such terms, would be a decision that would stand separately from the negotiation itself. Most of us would see it as an admirable, generous, and welcome decision, a concession that we might all applaud, but not one that we would necessarily have a right to expect. I will return to this below (see 5.4), but to conclude for the moment: the aspect of the negotiation then that is objectionable is the additional benefit – the egregious reward – that the exploiter expropriates from the exploited during the negotiation as a result of the latter's particular vulnerability.

What is this vulnerability? Vulnerability comes in many forms. As Benjamin Ferguson and Hillel Steiner (2018, 540) observe, "some form of vulnerability is present in nearly every transaction, yet not every transaction is an exploitation." To see this, let us imagine that our captain is also a drug smuggler. When he makes port in Italy, he sells heroin to drug addicts, offering his wares alongside other established drug pushers, all fiercely competing on price to secure their customers (i.e., there is no cartel or other quasi-monopoly). This market – and the captain's participation in it – clearly exploits the *habitual* vulnerability of those addicted to the heroin. But the negotiation is not itself exploitative. The addicts are free to walk away from our captain and buy their supplies from his competitors, at prices that reflect the intense competition among the suppliers. The nature of the trade is arguably exploitative, but not the negotiation itself. An exploitative negotiation requires not only an essential good or service, but *also* a monopoly, monopsony, or cartel in order to create extreme asymmetric negotiating leverage. And that monopoly power is absent here. As Alan Wertheimer (1996, 66) expresses this same point, "that B needs food in order to live does not give A (a supermarket) any special power over B if there is a competitive market among sellers of food."

Such 'special power' and resulting vulnerability is equally dependent though on the essential character of the goods involved. Let us now imagine that the captain is the *sole* available source (i.e., monopolistic supplier) of drugs for a party-going business executive, with a taste for cocaine that is casual and occasional rather than addictive. Again, the captain exploits the weakness of his counterparty, but arguably no more so than a local bakery exhibiting iced cakes to customers with a sweet tooth. The executive can always choose to walk away. For him the drugs, like the cakes, are desirable but not essential. As a result, again the negotiation is not itself an exploitative one. Indeed, if a negotiation were deemed exploitative simply because the vendor appealed to the weakness of his customer, as represented by the latter's desires and appetites, then almost every negotiation would be 'exploitative', rendering the concept all but meaningless – a case of what Ferguson (2021, 536) calls "the scope problem."

In short, an exploitative negotiation must involve a monopoly, monopsony, or cartel *and* an essential good or service. Only with the combination of both these elements can a party be said to be vulnerable *in negotiating terms* because it is only in this case that they are unable to walk away. Thus, neither case above is an exploitative negotiation: the heroin is essential in the eyes of the drug addicts, but the captain is not a monopoly supplier, nor part of a cartel. Conversely, the captain is indeed a monopoly supplier for the cocaine user, but in this case the drug is a lifestyle choice and inessential. The relevant vulnerability here is not the nature or condition of the counterparty per se but rather their ability or inability to walk away from the transaction. They are vulnerable to exploitative negotiation only if they are faced with a monopoly provision of an essential good or service that leaves them no practical option other than to agree.

Michael Walzer (1983) points to a similar conclusion when he writes that "needed goods cannot be left to the whim […] of some powerful group of owners or practitioners" (89). He argues that such circumstances of dire need and dominant power demand that these transactions be taken out of the sphere of the marketplace and placed into a different sphere of security and welfare, with commensurately different principles of distributive justice. To instead allow these transactions to progress uninhibited in the marketplace would be to permit "market imperialism," that is, "the dominance of the market outside its sphere" (120). The details of Walzer's thesis I leave aside. But the thrust of the argument seems right: that monopoly or cartel power over essential goods creates a vulnerability that calls for some protection against the potentially harsh implications of a monopolistic market, characterised in such cases by extreme asymmetric negotiating leverage. At a minimum, there is a requirement for a moral constraint on the part of the monopolist.

It is important to be clear that this vulnerability does not mean that such a counterparty is forced to participate entirely against their will. In one sense, at least, they agree willingly: they prefer to transact than not to transact. And this is perfectly rational, because they are, in their own eyes at least, net beneficiaries from the transaction. Otherwise, ceteris paribus, they would not agree to transact. But, to adopt Snyder's (2013, 354) terminology, this seems at best a "surface endorsement of a demeaning choice." It is simply that the excessive price they are asked to pay, and that they nevertheless accept, reflects their vulnerability, the fact that in practical terms they have no real choice other than to transact. Or, at least, their choice is very severely constrained. They do benefit but, as Wood (1995, 148) puts it, "being benefited and being exploited are often two sides of the same coin." And, as Wertheimer (1992, 223) says, "it is precisely because the exploitee stands to gain so much from the transaction (relative to the exploiter) that his bargaining position is comparatively weak." Our refugee family willingly surrender all their possessions because that is their only route to preserving their lives or at least some tolerable form of life. And yet, willingly though this is done, their decision is in another very real sense involuntary. Realistically, they cannot walk away.

Let us take stock. I have suggested quite a narrow set of necessary and sufficient conditions for exploitation as it relates to monopolistic negotiation: the combination of (1) a monopoly, monopsony, or cartel over (2) an essential good or service, that is subject to (3) discriminatory pricing. By discriminatory pricing, I mean terms that are egregiously and adversely out of line with those offered to others not acting under compulsion; and by the monopolistic provision of an essential good or service, I imply that the counterparty's agreement is in an important manner involuntary – they are made vulnerable. Intuitively we acknowledge that, when any one of these conditions is fully absent, a negotiation cannot be monopolistically exploitative. Such is the case where the captain offers the same price to all (there is no discriminatory pricing), where the captain competes with other sellers of heroin (there is no monopoly or cartel), and where the captain is the sole supplier to the occasional cocaine user (there is no essential good or service). These intuitive responses support the necessity claim: all three conditions must be satisfied. In many of these and similar cases, exploitation may occur despite the absence of one of these conditions, but that exploitation is not monopolistic in character and/or does not attach to the negotiation itself.

Meanwhile the sufficiency claim can be summarised as follows: the combination of a monopoly, monopsony, or cartel with an essential good or service creates extreme asymmetric negotiating leverage and therefore vulnerability. The further combination of that vulnerability with discriminatory pricing entails exploitation – taking unfair advantage. Intuitively, we judge that such monopolistic exploitation renders a negotiation pro tanto unfair, that is, morally impermissible.

5.4 FURTHER CLARIFICATION & QUALIFICATION

Significantly, I do not argue that this is a full account of exploitation. Clearly there are other aspects of a transaction which might be described as exploitative and which I have not sought to examine. For example, the very nature of the good or service, such as pay-day loans, addictive drugs, gambling, human organs, sex, and surrogacy; or the structure of society itself and the principles of distributive justice that underpin this (for further discussion on distributive justice see 6.1). My focus here is on the factors that make a *negotiation* exploitative and, specifically, in *monopolistic* terms. Even so, it may be objected that this emphasis on the use of monopoly power in discriminatory pricing of essential goods and services is still too broad and ill-defined. Each of these three elements is open to objection, or at least to a demand for clarification.

To take discriminatory pricing first, how does one determine whether a price is exploitative? What is a fair price? This has been debated since the time of Aquinas and before; and is far from straightforward to establish (Elegido 2009, 42, 45). It is especially difficult where a monopoly over an essential good or service is concerned. Let me first seek to eliminate the extreme positions.

At one extreme, there is marginal costing. Let us imagine a monopolistic price for a good or service that is essential to a counterparty, but beyond that counterparty's means. Might that still be considered exploitative, even if it were the *same* price as that offered to all other parties, including those for whom the good or service is inessential? Perhaps, it might be argued, such a good or service, because it is essential, should be offered at no more than the monopolist's marginal cost. This though would be a strange position to take regarding a commercial negotiation. It would after all require that the monopolist forgo any marginal profit at all and, presumably, incur a fully costed loss, once fixed costs have been taken into account (Elegido 2015, 559). Such an approach might be in line with a duty to aid others in dire need, but then such a duty would surely exist regardless of price – indeed regardless of any financial consideration at all. If, according to such a duty to aid, our captain is morally bound to take the refugee family, if necessary, for no recompense, then perhaps he must do so, but we would seem to have firmly departed the realm of legitimate negotiation. Our focus is the ethical constraints that attach to a negotiation, and specifically here to a monopolistic negotiation. A duty to aid, or similar appeal for charitable behaviour, may be entirely warranted, but it stands apart from such a negotiation. It lies in a different domain or, as Walzer (1983) would term it, a different "sphere." It seems then that such an obligation on the part of the negotiator, *qua negotiator*, can be at most supererogatory.

Conversely it is argued by some that price gouging, even in its more extreme forms, is morally permissible and indeed desirable. Matt Zwolinski (2008, 362-364), in particular, emphasises the signalling function of high prices in both stimulating an expansion of supply and depressing demand, other than where need is greatest. He suggests that the market-clearing price reflects "morally significant characteristics like intensity of need" (362). However, as Snyder (2009a, 279-284) points out, the market-clearing price reflects not only intensity of need but also ability to pay and can therefore exclude the poorest in a community, transgressing important values like "the norm of equal respect for persons" (280). As we saw earlier, Snyder is not concerned when this affects goods that are inessential, such as diamond tiaras or widescreen, high-definition televisions. But he does challenge the moral permissibility of such pricing when it threatens "equitable access to *sufficient […] goods to enable minimal human functioning"* (2009b, 304), that is, when it impacts the allocation of *essential* goods.

What then is the most reasonable measure of the 'fair price', when essential goods and services are the subject of a monopoly? Pricing at marginal cost unreasonably penalises the monopolist, while unfettered demands of the kind made by the Libyan captain of the refugee family appear equally unacceptable. The salient position between the two extremes would seem to be some form of market price and this is, indeed, favoured by most commentators (e.g., Wertheimer 1996, 230-236; Michel 1999, 186; Maitland 2002, 453, 471; Meng 2005, 19; Valdman 2009, 12; Elegido 2009, 32; Perrone 2014, 226; Elegido 2015, 557, 562-563). This does, however, create both epistemic and conceptual challenges.

The *epistemic* challenge arises because the very existence of a monopoly is bound to distort whatever kind of market exists. Establishing a market price is likely to be especially problematic in the case of longstanding monopolies where, as Christian Michel (1999, 189) suggests, reliance must be placed "not in the fact that there is competition but in the very *possibility* of competition." In most cases though there are sufficient transactions with a reasonably diverse group of counterparties, from which some kind of conclusion can be drawn.

But how do we *conceive* of a fair market price in such circumstances? Wertheimer (1996, 230-236) favours what he calls the "hypothetical market price." He variously defines this as the "price that an informed and unpressured seller would receive from an informed and unpressured buyer" (230), the "price at which neither party takes *special* unfair advantage of particular defects in the other party's decision-making capacity or special vulnerabilities in the other party's situation" (232), and the "price that would be generated by a competitive market"

(230). The concept of the price that would be paid by an informed and unpressured party, who has not been disadvantaged due to any special vulnerability, seems to draw very close to what we are seeking. It would appear to allow us to neutralise the effect of the vulnerability.

The reference to a competitive market though is potentially problematic. Elegido (2015, 558-559) points out that the equilibrium price in a fully competitive market is "the marginal cost" of production. In the case of a monopoly this is presumably the monopolist's own marginal cost of production. As argued earlier, pricing at this level would be unreasonably demanding of the monopolist, leaving them with no contribution to fixed costs, no reward for innovation, and no incentive to reinvest.

Elegido (2015, 562; also 2009, 32) suggests as an alternative an "open market standard," which requires that there be "a significant number of buyers" and, moreover, "a variety of buyers," but not a competitive market, let alone a perfectly competitive market. Such a variety of buyers seems important to ensure that the price reflects a level acceptable to a range of counterparties for whom the good or service is not essential. This would appear, then, to provide a suitable benchmark. Afterall, our concern is with the vulnerability entailed by a combination of monopoly and dependency (on essential goods and services). It is therefore natural to look at the price that would prevail in an open market, where there is a variety of other buyers, among them those for whom such goods and services are not absolutely essential – what Elegido (2009, 34) characterises as "the price [one] charges to the generality of [one's] customers." Where monopoly control of essential goods and services is concerned, the open market price, then, seems as close as we will get to a measure of the fair price.

While representing the best available solution, we must though acknowledge that this is not a perfect solution. A critic might ask what we are supposed to do if such a good or service is essential to all customers and priced equally exorbitantly across the board. Let us imagine, for example, that our Libyan captain never faces any competition, at any point in the day, and is thus able to charge every family that he transports $40,000. If every family pays the same, then it would appear there is no discriminatory pricing. Without discriminatory pricing I have argued there can be no monopolistic exploitation. And yet many will feel that this is monopolistic exploitation – on the grand scale! I would agree. But in reply I would argue that my position remains consistent because the $40,000 *is* discriminatory, at least *hypothetically*. It is substantially in excess of what we believe the captain would seek to charge his customers under more normal circumstances – for example, the price he would charge were he taking his customers on a fishing trip, even one in high seas that carried an element of risk. In other words, the price of $40,000 is discriminatory because it is egregiously out of line with what would be charged of a hypothetical customer for whom the trip was not essential. I accept that this appeal to a 'hypothetical customer' is a little untidy, but I do not believe it is ad hoc. It is consistent with the general concern over an egregious price that is specifically designed to exploit a particular vulnerability brought on by monopoly control of an essential good or service. On that basis, I believe the objection can be met.

There is a danger here that I must guard against. It could seem that in using such terms as the 'non-discriminatory price', the 'fair price', the 'open market price', and the price paid by a 'hypothetical customer', that I am appealing to some abstract measure of fairness – that is to some independent benchmark, formed separately from and prior to the negotiating process. A kind of judgment. That is not my intention. My focus remains, as previously argued (in Chapter 2), on a fair process. The challenge is how to make the process fair. The problem with monopoly control of essential goods and services is that the negotiating process is distorted – almost to the point where there is no real possibility of negotiating (see 5.1). It is distorted by the vulnerability created by the weaker party's essential need – an essential need that means they cannot in practical terms decline. They cannot walk away. The focus on an open market price is not an appeal to an independent benchmark but rather an attempt to create a parallel negotiating process free of this distortion. In looking at what a diverse range of buyers would pay, including those for whom the good or service were inessential, my aim is to remove the effect of the extreme vulnerability. The resulting open market price is intended to reflect the price that would have been agreed were the negotiating process not to have been distorted.

Let us step back for a moment. To place this in a broader context, the insistence on non-discriminatory pricing can be seen as another case of correcting for market failure. As argued earlier (see 5.2), some monopolies are in practice unavoidable. But they remain nonetheless, in economists' terms, cases of market failure. By seeking the price that would prevail in an open market, with a variety of buyers, including those for whom the relevant goods or services are not essential, we are in effect looking to reverse, to a degree at least, the impact of such market failure and thus to reflect something closer to ideal conditions. We are seeking to eliminate what Wertheimer (1996) terms the "supracompetitive profits" (303) and thus to "replicate the results of a more competitive market" (73). There is here a parallel with our previous commitments to the avoidance of lying (Chapter 3) and the seeking of mutual benefit (Chapter 4). These too placed an obligation on negotiators to promote something closer to ideal conditions, by in those cases eliminating the most egregious information asymmetry and not exploiting a counterparty that is failing to act in its own interest.

Moreover, there is a further parallel with regard to consent. The consent of a counterparty that is lied to is arguably compromised. On this view, avoiding lying helps to ensure that all parties are validly consenting to the same agreement (Chapter 3). Furthermore, the instances where a counterparty fails to act in its own interest could perhaps also be seen as cases of compromised consent (due to immaturity, inexperience, misinformation, gross error, conflict of interest, excessive benevolence, or extreme submissiveness). On this view, the pursuit of mutual benefit is an attempt to secure an outcome that the counterparty could rationally have consented to (Chapter 4). Similarly, here, the vulnerability experienced by a party whose essential needs are subject to monopoly control – the lack of any real option to walk away – may also suggest a consent that may be, in some manner, impaired. As Wertheimer (1996, 27) observes, "it might be argued […] that choices made under conditions of desperation or from an inequality of bargaining power are not appropriately consensual." In contrast, by seeking *not* to exploit such vulnerability and instead offering open market terms, the counterparty's consent is made that much more meaningful.

It is important though not to overstate this concern over the validity of consent. Just because a party chooses an outcome in the face of severely constrained options, does not mean they necessarily fail to consent. Valdman (2009, 3) cites the case of a patient, with an appendix that is about to burst, who faces a choice between urgent surgery and imminent death. The patient may in practical terms have no option but to consent to the surgery and yet we do not, as a consequence, deem their consent invalid. By parallel argument, there must be some doubt as to whether the simple constraining of choice through the monopoly provision of an essential good or service can alone be said to invalidate consent.

Nonetheless, by insisting on a non-discriminatory price – expressed as that which would prevail among parties for whom the good or service were not essential – there is a sense at least, in which we draw closer to the terms that might have been expected were that consent to have been unimpaired. The moral imperative on a monopolist, negotiating with a vulnerable counterparty, to adopt non-discriminatory, open-market pricing is indeed consistent with both our previous two constraints on a fair negotiation. All three drive in the same direction: the mitigation of market failure; and the bolstering of consent.

Let us turn now to the other conditionalities: the monopoly provision of an essential good or service. The implication here is that such combined circumstances entail that the counterparty's agreement is involuntary, in the sense that they have no real choice in the matter. Now, it might be objected that, as I have already conceded, almost no action is entirely involuntary; and, consequently, that there is no clear distinction between what is voluntary and what is not. We are left with a gradation, a continuity between the two, the result of which is that there can be no strict demarcation of what constitutes vulnerability and therefore an exploitative negotiation. I accept this. As is so often the case with moral philosophy, the clarity of the theoretical must ultimately make way for the lesser certainty of the applied. But the precept remains clear: to the extent that a party holds monopoly power, or near monopoly power, over a good or service, that is in practical terms essential to a counterparty, then the monopolist must not seek an excess reward by charging a price over and above that demanded of others for whom the good or service is inessential. In other words, the party must not use discriminatory pricing to exploit a counterparty that is made vulnerable by a combination of monopoly and desperate need. Clearly this cannot apply to competitive markets, nor to inessential goods and services. It does though come into force where there is monopoly control, or at least extreme asymmetric negotiating leverage, over goods and services that are, for all practical purposes, essential.

In the workplace, one form that this extreme asymmetric negotiating leverage can take is domination: the near total effective control of an employee by her employer. Joshua Preiss (2019, 886), in particular, highlights the plight of migrant workers, who have incurred debts to their employers in order to move to a foreign country, where that employer represents their sole opportunity for employment. In this case, their "exit options are virtually non-existent." They cannot walk away. The employer is the sole purchaser of their labour. This is a clear case of monopsony control (by a single employer) over an essential good (employment) and thus justifies moral (and arguably legal) protection against the imposition of discriminatory terms and conditions. Again, a vital element of this is the essential nature of what is at stake.

But what makes something essential? First of all, we should be clear that these essential goods and services do *not* equate with what John Rawls (1971) calls "primary goods". The latter Rawls defines as "things that every rational man is presumed to want" (62, also 90-95). The fact that all rational men, in his view, want for example more "income and wealth" (62) cannot make such things essential without limit. Essential goods and services are, for our purposes here, something closer to what Harry Frankfurt (1987) calls "urgent needs" (35). Consistent with his "doctrine of sufficiency" (22), they are measured in absolute not relative terms. The relevant moral concern for those with unsatisfied essential or urgent needs is not that they have "*less*" but rather that they have "*too little*" (32).

So, what things are truly essential? At the extremes this is usually clear enough. Most of us would consider sufficient food and water, clothing and housing, heat, light and sanitation, healthcare and education, employment opportunity, and freedom from harassment to constitute essential needs. Equally, we would not include champagne, caviar, portraiture, diamond tiaras and widescreen, high-definition televisions. But what about those things that might be considered to fall somewhere in between, such as basic transportation, access to the internet, enjoyment of the countryside and occasional holidays with family? Clearly there is room for considerable debate and subjectivity. It is apparent that the set of essential goods and services will vary both with time and across cultures. Electricity might be considered essential in the modern developed world but is of no value to certain indigenous forest dwellers and was unknown to those of the Napoleonic era. Even within one timeframe and culture, what one person deems essential, another may feel ambivalent about or have no need for. As Elizabeth Anderson (1999, 320) points out, a wheelchair is essential if you have no legs, but of no obvious use if you are fully mobile. Following Amartya Sen and Martha Nussbaum, perhaps we should think of essential goods and services as those required to support certain "basic capabilities" (Sen 1982, 367-369) or "central capabilities" (Nussbaum 2011, 17-45) needed for a flourishing life. The critical underlying question though remains the same: do the parties to the negotiation have a reasonable option of walking away or are they in practical terms forced to agree to whatever is demanded of them? Are they, to use Hallie Liberto's (2014, 627) expression, "caged?" The key once more is the sense of vulnerability that accompanies a lack of any real alternative.

It is clear from this discussion that my claim must be qualified. These three components go a long way to explaining our intuitive concern around monopolistically exploitative negotiations. But none of them is precise – at least not as we encounter them in the real world. Neither the noun 'monopoly', nor the adjective 'essential', normally permit of degree. However, as we have seen above, in reality they both characterise an extreme point on a continuum. And their combination – vulnerability – is clearly variable by degree (Palmer & Hedberg 2013, 412), as of course is discriminatory pricing. Similarly, while it is commonplace to assume that an action is either morally permissible or morally impermissible, but not somewhere in between (Ferguson 2016b, 958), in reality most of us experience a spectrum of judgment from 'clearly permissible' to 'clearly impermissible', often with an ill-comprehended and much-contested dividing line between the two – what Campbell Brown calls "degrees of rightness" (2016, 21-23) and "degrees of wrongness" (2020, 1371-1372). So, perhaps we should look at a monopolistic negotiation through a more nuanced lens. On this view, (1) the more strictly monopolistic the negotiation, (2) the more truly essential the good or service, and (3) the more egregious the discriminatory pricing, then the greater the unfairness of unconstrained pursuit of maximal benefit by the more powerful party, and thus its degree of moral impermissibility.

The necessity claim would still suggest though that an element of all three is essential to establishing exploitation in the context of a monopolistic negotiation. As I have sought to demonstrate, a complete absence of any one component undermines the case for a monopolistically exploitative negotiation. If there is a highly competitive market, then a party is always free to negotiate with another; if the goods concerned are completely inessential – frivolous luxuries – then the party can always choose to decline, without suffering any great loss; and, if the pricing is non-discriminatory (i.e., not egregiously out of line with that determined by an open market and accepted by all, including those not acting under compulsion), then there can be few grounds for complaint – at least as regards the ethics of the negotiating process.

But to insist on the absolute – absolute monopoly, absolutely essential goods, or absolutely egregious pricing – as necessary and sufficient grounds for exploitation is too stringent. These factors work in combination and, to some extent, the severity of one component may permit of exploitation, to a degree at least, even if another component is less than absolute. Once again, it is these three factors, *taken together*, that offend our moral sensibility.

5.5 OBJECTIONS

Let me turn to potential objections:

5.5.1 Challenging Cases

These come partly in the form of contrary cases and for two of these I am indebted to Michael Otsuka, one in respect of the necessity claim, the other in respect of the sufficiency claim. To take the former first:

> *Exploitative Cinema*
>
> A cinema uses predatory pricing to drive the one other cinema in town out of business, only then to negotiate a much higher price with the local film clubs for their evenings out.

This feels monopolistically exploitative and yet it does not involve an essential good or service. This would appear to contradict the necessity claim that monopolistically exploitative negotiation is only possible where essential goods or services are involved.

However, it seems to me that, if a wrong is committed here, then it relates to the initial anti-competitive practice, represented by the predatory pricing, rather than to any subsequent negotiation. The wrong is committed *before* any negotiation unfolds. The use of such bargaining power is not wrong per se. If a vendor is the sole provider of a non-essential service, then presumably it is that vendor's right to choose whether they provide that service or not, and on what terms. They are not obliged to provide the service. In this case, nobody is obliged to run a cinema. They might choose instead to turn the building into an art gallery or an exhibition venue. If they do choose to operate as a cinema, and nobody chooses to compete against them, then it seems they must be free to price their service accordingly. As long, that is, as they permit open competition, should others choose to compete. And this, of course, is where the morally impermissible conduct occurs. What causes offence here is not the use of such a strong negotiating position, but the manner in which that competitively unassailable position is achieved in the first place. The moral impermissibility attaches not to the negotiation, but rather to the predatory pricing that predates the negotiation. So, our intuitive sense that there is 'foul play' involved here is well grounded, but the root cause of that unease is the anti-competitive practice, not a monopolistically exploitative negotiation (since the conditions required for the latter are not fulfilled).

The second case challenges the sufficiency claim:

> *Non-Exploitative Pharmaceutical Company*
>
> A pharmaceutical company, with monopoly control of a life-saving vaccine, negotiates a substantially higher price from richer countries in order to offer the same vaccine at a discount to poorer countries.

This does not appear exploitative, at least not in a morally reprehensible manner, and yet all three sufficient conditions of wrongful exploitation are present: a monopoly, an essential good, and discriminatory pricing. So, the conclusion that sufficient conditions for an exploitative negotiation have been met appears to conflict with our intuitive judgment that this is, however, not exploitative, at least not in a pejorative sense.

I think the explanation is that there are two different negotiations going on here. The negotiation with the richer countries *is* exploitative and, were it to stand alone, simply enriching the pharmaceutical company, it would indeed be morally reprehensible. But, as argued at the outset, monopolistic exploitation is only wrong pro tanto and, in this case, the wrong is outweighed by the greater good/obligation of supplying a life-saving vaccine to those who otherwise could not afford it. This is a case of what one might call Robin Hood syndrome: the rich are being robbed (a pro tanto wrong), but only in order to offer succour to the poor (perceived, at least, to be a greater good/obligation that outweighs that pro tanto wrong). The 'all things considered' outcome is therefore morally acceptable. Thus, a view that the pricing of the vaccine to richer countries is monopolistically exploitative pro tanto is consistent with our intuitive judgment that this is morally permissible 'all things considered'.

5.5.2 The Nonworseness Claim

There is, however, a more generalised objection. The claim that moral reprehensibility attaches to monopolistic exploitation would appear to fall foul of what Wertheimer (1996, 289-293) has coined "the nonworseness claim." This objection argues that if it is morally permissible for a party not to interact at all with a counterparty, as is thought to be typical of most negotiations (Kates 2019, 27), then it cannot be morally worse for that party to interact with that counterparty on terms that are consensual, mutually beneficial, and free of negative externalities. Any such agreement – even a monopolistically exploitative one – must be better for the counterparty than no agreement at all; and thus, if no agreement is morally permissible, such a superior outcome must also be morally permissible (Powell & Zwolinski 2012, 470).

The nonworseness claim appears to challenge the intuition many have that there is something unfair, and therefore morally impermissible, about negotiations like Desert Hiker or that between our Libyan captain and the refugee family. There is an apparent conflict between welfare and fairness. Welfare considerations argue for any transaction that is mutually beneficial (the nonworseness claim), and yet concerns about fairness emerge if the resulting benefit enjoyed by the controlling party seems disproportionate (the conflicting intuition). Attempts to resolve the "problem" (Horton 2019, 469), "puzzle" (Arneson 2013, 395) or "paradox" (Ferguson 2016b, 957) of this "exploitative benefiting" (Meyers 2004, 320) have been wide ranging. They have drawn, for example, on a less restrictive, weakened version of the nonworseness claim (Ferguson 2016b, 965), on the distinctive nature of fairness driven by considerations other than outcomes (Meyers 2004, 320-322; Bailey 2011, 246), on the obligation to share mutually created gains (Kates 2022, 7-11), on the implications of "complicity" and its "expressive effects" in validating and perpetuating "structural injustice"

(Malmqvist 2017, 480-485), and on the value of the heuristic moral rules advocated by "multi-level act consequentialism" (Arneson 2013, 407).

Rather than debate these strategies, though, I would prefer to focus on an argument that is more specifically related to *monopolistic* exploitation, as I have defined it above. I will draw on an account proposed recently by Brian Berkey (2020, 244-248; and especially 2021, 48-52), which I believe fails in its generality, but does carry weight when applied to the more circumscribed conditions represented by the monopoly provision of essential goods and services. (A further, somewhat similar argument to Berkey's is offered by Joe Horton (2019, 472-478) and a more qualified one by Tadhg Ó Laoghaire (2022)).

Berkey challenges a core premise of the nonworseness claim, that it is morally permissible for a party not to interact at all with a counterparty. For example, he suggests that it is, on the contrary, obligatory for multinational enterprises ("MNEs") to source their production in less developed, less wealthy countries. Such MNEs, in his view, have "positive duties to the global poor" (2021, 48) that oblige them to employ workers in developing countries. Clearly, then, if Berkey is right and noninteraction is potentially *im*permissible, it no longer follows that an interaction, which is better than noninteraction, is necessarily permissible. After all, it is possible for a course of action to be 'no worse' than something that is impermissible and yet still itself be impermissible. On Berkey's account, the nonworseness claim is thus undermined.

Berkey's arguments for this assertion are, however, contentious. His defence based on the claim that this alone explains the paradox above might be considered question-begging, given that the paradox is what he set out to explain in the first place. His further arguments based on the resources available to MNEs, the priority that he believes should be afforded to the least well-off, and the central importance of employment opportunities, all seem less than decisive. Meanwhile his final and, in his own view, "most powerful" argument appears to be worryingly circular. He derives the MNEs' positive duty to engage with the global poor from the claim that the badly off in developing countries are the victims of a "global structural injustice," from which MNEs benefit. He goes on to explain that these MNEs benefit because they "typically locate production sites in poor countries" (2021, 51). This amounts to saying that MNEs have a positive duty to locate production facilities in poor countries, because they locate production facilities in poor countries. This feels logically problematic.

Furthermore, for many people, there will be something intuitively unpersuasive about Berkey's putative obligation on MNEs to locate their facilities overseas. Such offshoring, at the expense of domestic employees, is certainly not applauded by all. Indeed, as Michael Kates

(2019, 27) observes, "many believe that keeping production in the developed world is something to be praised."

However, my purpose here is not to debunk Berkey's argument. Instead, I want to suggest that Berkey's challenge to the nonworseness claim *can* be upheld if confined to a narrower set of circumstances. Specifically, I argue that there *is* a positive obligation on one party to interact with another, *if* that party is the only available provider of goods or services that are essential to that other party, assuming always that such interaction does not come at an unreasonable cost to the former party. I do not claim that this is a requirement of a fair negotiation. Indeed, it has very little to do with negotiation. This obligation is instead one of beneficence. It stems from a duty to aid.

The easiest way to illustrate the importance of monopoly control of essential goods and services, in underpinning this obligation, is via a direct appeal to our intuitive judgment. Let us return to our Libyan captain. As the owner of the only means of escape for the family, from a high risk of serious injury or death, most of us will intuitively feel that the captain is *morally bound* to assist them, as long as there are not unreasonable additional risks to him or his boat in doing so. If, on the contrary, the refugee family had a choice between various means of escape, including a fleet of other boats, a light aircraft, a route overland etc., this intuition would either cease to be the case or would, at least, be significantly diminished. A similar distinction can be discerned in Peter Singer's (1972) renowned case of the child drowning in the nearby pond. As the only available rescuer, we would feel *morally bound* to intervene, especially if the cost were merely the spoiling of a suit.

Berkey (2021) himself highlights the importance of monopoly control of essential goods and services, in underpinning this obligation, without perhaps intending to do so. He believes the MNEs' obligation to locate production facilities in poorer countries is "at least roughly, analogous" (50) to a thought experiment he calls *Exploitative Provision of Treatment 2* (47). But there are very substantial differences between the two. In the thought experiment, A has the only dose of a drug D that is required by B and nine others, all of whom will die without it. Berkey reasonably concludes that A is morally bound to interact with at least one of the ten to save their life. But in this case A is in *monopoly* control of an *essential*, indeed life-saving good, D. By contrast, in the case of MNEs considering where to locate, such MNEs are typically one of many such enterprises and are in control of something of great importance, but arguably not of life-saving importance. Whether employment in a particular sweatshop, instead of employment elsewhere, is absolutely essential or not is at least debatable and may depend on circumstances. Our intuitive judgment that A is morally bound to assist B (or one of the nine others), thus reflects important elements of the thought experiment that are distinct from the decision faced by MNEs – specifically the *monopoly* status of A and the *essential* nature of D.

To conclude then, I do not argue, as Berkey does, that the nonworseness claim can be defused by a broadly based appeal to a positive obligation to interact with counterparties. However, I do argue that there is such an obligation in the much more limited case of monopoly control of essential goods and services, where this does not entail an unreasonable cost to the monopolist. We recognise this intuitively as an obligation that falls on those who find themselves in sole control of those things that are absolutely vital to others' well-being. This does effectively defuse the nonworseness objection as far as it relates to my much narrower purpose in this chapter, that is, my goal of establishing the required conditions for *monopolistic* exploitation and its moral impermissibility.

5.6 A NOTE ON FAIR DIVISION

Before closing, I want to touch briefly on the subject of fair division. Monopolistic exploitation can be seen as an extreme case of *un*fair division. The transaction involved is mutually beneficial, but the share of the cooperative surplus that goes to the monopolist is seen as disproportionate (although usually much less than half, because the goods or services being negotiated are of such high value to the exploited party). The resulting division is felt to be unfair. I want to suggest that this sense of unfairness is consistent with the three criteria we have established above for a monopolistic exploitation.

This can be illustrated with reference to the ultimatum game. Under the rules of this game, a sum of money must be divided between two participants. One participant proposes a split. The other participant can either accept or reject this. If the proposal is rejected, then neither participant receives anything. The rational solution, suggested by rational-choice game theory, is that the proposer offers the smallest positive amount available, and the responder accepts this (on the basis that something is better than nothing). But extensive empirical research has repeatedly confirmed that this rarely happens. Instead, proposers in general offer something much closer to an equal division (widely considered, in this case, to be a salient benchmark of fairness) and responders regularly refuse anything that strays too far from that (Guth, Schmittberger & Schwarze 1982, 374-375, 384-385; Fehr & Schmidt 1999, 826; Heinrich et al. 2004, 18-19). A similar pattern can be seen in other games, like divide-the-dollar where participants invariably opt for an even split (Alexander & Skyrms 1999, 589).

Brian Skyrms (2014 [1996], 29) believes that participants in these games "rather than maximizing their expected monetary payoff, are implementing norms of fairness." The corollary of this is that, when in such circumstances a party demands a highly unequal share, skewed in their favour, we exhibit what Ernst Fehr and Klaus Schmidt (1999, 819) call "inequity aversion." We judge such demands to be unfair. But why is that?

Extensive efforts have been made to explain such behaviour. I will not attempt to add to them here. However, I would like to suggest that this inequity aversion is at least consistent with a concern over the abuse of monopoly power. In any case of extreme, *un*fair division, the three criteria for monopoly abuse, laid out above, are very likely to be present.

First, the ultimatum game is inherently monopolistic and monopsonistic since there is only one party on each side of the transaction. Neither the proposer nor the responder face competition. Moreover, it is generally recognised that the proposer – the monopolist – is in much the stronger position.

Second, in the case of an extreme inequality of division, it is clear in this case that the terms are discriminatory, in the sense used above. They do not reflect what might reasonably be expected in an open market, where multiple proposers and responders freely interact. This is probably obvious. After all, few if any responders would accept an egregiously adverse division if there were a near limitless number of other proposers with whom they could transact. There is, though, some further, at least circumstantial, support for this conclusion – if needed – from the field of evolutionary game theory. Although not yet conclusive, increasingly sophisticated computer models do suggest that there is a natural equilibrium, that is, a 'stable market', based in this case on equal division (Alexander 2007, 237).

Finally, where a responder is prepared to accept an extremely adverse proposal, for example 90:10, it seems reasonable to assume that they have no real choice, suggesting that the goods at stake, represented by the '10', must be in some way essential to them. At least that would, I think, be most people's presumption. A die-hard exponent of rational-choice game theory might argue that this reflects extreme rationality, but most of us will feel it indicates, instead, extreme vulnerability.

So, in summary, the three elements behind abuse of monopoly power all seem present in egregiously unfair division: monopoly power, essential goods or services, and discriminatory pricing.

5.7 CONCLUSION

My claim, based on our intuitive judgments, is that a monopolistic negotiation becomes exploitative if, and only if, there is discriminatory pricing of essential goods and services.

With this lens, let us examine again our opening scenario. The Libyan captain exploits the refugee family because he demands of them a discriminatory price for an essential service that only he can offer. Each of these elements is morally permissible standing alone. It is not the fault of the captain that his boat is the last left in the port. Nor is he to blame for the plight of the refugee family that makes their evacuation so essential. And, under normal circumstances, the price the captain chooses to ask for the use of his fishing boat is legitimately for him to decide. But the combination of these three factors changes the moral dynamic. As a monopolistic supplier of an essential service, we intuitively recognise that he has a moral obligation not to employ discriminatory pricing. To do so, is to combine the three elements above in a manner that is an abuse of monopoly power and morally unacceptable.

By the same criteria, the captain's negotiating tactics with the heroin addicts and the occasional cocaine user are *not* themselves exploitative, because, in the first case, the captain is not an exclusive (i.e., monopolistic) supplier and, in the second case, the goods are inessential. In both cases, the buyers are, in a very real sense, free to walk away. They are not vulnerable in negotiating terms.

In sum, to be monopolistically exploitative, a negotiation must have, and need only have, three elements: (1) a monopoly, monopsony, or cartel, (2) an essential good or service, and (3) discriminatory pricing. The combination of a monopoly, monopsony, or cartel with an essential good or service creates extreme asymmetric negotiating leverage and therefore vulnerability. The further combination of that vulnerability with discriminatory pricing entails exploitation – taking unfair advantage. Intuitively, we judge that such monopolistic exploitation renders a negotiation pro tanto unfair, that is, usually morally impermissible. Accordingly, the obligation to avoid monopolistic exploitation represents the third requirement of a fair negotiating process and resulting outcome.

In Part 2, we have reviewed the three necessary constraints required for a fair negotiating process and resulting outcome. Let us take a moment to recap the argument thus far.

As we have seen, negotiation is a feature of the non-ideal world and specifically of imperfect markets. Where there is a perfect market with an unlimited number of buyers and sellers, there is a *single* price, and a single price leaves no room for negotiation. At the opposite extreme, where there is a single seller of an essential good or service, there is again, in theory at least, a single price, being that dictated by the monopolist, and there is again little if any room for negotiation. The practice of negotiation occupies the extensive twilight zone between these two extremes. Where there is a negotiated agreement, there is typically a significant cooperative surplus and, instead of a single price, a *range* of possibilities as to how that cooperative surplus may be divided between the parties (Wertheimer 1996, 61, 139, 217-218). The purpose of negotiation is to determine exactly where in that range agreement is reached.

I have argued that the fairness of such a negotiated agreement is wholly determined by the fairness of the process that leads to that agreement; and that the fairness of the process is itself wholly determined by the adequacy of the constraints applied to that process. A negotiated agreement itself implies some degree of consent. But beyond that I have identified three pro tanto constraints on negotiators.

The first constraint is a rather basic one: that people observe the same standards of honesty that they would in other social interactions. This entails a pro tanto prohibition on lying. The case for this is based on consistency. There is no 'special exemption'. A party is no more entitled to lie in a negotiation than they are to lie in any other context.

The second pro tanto requirement is that the cooperative surplus be shared in some form. It is not acceptable for one party to demand the entire surplus and more, leaving the counterparty with less than none. Where this occurs intentionally, we intuitively judge the counterparty to have been unfairly taken advantage of, whether due to a lack of sufficient rational capacity, gross misinformation and error, conflict of interest, excessive benevolence, or extreme submissiveness.

The third pro tanto requirement is that this division of the cooperative surplus must not reflect an abuse of extreme asymmetric negotiating leverage, that is, a monopoly, monopsony, or cartel combined with essential goods and services, and discriminatory pricing. Again, we intuitively judge such treatment of the vulnerable to be exploitative.

These constraints have been independently derived and might seem unrelated, even ad hoc. However, I have tried to demonstrate, on two grounds, that there is a pattern here, an element of common reasoning. First, these constraints can be seen collectively as an attempt to correct, at least to some degree, the effects of market failure, that is, of non-ideal market conditions. The pro tanto prohibition on lying seeks to eliminate the most egregious cases of information asymmetry. The pro tanto insistence on mutual benefit looks to ensure that the result of any negotiation is a Pareto improvement. And the pro tanto stipulation against the abuse of monopoly aims at eliminating extreme market power and thus promotes negotiating terms that more closely reflect those characteristic of open competition. In all three cases, these constraints drive the negotiated terms closer to those one might expect under the ideal conditions of a perfect market.

Furthermore, each of these three constraints can be seen as a step towards perfecting consent. As I have argued, those who are lied to, grossly misinformed, in evident error, compromised by conflict of interest, beholden to other-regarding concerns based on convention, or the victim of an abusive monopoly may all perhaps be thought to exhibit an imperfect form of consent. By contrast, in refusing to lie, negotiating parties increase the likelihood that they are, in a meaningful way, all consenting to the same deal. By seeking mutually beneficial transactions, they help ensure that the resulting deal is one that all parties could rationally have consented to. And by avoiding the abuse of extreme asymmetric negotiating leverage, they reduce the risk of coercing a vulnerable party that has little choice but to accept the terms demanded. As acknowledged earlier, such an analysis may place too heavy a burden on the meaning of 'consent', but it does at least suggest that the three constraints we have identified are all pulling in the same direction.

These three constraints are each *necessary* for a fair negotiating process and resulting outcome. But are they also jointly *sufficient*? In Part 3, I will argue that they are. In particular, I will seek to refute the argument for any additional constraint on negotiators aimed at promoting distributive justice (Chapter 6).

PART THREE: SUFFICIENT CONSTRAINTS

Chapter 6. The Limits of Obligation

ABSTRACT: Is a fair negotiating process subject to a fourth constraint, one with respect to the broader demands of distributive justice? I argue that it is not, drawing on two important distinctions: (1) a division of moral labour that makes society as a collective whole, not individual negotiators, primarily responsible for distributive justice, and (2) a narrower obligation of beneficence that does place on the individual a duty to aid, but one which stands independently of any negotiation, and falls well short of any wider demands for distributive justice. In exceptional circumstances, this duty to aid can, though, take precedence and require that negotiations cease. It may be objected that, without a broader constraint addressing distributive justice, the morally permissible negotiated outcome remains underdetermined. However, I argue that this is not a weakness, but rather a strength, of the approach. It more closely reflects the character of the negotiating process.

Beyond the three constraints discussed thus far, are there, then, any further constraints that are required to render a negotiating process and resulting outcome fair? I seek to address this question in 6.1, looking in particular at the claims of those who would add a fourth constraint reflecting the broader demands of distributive justice. I argue against such a further constraint by drawing on two significant distinctions: (1) a distinction between the obligations of the individual negotiator and those of society as a collective whole, and (2) a distinction between the individual *qua negotiator* and the individual *qua benefactor*.

In 6.2, I aim to refute a remaining argument in favour of such a fourth constraint, namely that it is required in order to avoid, in moral terms, an underdetermined outcome. I argue that the underdetermined outcome inherent in my approach is not a weakness but rather a strength. It more closely reflects the character of the negotiating process.

6.1 DISTRIBUTIVE JUSTICE

It seems reasonable to assume that any additional constraint would, in some fashion, further determine how the cooperative surplus be fairly shared. The obvious candidate would be some broad measure of distributive justice. We must then address the following question: does a negotiator have a further obligation to take into account the wider circumstances of their counterparty and, if necessary, adjust their negotiating approach to compensate for any material disadvantage on the counterparty's side? Is there, in effect, a fourth constraint – related to the wider demands of distributive justice – that we have yet to acknowledge? I will argue that there is not.

We should be clear, though, as to what we mean by disadvantage. When we talk here of 'disadvantage', we do not refer to those factors, such as an informational disadvantage or a conflicted interest or an unduly benevolent disposition, all of which were discussed at length in Chapter 4. Nor do we mean those conditions of total domination that characterise some relationships between particular employers and their migrant workers. These latter cases constitute monopsony control of an essential need (to sell one's labour) and are covered by our third constraint, discussed in Chapter 5 (see 5.4 in particular). These disadvantages and/or cases of domination relate directly to the negotiation itself and have already been addressed above. Instead, we have in mind background disadvantages that are only indirectly related to the negotiating process itself – most notably, any unjust differences in the wealth and welfare of the parties involved. These differences in wealth and welfare are not the product of any particular ongoing negotiation, but rather the result of past injustice – in many cases, the prior workings of a market driven to achieve efficiency gains rather than deliver distributive justice (Heath 2018b, 26-27, 31).

6.1.1 Division of Moral Labour

In considering this question, I want to draw on two important distinctions. The first is what John Rawls (1993, 268; 2001, 54) calls an "institutional division of labor." We generally view some moral obligations as attaching to us as individuals and others as primarily related to society as a collective whole. As argued earlier (see Chapter 1), a single negotiation cannot be expected to right all the world's wrongs. Such individual negotiations are played out against a backdrop of widespread and often undeserved inequality – what Alan Wertheimer (1996, 299) terms "the injustice of the background conditions." But there must be a limit as to how much responsibility for addressing that inequality can be placed on the individual parties involved in a particular negotiation. In this regard, Rawls (2001, 54; also 1993, 268-269) draws a clear and important distinction between the "principles required to preserve background justice," being the responsibility of those involved in designing society's institutions; and the "principles that apply directly to particular transactions between individuals and associations." Crucially, the latter are a case of pure procedural justice, requiring only that the rules laid down in the "basic structure" of society be observed. The wider task of maintaining "background justice," through for example progressive taxation and the provision of welfare support, is the responsibility of that basic structure and therefore of society as a collective whole. (For other quite similar views see Wertheimer 1996, 216; Maitland 2002, 459-461; Scheffler 2005, 248; Meckled-Garcia 2008, 252-258; Heath 2014, 10; Flanigan 2018, 84; Brennan 2019, 15, 22, 25, Ferretti 2022, 241-253).

There are, I suggest, both principled and practical arguments for this division of moral labour. The principled case recognises that the injustice evident in the background conditions is not the responsibility of any one party – let alone the particular party that might at any given time be engaged in a specific negotiation. Instead, it is the responsibility of many past parties, reflecting factors, such as wrongful enclosure of common lands, unequal educational opportunities, and racial discrimination, that may have occurred over decades or even centuries. To place an onerous obligation on just one negotiator to take responsibility for and address these past wrongs – wrongs for which that negotiator is in no manner individually blameworthy – would appear grossly unjust. The obligation to right these past wrongs can fairly fall only on society as a collective whole.

Ian Maitland illustrates this point well with reference to his study of the pharmaceutical industry. He accepts that people with AIDS have a right to drugs like AZT at an affordable price. But he rejects the implication that Burroughs Wellcome has an obligation to negotiate the sale of AZT at an affordable price. If there is a gap between a fair negotiated price, "playing by the rules" (i.e., procedurally fair), and the price affordable to all, then the responsibility for closing that gap belongs not to the individual negotiating party but to society as a collective. After all, such inequalities in wealth across the population are hardly the fault of any one pharmaceutical company. Thus, "to demand that Burroughs Wellcome shoulder a burden that rightly belongs to all of us is compassion on the cheap. It makes moral free-riders out of all the rest of us." The obligation must fall not on individual actors, but on "the community as a whole" (Maitland 2002, 460; see also Elegido 2009, 40).

In addition to this principled argument, there are at least two practical arguments in defence of the division of moral labour. The first is epistemic. It is unrealistic to assume that individual parties to single negotiations can determine accurately the varying endowments – background advantages or disadvantages – of multiple counterparties. By addressing this at a collective level, through the application of broadly agreed principles, we are collectively able to render the task of delivering distributive justice informationally feasible or, at the very least, significantly less demanding. As Rawls (2001, 54; also 1971, 87-88; 1993, 266-268) puts it, "this allows us to abstract from the enormous complexities of the innumerable transactions of daily life and free us from having to keep track of the changing relative positions of particular individuals" (see also Meckled-Garcia 2008, 253-257; Miklós 2011, 165-166; Ferretti 2022, 254-256). In short, the task is only epistemically manageable at the collective level.

The other practical argument acknowledges the pressure of competition. If an individual enterprise were obliged consistently to accept terms below those readily negotiable, then they would, very likely, soon find themselves competitively disadvantaged and ultimately go out of business, succumbing to Darwinian extinction. The policy would be self-defeating. Reasonably competitive terms are thus essential in order that the "*enterprise may persist in time to perpetuate its benefits to the consumer society*" (Meng 2005, 19; see also Wertheimer 1996, 218; Young 2004, 369; Vogel 2005, 44; Snyder 2008, 398, 400; Ferretti 2022, 252).

The strength of this second practical argument does, of course, depend on the significance of the relevant cost within the overall cost structure of any particular product or service (Vogel 2005, 108-109). For example, critics of sweatshop labour argue that substantial increases in employee wages and welfare would translate into only very modest increases in the overall cost and therefore price of the final garments and footwear, with consequently minimal impact on employment levels (Arnold & Bowie 2003, 238-239; Coakley & Kates 2013, 557-558; Kates 2015, 200-204; Berkey 2021, 52). They observe that, at least in some instances, the end customers are prepared to bear this extra cost, citing for example the purchase of 'fair trade' products (Vogel 2005, 103-106; Elegido 2015, 568) and the use of such informational aids as 'Rugmark' (Vogel 2005, 102-103) and the 'Fashion Transparency Index' (Ferguson & Ostmann 2018, 313). However, such cases would seem to be the exception rather than the rule, leaving "a major gap between what consumers say they would do and their actual behavior" (Vogel 2005, 48, also 93). As a result, supporters of sweatshops warn that insisting on minimum wage levels makes individual manufacturing locations uncompetitive and thus threatens to "unemploy a significant number of workers" (Powell & Zwolinski 2012, 456; see also Zwolinski 2007, 697; Kristof 2009; Sollars & Englander 2018, 16-19; Miklós 2019, 61). A recent study of Indonesian sweatshops seems to offer some support for this view, suggesting a 29% reduction in employment levels resulting from anti-sweatshop activism (Makioka 2021, 650). However, this empirical debate remains largely unresolved, at least as far as it relates to sweatshops (Kates 2015, 195; Bowie 2017, 191; Flanigan 2018, 91-93).

The broader point of the need for competitiveness remains though. While it may be possible, and indeed desirable, to offer better-than-market terms in some marginal cases, a general policy of unilaterally pursuing distributive justice, at the expense of longer-term profitability, would likely prove unsustainable. For example, to raise the living standards of Bangladesh's 4 million garment workers to the level of textile workers in the developed world would almost certainly undermine that country's commercial attractiveness and the competitiveness of any firm based there. Such a transformation would require changes to the basic structure that are well beyond the scope of any individual negotiator (Brennan 2019, 16-17).

This is not to say that those involved in such negotiations owe *no* responsibility for the nature of the background conditions. Each of us bears a responsibility, as members of a wider society, to play our part in the *collective* effort to correct the injustices inherent in the system – a system under which we, in the main, enjoy great advantages. Thus, we each have a share of what has been termed the "responsibility pie" (Dahan, Lerner & Milman-Sivan 2023, 1026). Reflecting this obligation, Iris Young (2006) argues for what she calls a "social connection model" (119). This requires that all those involved in – that is, connected to – a particular trade, exercise a "*political* responsibility" (123) to seek change in the basic structure that addresses the unfairness of the background conditions. But, consistent with Rawls' division of moral labour, this is, as Young says, a "*shared responsibility*" (122). It extends to all those advantaged by such conditions, including for example customers and competitors. This duty derives from the closeness with which all the parties are connected to each other and their mutual commitment to a common system of production and exchange (119). Other more recent commentators mount similar appeals for such a collective obligation to seek change at the political level. Nien-hê Hsieh (2009, 251; also 2017, 305-308) proposes a corporate obligation to promote "well-ordered social and political institutions," Yossi Dahan, Hanna Lerner and Faina Milman-Sivan (2023) argue for a duty to bridge "the 'responsibility gap'" (1029), based on what they call "A Labor Model of Shared Responsibility" (1026), and Thomas Ferretti (2022) makes a case to the AI business community for "political action" (261) in support of "legitimate government intervention" (260).

This broad obligation on us all to play our part in addressing the "structural injustice" (Young 2006, 111) of the society and world order, within which we enjoy great advantage, seems entirely compelling to me. But it is a collective obligation that we all share as members of a society, grounded in a liberal, capitalist system of free markets. It bears no specific relation to the ethical constraints on any particular negotiating process.

However, this division of moral labour that I, and many of those above, are advocating faces potential objections, or at least attempts to blunt its force and application. Juan Elegido

(2015, 568), in particular, draws a distinction between a general obligation to change the unjust background conditions and an obligation not to exploit those injustices in specific cases. The former he accepts is unachievable for any individual party, but the latter he sees differently: "It is one thing," he argues, "not to attempt to make up to every single disadvantaged party with whom one interacts the bad effects of the injustices she suffers, and a very different one to take advantage oneself of these injustices." In a similar vein, Abraham Singer (2018, 113) insists that "managers ought to restrain themselves from profiting off" what he calls "justice failure."

But it is hard to see how this distinction of Elegido's and Singer's works in practice. To return to the case of sweatshop workers, their competitiveness is in large part due to the low wage structure prevalent in their local communities, which is itself arguably a function of the injustice of the background conditions. Singer (2018, 113) admits that the normative implications of his argument are "unavoidably vague and indeterminate." But an obligation not to take advantage of these unjust background conditions – or justice failure – would presumably require, once again, paying a wage that raises the living standards of sweatshop workers to those prevalent in the developed world. As already argued, this would likely leave most of these sweatshop locations uncompetitive. Jeremy Snyder (2008, 393) captures something of this same concern when he warns that to ask employers "to imagine away all of the global injustices that have made their employees poorly educated, largely unskilled, and generally ripe for exploitation, and then set a fair price according to the price that would be reached in this imagined world" would be to place on those employers "an enormous burden […] for rectifying global injustice." No one employer will have the capacity to shoulder such a burden. And the obvious alternative for employers of not engaging at all would likely leave sweatshop workers still further disadvantaged.

András Miklós (2019) offers a similar line of argument to that of Elegido and Singer, with a more precise formulation; but again, this does not fundamentally undermine the case for the division of moral labour made above. Miklós insists that, when negotiating wages, employers are exploitative – "in the sense of taking unfair advantage of an underlying injustice" (60) – if they demand a surplus, "in an unjust state of affairs," that exceeds that they could maximally obtain "in a just state of affairs" (62-64). He defines 'surplus' in this context as "the difference between the employer's willingness to pay and the actual wage paid to the worker" (60). Presumably Miklós has in mind here some measure of profit or marginal contribution.

There is no reason, though, to assume that the employer's surplus, thus described, *is* any greater under a system reliant on unjust background conditions than it would be under a system with just background conditions. In an efficient market, the impact of cheaper labour, reflecting unjust background conditions, is to reduce the cost for all producers. The result of this should be not to inflate the employer's surplus, but rather to (1) competitively drive down the price paid by the end consumer (2) as a result, increase consumption and (3) as a further result, drive up employment levels (and to some extent wages, although these may be determined by the wider market for labour). As an empirical matter, there is evidence that just this has happened in the clothing sector (Vogel 2005, 93-95). Low-cost labour, and the resultant cheap 'disposable' clothing, have fuelled a growth in affordable fast fashion, which has increased overall consumption and driven a further expansion of manufacturing in developing countries.

If the employer enjoys a greater surplus under this system, with its unjust background conditions, that can only imply some level of market *in*efficiency, for example, through the presence of a labour monopsony or cartel. Such a surplus would certainly represent a market failure. But, if Miklós's argument is simply that employers should not drive down wages and generate an 'artificially' high surplus, by exploiting market failures of this kind, then we are in agreement. Indeed, I have already argued in Chapter 5 that the exercise of extreme negotiating leverage (in this case, monopsony power) over essential goods and services (such as employment) to offer discriminatory terms (here, wages below those that would be available in an open market) is morally reprehensible. I have further argued (see 5.6) that the three constraints, that I have proposed, on a fair negotiating process, can in part be seen as attempts to correct for just such market failure.

In sum, then, it is morally impermissible for employers to exploit market failure – in the form of a monopsony or cartel – to drive down the wages and working conditions of employees, made desperate by the structural injustice of their circumstances. But this is to concede no more than has already been argued for in the discussion of monopoly control of essential goods and services in Chapter 5. To insist that individual employers go further than this and not 'take advantage' of such factors as the low wage structure, prevalent in developing economies, would be to threaten their viability and the prospects for their workforce. As Snyder implies above, this low wage structure reflects ingrained economic and social disadvantages such as lack of education, paucity of skills, ill health, weakness of local infrastructure, additional costs imposed through corrupt and inefficient bureaucracies, extended supply chains to/from major markets etc. For individual employers to unilaterally negotiate with employees, as though such disadvantages did not exist, would, if taken to the limit, imply offering them a standard of living equivalent to that of similar workers in the developed world. Few, if any, individual employers could do this and remain competitive. And, even if the rare exception could, would it be either fair or prudent to place this burden on such an employer, simply because their superior performance made such a concession possible?

Any solution to these challenges needs to be based on a collective action. The problem for sweatshop workers is not so much the behaviour of individual employers but rather the structure of the market as a whole. Of course, an individual employer could unilaterally increase wages, but in the absence of other employers doing the same, that would likely mean reducing that employer's surplus *below* that comparable with the maximal surplus available given just background conditions, potentially leading to competitive attrition. It is possible that certain individual employers might have sufficient margins to allow for significantly above-market wages. In a competitive market this would presumably reflect such factors as prior capital investment, scale benefits, greater efficiency, innovative and highly differentiated products, or simply superior management. But to insist that the obligation not to 'take advantage', and instead to pay significantly above-market wages, should fall on these industry leaders alone would be to penalise good performance and, by implication, to subsidise failure. This would distort the natural forces of competition and would seem both imprudent and unjust. The only sustainable way of driving up sweatshop workers' wages would be to regulate labour rates and force up prices to the end consumer. But clearly this would require a collective action and simply underlines the need for the division of moral labour that I am advocating.

I should point out that Miklós himself is not opposed to the idea of a division of moral labour. In prior work (2011), he highlights "the role of the institutions of the basic structure in regulating strategic interactions among individuals" and emphasises "the collective nature" that characterises the "requirements of justice" (171), arguing that "we need political institutions to coordinate individual conduct" (182). He thus acknowledges the significance of the collective action problem that requires a political, or at the very minimum an industry-wide, response to address the injustice of these background conditions.

The appeal, then, to a division of moral labour is compelling. Negotiators, who are in an 'advantaged' position due to background injustices – in the case of sweatshop workers, their employers – find themselves with a *collective action problem*. Either they each must seek the most advantageous terms available, within a relatively narrow set of accepted side constraints, such as those already reviewed above (in Part 2), or they must act together to pursue broader, more diffuse goals of distributive justice. But the latter requires a collective effort that is achievable only at the level of society as a whole – or at least of an industry. It requires, as

Young and others suggest, a political will to address the injustices inherent in the basic structure. This is not a task that can, or should, be left to individual negotiators.

6.1.2 Obligation of Beneficence

However, there are occasions when we feel that individuals involved in negotiation do have a responsibility, as individuals, to take background factors into consideration and this is where the second distinction comes in. Let us recall the case of the Libyan captain in Chapter 5. If the refugee family had no money or possessions at all, would we still expect the captain to help them escape the port? Or would we argue that this was not his responsibility but that of society as a whole? I think the former. We would feel that the captain had an individual obligation to help the family notwithstanding their inability to pay – an obligation grounded in their dire, that is extreme and urgent, need combined with his unique ability to assist (without incurring unreasonable additional cost or risk). This obligation arises not because the captain has acted unjustly, nor because he is materially better off than the family. It arises because the family must be helped, and he alone can help them.

But, as discussed in 5.5.2, this sense of individual obligation would have nothing to do with the captain's role as a *negotiator*. His putative obligation would derive from a duty to aid. This would be the case whether he was involved in a negotiation or not. It would be an entirely separate obligation. If the captain knew for certain that he could get at least $1,000 a head for those five seats in his boat, then the moment he dropped his price for the refugee family below that level, he would no longer be negotiating. He would instead be contemplating a (surely commendable) duty to aid, that is, an act of beneficence.

This, then, is the second important distinction: between the individual obligation of the negotiator *qua negotiator* and the individual obligation of the negotiator *qua benefactor*. Each of us is subject to an array of moral obligations. One of those is to negotiate fairly on the terms of cooperation that we seek with others. But there are further moral obligations that have little to do with such terms of cooperation. They are unilateral obligations, and they include a duty to aid those in extreme and urgent need – a duty of beneficence. Such an obligation exists entirely independently of any negotiation. It involves not a trade, but a gift, albeit a gift that is owed.

There is a potential objection to this clear distinction, an objection that, I believe, should be resisted, but which nonetheless warrants examination. Snyder, in particular, seeks to blur this distinction by tying directly together the "duty of beneficence" and the cooperative relationship (2008, 395-400; 2009b, 306; 2010, 196-200; 2013, 348-350). He argues that the broadly accepted "imperfect duty of aid, charity, or beneficence" is transformed by close "interaction" and "a relationship of use" into a much more specific "perfect form of the duty of beneficence" (2008, 395-396). This duty is made all the more pressing, the closer the relationship and the greater the degree of dependency. The requirements of this enhanced duty of beneficence are quite specific. In the case of sweatshops, "employers are required to cede as much of their benefit from the interaction to their employees as is reasonably possible toward the end of the employees achieving a decent minimum standard of living" (396). Snyder offers, however, clearly proscribed limits to this obligation as to affordability: it does not require such employers to give up a "flourishing" standard of life, although it can require them to give up anything in "excess" of that (398).

Much of this argument seems reasonable. Many would both acknowledge an obligation of beneficence and accept that this obligation is strengthened by proximity and the closeness of any relationship (see, for example, Tadhg Ó Laoghaire's (2022) 'linkage-based dependence account'). This reflects our intuitive tendency to prioritise the child drowning in the nearby pond over a similar child starving on the other side of the world (Singer 1972) – what Maitland (2002, 460) calls "the proximity theory or Good Samaritan theory of moral obligation." But these widely accepted facets all relate to a general obligation that, once again, stands regardless of whether the obligator is party to a particular negotiation or not. This obligation derives from a combination of dire need and the ability (through proximity) to render assistance.

The specific application that Snyder makes of this obligation to the requirements of an employer-employee negotiation is much more contentious. He argues that, if necessary, there must be a sacrifice of "luxury" and "opulence" (Snyder 2008, 398) on the employer's side. But who does Snyder intend by the employer in this instance? Those negotiating the terms of sweatshop labour are likely to be manufacturing subcontractors, locals themselves (Meyers 2004, 329; Berkey 2021, 45, 53), whose living standards will probably fall well short of opulence by any global standard. As for the multinational enterprises ("MNEs") that lie further up the supply chain, they will typically be publicly owned. Their shareholders may largely comprise pension funds, with beneficiaries that include teachers, firefighters, and other modestly rewarded constituencies. Meanwhile the "excess" wealth, if any, of the CEO at head office is unlikely to go very far towards furthering living standards across the company's labour force as a whole. And, anyway, it is not within the power of those negotiating to surrender that putative excess. This seems to breach the principle of 'ought implies can' (Sollars & Englander 2018, 26).

The counter-argument above is an essentially practical one, but there is also a more principled objection to Snyder's proposal. This is well articulated by Benjamin Powell and Matt Zwolinski (2012, 468-469). They observe that, by offering such employment opportunities, MNEs are already conferring a considerable benefit. Of course, *both* sides contribute to creating the cooperative surplus (Kates 2022, 7-11; Faraci 2019, 177). However, even the fiercest critics of sweatshops do acknowledge that sweatshop workers are generally *substantially* better off with than without sweatshop employment – as is evidenced in China by the "tens of thousands of young men and women," who travel each year from rural areas to the cities in search of such employment (Kates 2015, 191). But then, as Powell and Zwolinski (2012, 468) go on to ask, "why should the very act of providing such a benefit impose on MNEs a moral obligation to confer an even greater benefit?" There is, as they say, something implausible about this (Berkey 2021, 54).

The point can be pressed still further. It is reasonable to question why MNEs have such an augmented obligation to their employees, when the worst-off in the communities, where they operate, are most likely those they do not employ! Would they not be better directed to pay their employees no more than the market rate and offer anything over and above that to those to whom they can*not* offer employment (Flanigan 2018, 85-86)? After all, as Thomas Carson (2013, 132) points out, "candidates for sweatshop jobs are rarely at risk of starvation." Indeed, it has been suggested that they may represent something of a "labor aristocracy" (Vogel 2005, 100). By contrast, the very poorest countries in the world do not receive investment from MNEs and their populations face a very real fight for survival.

One might also challenge Snyder's assumption that sweatshop employers are alone responsible for meeting the dependency requirements of their workers. Any individual stands in multiple relationships of dependence, so what about the responsibility of family and of the community? Or of other employers, who could equally well provide sweatshop workers with an income? There seems to be a confusion, in Snyder's argument, between sweatshop workers' right to those goods necessary to meet their basic needs, and an exclusive obligation on the part of particular employers to provide for those needs. The former does not obviously entail the latter. In this respect, Kates (2019, 32), more generally a critic of sweatshops, comes to their defence through analogy. He draws an illuminating parallel with American workers: "from the fact that US citizens have, say, a right to a minimally decent level of health care it most certainly does not follow that employers should be the only ones to provide them with it."

Let us then put Snyder's objections to one side. Underlying all this debate is a general view that sweatshop workers are better off with sweatshops than without them. The widely held belief that "globalism is a better promoter of economic welfare (for the poor)" (Ferguson & Thompson 2021, 116) seems to apply equally to consumers and producers. Consequently, while those negotiating with counterparties in the developing world are certainly under an obligation not to abuse extreme asymmetric negotiating leverage (see Chapter 5), it is in everybody's interest that they enter agreements that are competitively viable. There may be occasions when a duty to aid calls for an act of beneficence, but this cannot sensibly be allowed to undermine that competitive viability.

Importantly, such a duty to aid is, by its nature, exceptional. It is grounded not in a requirement for reciprocation, nor in a belief in the merits of cooperation (negotiated or otherwise), but rather in a strong intuitive sense that, in the particular case of *dire need*, we are obliged unilaterally to offer aid, always assuming that we are in a position to do so without incurring unreasonable cost (note again the parallel with the argument in 5.5.2). Under such circumstances, the negotiator is obliged to cease negotiating – at least, to cease doing so in any substantive sense. The negotiator is no longer permitted to seek the most attractive terms available; nor are they permitted to decline to participate, to walk away. Instead, they must attend to this overriding duty to aid. In Peter Singer's (1972) terms, they must sacrifice their suit and wade into the pond.

It is important to emphasise though how very different this duty to aid is from any broader obligation to pursue distributive justice. It is a response not to inequality or injustice, but instead to dire need. It is best seen not as a broad constraint on the negotiating process, but rather as a competing alternative to that negotiating process, an obligation that on occasion takes precedence over any attempt to maximise one party's interests in contention with another. And, to state perhaps the obvious, this duty to aid those in extreme and urgent need does not of itself imply any further duty to equalise the satisfaction of a wider range of desires and preferences that do not themselves reflect extreme and urgent need.

Notwithstanding the above, it might be objected that the occasions when a negotiator is obliged to cease negotiating and act beneficently are not limited to those of dire need alone but arise also in certain cases of inequality. For example, G.A. Cohen (2000, 134-147; 2008, 129-146) argues for an individual obligation to promote distributive justice, and thus to compensate for any unjust inequality, which for Cohen means virtually any inequality (2000, 124; 2008, 119). He believes this obligation extends beyond the collective responsibility for Rawls' basic structure to the actions of individual parties including, by implication, the bargaining strategy of individual negotiators. Perhaps there are exceptional moments when, in the absence of dire need, an inequality in circumstances is nonetheless so marked that an act of generosity might be called for. However, this would be arguably supererogatory. A *general obligation* on an advantaged party not to negotiate, but instead to unilaterally surrender the cooperative surplus to any less advantaged counterparty, would be unsustainable.

First, such a general principle would create some troublingly *un*intuitive outcomes. For example, there would be almost nothing of a financial nature that a multi-millionaire would be permitted to negotiate. If the multi-millionaire wished to sell a house or a car, they would presumably be required to consider the wealth and welfare of their potential counterparties, determine that negotiation was not permitted due to the likely inequality, and hand over the assets for little or nothing. While I believe most of us would acknowledge a duty to aid in cases of extreme and urgent need, I don't believe that this intuition would extend to such cases of simple inequality, absent such need. An act of generosity in these circumstances would typically be viewed as supererogatory, not obligatory. Furthermore, the charitably minded might ask whether, if such a 'gift' were to be made, it would not be better directed to more *needy* parties outside the negotiation.

Second, in almost all transactions, there is some degree of inequality between the parties. If the better endowed party were always obliged to cease negotiating and to unilaterally surrender the cooperative surplus to their less endowed counterparty, there would be very little room left for negotiation. Instead, the morally permissible outcome to any exchange would be determined by the relative wealth and welfare of the parties. Once again, the process of negotiation would be supplanted by that of judgment (see Chapter 1). And for the reasons articulated earlier – the epistemic challenge of establishing relative wealth and welfare, and the collective action problem in the face of intense competition – it is not easy to see how that judgment might be reached and implemented in a sustainable way. If we remain of the view that negotiation is an important social construct, with a vital role to play in establishing terms of cooperation and thus in securing the gains of trade, then we cannot apply such a tight set of moral constraints that there is little occasion remaining when negotiation is permitted, and no room left to negotiate when it is. In short, a general principle of distributive justice, beyond a simple duty to aid, is a constraint too far.

6.1.3 Implications

So, to conclude, if a wealthy party succeeds in securing a majority share of the cooperative surplus, in negotiation with a much less well-off counterparty, is that fair or unfair? While it may perhaps be unfair in some broader sense, it is doubtful that any such unfairness attaches to the negotiating process itself. The situation might call for a more empathetic, more charitable approach, on the side of the wealthy party, but that would be either supererogatory or reliant on a separate and exceptional obligation to aid others in time of dire need. It could not reasonably be seen as a requirement for a fair *negotiation*. And it would exist regardless of the negotiation. Such an obligation is a function not of any particular transaction, but of the broader relationship between the parties as fellow human beings.

We have then two reasons to doubt whether the fairness of a negotiating process, and therefore of its outcome, should be subject to a requirement of distributive justice. First, there are both principled and pragmatic arguments that suggest such a requirement should not and cannot normally lie with the individual negotiator. Instead, demands for distributive justice call for a collective effort on the part of society as a whole to address the deficiencies in the basic structure. This reflects an important division of moral labour, the need for which is accentuated by a collective action problem.

Second, to the extent that a negotiator is nonetheless individually obliged to unilaterally offer an increased share of the cooperative surplus to a counterparty, this can only be in the exceptional circumstances of dire need and derives from a separate moral imperative that stands independently of the negotiating process itself. The obligation to transfer value in this fashion exists regardless of whether there is a negotiation or not. It is a unilateral obligation of beneficence, a duty to aid. Once a negotiator acknowledges and follows such a duty to aid, they have ceased to negotiate in any meaningful sense of the word.

This may all seem rather theoretical, but let me illustrate these various arguments with the simplest of practical cases:

> *The Grocer*
>
> A grocer at a country market stall negotiates with her customers over the price of her vegetables. Should her negotiating position reflect concerns of distributive justice? At the very margin perhaps. We are all susceptible to the occasional gesture. But let us consider the implications of any substantial attempt to compensate for unjust background conditions. First, the grocer may ask herself what fault she has committed that leaves her obliged to compensate certain of her customers. Then, she will face the challenge of identifying who are the neediest among those customers. This may not be straightforward. Some may seek to maintain appearances, while others may be more willing to come forward for help. And then she may wonder whether the most in need really are those who visit her stall at all. Perhaps instead they are those too poor to visit the market, or those who do not live in her rather affluent county. Might her good

intentions be better served by reaching out somehow to these others? Notwithstanding these reservations, she reduces the price at her stall for those she feels to be less well off. Of course, in a competitive market and in order to remain profitable overall, this requires her to charge her better off customers a price in excess of the market rate. These customers, being either unaware of her charitable aims or unsympathetic with them, start to drift away to her competitors. Eventually she is left with only those unable to pay a market price. Her business has ceased to be viable. How might this situation be corrected? There are two obvious solutions. First, society, in the guise of its political representatives, could step in and insist on an element of progressive taxation to fund welfare support for those otherwise unable to pay a market price for their vegetables. Second, the grocer could take those vegetables that she has left over at the end of the day and donate them to a food bank for any in dire need.

These two solutions are in effect the division of moral labour and the obligation of beneficence. Neither entails that an individual negotiator be morally obliged to reflect distributive justice in their *negotiating* position, although the obligation of beneficence does require on occasion that the negotiator cease negotiating and instead offer aid.

Fundamental to these two distinctions that I have drawn is the nature of the negotiating process as one of pure procedural justice. This marks an important difference from the more foundational values that characterise the basic structure and a duty of beneficence – values such as those of distributive justice. As Benjamin Ferguson (2016b, 966) writes, "we must distinguish fairness as a distributive concept from fairness as a procedural concept." The thrust of my argument is that, when it comes to negotiation, the relevant test of fairness is entirely procedural. If, to the contrary, the division of the cooperative surplus were to be determined strictly by such independent demands as those of distributive justice, then as argued at the outset (see Chapter 1), the process would be more accurately characterised as one of judgment or binding arbitration (imperfect procedural justice), not one of negotiation (pure procedural justice).

This is the case for any formulaic division of the cooperative surplus. If we return to our previous discussion of David Gauthier's (1986, 137) minimax proportional concession, or an even simpler equal division of the cooperative surplus, either will give us a basis for some kind of adjudication. But, as we have seen earlier (see Chapter 1), neither can constitute a negotiation. Instead, such solutions have the character of judgments. The negotiator is left redundant. What is required is a judge. Any outcome is not reached by the parties themselves.

Instead, it is imposed upon them, according to an external benchmark of right and wrong, being the independent basis of the judgment, whatever that might be.

6.2 AN UNDERDETERMINED OUTCOME

In contrast to a judgment, when it comes to negotiating, the demands of fairness do not prescribe a single, specific outcome. Instead, they are reflected in a series of side constraints, encapsulated in the designated procedure. Negotiators are left free to pursue their own goals subject only to the side constraints established above: they may not behave dishonestly and, in particular, may not lie; they may not seek an outcome where their counterparty is left worse off than they would be in the event of a failure to agree, that is, than in the disagreement position; and they may not exploit extreme asymmetric negotiating leverage through a combination of a monopoly, monopsony or cartel with essential goods and services, and discriminatory pricing.

It may be objected that, in most negotiations, the constraints above leave the morally permissible outcome underdetermined. That is, there will be a range of different outcomes that might all be considered fair. But, as we have seen earlier, such a range is in the very nature of a negotiation. In a negotiation there is no fair solution, that stands independently of the process by which it is reached. If there were, there would again be no purpose to the negotiating process. The solution is the product of the process. It is constructed – and *fairly* constructed, when acknowledging adequate constraints. This reflects how we commonly perceive a negotiation, as something clearly distinct from a judgment, which in contrast *is* designed to reveal a single fair outcome based on an independent benchmark of fairness, whatever that should be.

The fact that the approach to negotiation, argued for here, generates not a single fair solution but a range of possible fair solutions is not a weakness but, on the contrary, a strength. It properly captures the character of a negotiation as an exchange between advocates of partially conflicting interests and objectives – importantly, an exchange that has no predetermined outcome. Fairness, in this context, is not a single independent standard, waiting to be discovered. It is instead one of a range of possible outturns, the eventually agreed outturn being the product of a negotiating process, itself judged fair by the criteria, that is by the side constraints, detailed above.

But, if these side constraints do not fully determine the outcome of a negotiation, how *is* that outcome finally determined? The answer can only be through bargaining power and negotiating skill. Let us examine these briefly.

Bargaining power reflects many attributes, central among them the value that a party brings to an agreement: the relative scarcity of the product or service they offer, and the value of that product or service to the counterparty. It does not seem inherently wrong that such considerations should determine, to a large degree, the division of the cooperative surplus. This reflects no more than the standard rules of supply and demand, together with the principle of comparative advantage. If, for example, a company develops a new tool that enables a wide range of customers to manufacture their products at a substantially lower cost or with significantly enhanced capability, it does not seem unreasonable for the innovator to demand a substantial share of the resulting cooperative surplus – that is, as long as the innovator behaves honestly, shares at least some of the cooperative surplus with its counterparties, and does not use monopoly control of essential goods or services to exploit a vulnerable counterparty through discriminatory pricing.

Negotiating skill is perhaps more contentious. But we need to be careful what we mean by negotiating skill. In the popular mind, such negotiating skill is often equated with deviousness, even dishonesty. We should remind ourselves, though, that we have already established a side constraint in this regard: parties to a negotiation are required to meet the same standards of honesty that would be expected of them in other contexts. Specifically, there is a pro tanto prohibition on lying.

In practice, negotiating skill is very largely about creating and communicating value (Fisher, Ury & Patton 2011 [1981]; Shell 2018 [1999]). As we have seen, one of the problems with lying is that it inhibits such creativity and communication and is thus largely self-defeating. In contrast, a skilled negotiator knows that one of the most effective ways of enhancing their own position is to find new sources of value for their counterparty and to clearly communicate that value.

Nonetheless, negotiating skill may strike some as an arbitrary, and therefore arguably an unfair, determinant of a negotiating outcome. Such a skill, whether naturally endowed or advantageously acquired through prior experience, might be seen as "the outcome of the natural lottery" or the result of "historical and social fortune" (Rawls 1971, 74), a function of what has been termed "brute luck" (Dworkin 2000 [1981], 73; Cohen 2011 [1989], 5). Luck egalitarians might therefore argue that negotiating skill should not be allowed to determine the outcome of a negotiation, favouring one party at the expense of another.

However, once again there is here a question of appropriate domain. If there is a legitimate egalitarian concern that some parties are simply less skilled at negotiation than others, then it seems that any practical solution, any attempt at a compensating equalisation, would need to be made at the societal level, not at that of each and every negotiation. Even so, it is not easy to see how this would be done.

Moreover, if the unequal distribution of talents is felt to render any negotiation unfair, then the same must presumably be the case with other professional capabilities and business interactions. If one adopts this latter view, then all managerial skill, rather than being used to create competitive advantage for a particular party, ought instead to be shared among all market participants. In other words, if we are to be consistent, then an egalitarian insistence that negotiating skill should not be permitted to render a competitive advantage to one party must presumably entail the same for other business skills such as strategy formulation, team building, product development, component sourcing etc. All these endowed skills and talents would need to be shared in some way. This would represent a radically different economic model to that currently based on open market *competition*. It seems then that, if we remain committed to free markets and open competition, we must accept elements of inequality in these interactions, including those involving negotiation.

6.3 CONCLUSION

There is then good reason to believe that the fairness of an agreement, reached with the consent of all parties, can be ensured if the negotiating process is subject to three side constraints, being an obligation on the parties to (1) behave honestly (2) seek mutually beneficial outcomes, and (3) avoid the abuse of extreme asymmetric negotiating leverage. Such side constraints are both necessary and *sufficient* conditions for a fair process and resulting outcome.

In particular, there is no requirement for a further condition of distributive justice. To the extent that such an obligation exists at all, it generally rests with society as a collective whole, and only rarely with the individual negotiator. It is a response to an obligation to address a collective action problem through contributing to the shaping of appropriate institutions, and abiding by them once they are in place. On those exceptional occasions that direct responsibility does fall on an individual, it takes the form of an obligation of beneficence, or duty to aid, that is clearly distinct from any constraints that legitimately apply to the negotiating process itself. Under such circumstances the individual concerned ceases to be a negotiator and becomes instead a benefactor.

It may be objected that in the absence of any further constraints, the morally permissible negotiated outcome remains underdetermined. However, this is not a weakness, but rather a strength, of the account. It more closely reflects the character of the negotiating process.

Chapter 7. Conclusion

I have argued for the following:

(1) Negotiation is a *social construct* whereby the parties seek to resolve differences and promote cooperation through a process rooted in *mutual consent*.

(2) Negotiation is a feature of *imperfect* markets. Where there is a negotiated agreement, there is typically a significant cooperative surplus and, instead of a single market price, a range of possibilities as to how that cooperative surplus may be divided between the parties. The purpose of negotiation is to determine exactly where in that range agreement is reached.

(3) The fairness of a negotiated agreement is wholly determined by the fairness of the process that leads to that agreement; and the fairness of the process is itself wholly determined by the adequacy of the constraints applied to that process. This makes negotiation an instance of *pure procedural justice*. In this regard, negotiation is to be clearly distinguished from judgment or binding arbitration.

(4) Negotiated agreement itself implies some degree of consent. But beyond that there are pro tanto obligations on negotiators to (a) *behave honestly* and, in particular, not lie, (b) *seek mutually beneficial outcomes*, whereby their counterparties are left no worse off than they would be in the event of a failure to agree, and (c) *avoid the abuse of extreme asymmetric negotiating leverage* that characterises a combination of monopoly, monopsony or cartel with essential goods and services, and discriminatory pricing.

(5) These constraints are both *necessary* and *sufficient* to deliver a fair process and, therefore, a fair outcome.

(6) Such constraints can be seen, in part, as attempts to counter *market failure* and compensate for *shortfalls in consent*.

(7) A further constraint, related to the broader demands of *distributive justice*, is *not* called for. Such a requirement is the primary responsibility of society as a whole. To the extent that it does fall on any individual, it takes the form of a separate and narrower duty to aid that can, in cases of dire need, take precedence and require that negotiations cease.

(8) The fact that these moral constraints leave the outcome of most negotiations less than fully determined is not a weakness, but a strength of this account. It properly captures the character of a negotiation as a cooperative but also partially adversarial process,

during the course of which a fair outcome is *constructed* – importantly, an outcome that is *not predetermined*.

I have argued for these positions based on the consistent application of widely shared intuitive reasoning. Nonetheless, I suspect that, many will disagree with these conclusions. Those that disagree most violently will typically fall into one of two opposing camps. From the left, egalitarians will object to the lack of any constraint tied to distributive justice. With this in mind, they will seek to further constrain the negotiating process, at the extreme to the point where there is just one single 'fair', 'equitable' solution. But, as I have argued, this amounts to abandoning the negotiation process altogether and to substituting for it a judgment. This is at best a utopian ideal, because in most instances no such judgment is available and acceptable to all. It requires either an unusual level of acquiescence on the part of all concerned or some power of enforcement. In the real world there is commonly no alternative other than to negotiate.

From the right, the most committed libertarians will pull in the opposite direction. They will seek to reduce the constraints advocated here and thus to increase still further the set of morally permissible outcomes. In particular, they may reject the obligation to consider the interests of a counterparty and, perhaps also, the prohibition on exploiting extreme negotiating leverage. In effect, they will argue for something near a 'free for all'. But, in doing so, they must justify some very counter-intuitive outcomes. For example, they must explain the implied moral permissibility of the husband, taking advantage of a patriarchal norm to negotiate from his wife one of her kidneys; and of the Libyan captain, exploiting the plight of the refugee family.

What I have sought to do here is to steer a course between these two extremes. I consider the moral character of negotiation to be neither totally prescriptive nor totally permissive. Instead, it is best seen as constrictive. In a typical negotiation there is a range of fair outcomes, that range being determined by a number of ethical constraints. Those constraints require us to treat each other honestly, to seek each other's mutual benefit, and to refrain from exploiting another's extreme vulnerability.

In all three cases, these constraints move us towards a more ideal world. They each help, at least to a degree, to rectify market failure and to restore an otherwise compromised consent. In so doing, they promote, within the negotiating setting, many of those qualities that we associate with the ideal of a perfect market. The prohibition on lying enhances informational transparency (Chapter 3); win/win negotiating entails Pareto improvement

(Chapter 4); and the constraint on monopoly power engenders terms closer to those we might expect from open market competition (Chapter 5). Thus, these constraints help deliver an outcome, in a negotiating context, closer to that commonly associated with the implicit morality of a free and efficient market.

The same pattern is evident when it comes to consent. An honest negotiation increases the likelihood that the various parties are, in a meaningful way, all consenting to the same deal (Chapter 3); a mutually beneficial transaction helps ensure that the resulting deal is one that all parties could rationally have consented to (Chapter 4); and the prohibition on the abuse of extreme asymmetric negotiating leverage reduces the risk of coercing a vulnerable party that has little choice but to accept the terms demanded (Chapter 5). In this way, all three constraints bolster the consent that underpins the very concept of a free market.

But, while this common tendency towards a theoretical ideal gives these constraints an important coherence, it should not distract us from the intensely practical objective. Negotiation exists because it is useful, indeed essential. It is a socially constructed process, an instance of pure procedural justice. It is no more, and no less. As such, any arguments concerning its moral permissibility need to be practical, that is, action-guiding. Negotiators need to know what they are morally permitted or not permitted to do. So, let us express these moral imperatives once more in the simplest possible terms: negotiators should behave honestly with each other, they should seek an outcome that is beneficial for all parties, and they should not exploit another party's extreme vulnerability.

That does not seem much to ask. But, I think, that is also all that one can reasonably ask. At least, all that one can reasonably ask in the context of a negotiation.

REFERENCES

Abeler, Johannes, Daniele Nosenzo & Collin Raymond (2019) "Preferences for Truth-Telling" *Econometrica 87*(4):1115-1153.

Adler, Robert (2007) "Negotiating with Liars" *MIT Sloan Management Review 48*(4):69-74.

Aguire, Iñaki, Simon Cowan & John Vickers (2010) "Monopoly Price Discrimination and Demand Curvature" *The American Economic Review 100*(4):1601-1615.

Akerlof, George (1970) "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism" *The Quarterly Journal of Economics 84*(3):488-500.

Alexander, Jason (2007) *The Structural Evolution of Morality*. Cambridge: Cambridge University Press.

Alexander, Jason & Brian Skyrms (1999) "Bargaining with Neighbours: Is Justice Contagious?" *The Journal of Philosophy 96*(11):588-598.

Allhoff, Fritz (2003) "Business Bluffing Reconsidered" *Journal of Business Ethics 45*(4):283-289.

Anderson, Elizabeth (1999) "What is the Point of Equality?" *Ethics 109*(2):287-337.

Applbaum, Arthur (1999) *Ethics for Adversaries: The Morality of Roles in Public and Professional Life.* Princeton, NJ: Princeton University Press.

Aquino, Karl (1998) "The Effects of Ethical Climate and the Availability of Alternatives on the Use of Deception During Negotiation" *International Journal of Conflict 9*(3):195-217.

Aquino, Karl & Thomas Becker (2005) "Lying in Negotiations: How Individual and Statistical Factors Influence the Use of Neutralization Strategies" *Journal of Organizational Behavior 26*(6):661-679.

Ariely, Dan (2009) [2008] *Predictably Irrational.* Rev. and exp. ed. London: Thorsons.

Arneson, Richard (2013) "Exploitation and Outcome" *Politics, Philosophy & Economics 12*(4):392-412.

Arnold, Denis & Norman Bowie (2003) "Sweatshops and Respect for Person" *Business Ethics Quarterly 13*(2):221-242.

Arrow, Kenneth (1951) *Social Choice and Individual Values.* Mansfield, CT: John Wiley & Sons.

Arrow, Kenneth (1973) "Social Responsibility and Economic Efficiency" *Public Policy 21*:303-317.

Augustine, St. (2016) [395] "De Mendacio" (tr. H. Browne) in Kevin DeLapp & Jeremy Henkel (eds.) *Lying and Truthfulness*, 6-35. Indianapolis, IN: Hackett.

Badaracco, Joseph Jr. (1997) *Defining Moments: When Managers Must Choose Between Right and Right.* Boston: Harvard Business School Press.

Badaracco, Joseph Jr. (2002) *Leading Quietly: An Unorthodox Guide to Doing the Right Thing*. Boston: Harvard Business School Press.

Bailey, Adam (2011) "The Nonworseness Claim and the Moral Permissibility of Better-Than-Permissible Acts" *Philosophia 39*:237-250.

Banai, Moshe, Abraham Stefanidis, Ana Shetach & Mehmet Ferhat Özbek (2014) "Attitudes Toward Ethically Questionable Negotiation Tactics: A Two-Country Study" *Journal of Business Ethics Vol 123*(4):669-685.

Barnes, John (1994) *A Pack of Lies: Towards a Sociology of Lying*. New York: Cambridge University Press.

Barry, Brian (1989) *A Treatise on Social Justice Vol I: Theories of Justice.* Los Angeles, CA: University of California Press.

Barry, Brian (1995) *A Treatise on Social Justice Vol II: Justice as Impartiality.* Oxford: Oxford University Press.

Beach, John (1985) "Bluffing: Its Demise as a Subject unto Itself" *Journal of Business Ethics 4*(3):191-196.

Bebchuk, Lucian & Jesse Fried (2003) "Executive Compensation as an Agency Problem" *The Journal of Economic Perspectives 17*(3):71-92.

Becker, Lawrence (1980) "Reciprocity and Social Obligation" *Pacific Philosophical Quarterly 61*(4):411-421.

Berkey, Brian (2020) "The Value of Fairness and the Wrong of Wage Exploitation" *Business Ethics Quarterly 30*(3):414-429.

Berkey, Brian (2021) "Sweatshops, Structural Injustice, and the Wrong of Exploitation: Why Multinational Corporations Have Positive Duties to the Global Poor" *Journal of Business Ethics 169*:43-56.

Bernheim, Douglas (2021) "In Defense of Behavioral Welfare Economics" *Journal of Economic Methodology 28*(4):385-400.

Betrand, Marianne & Sendhil Mullainathan (2001) "Are CEOs Rewarded for Luck? The Ones Without Principals Are" *The Quarterly Journal of Economics 116*(3):901-932.

Bicchieri, Cristina (2010) "Norms, Preferences, and Conditional Behaviour" *Politics, Philosophy & Economics 9*(3):297-313.

Binmore, Ken (1994) *Game Theory and the Social Contract Vol 1: Playing Fair.* Cambridge, MA: The MIT Press.

Binmore, Ken (1998) *Game Theory and the Social Contract Vol 2: Just Playing*. Cambridge, MA: The MIT Press.

Binmore, Ken, Joe Swierzbinski, Steven Hsu & Chris Proulx (1993) "Focal Points and Bargaining" *International Journal of Game Theory 22*:381-409.

Blodgett, Timothy (1968) "Showdown on 'Business Bluffing'" *Harvard Business Review* May-June 1968:162-170.

Blosser Reed, Homer (1916) "The Morals of Monopoly and Competition" *International Journal of Ethics 26*(2):258-281.

Boatright, John (1992) "Morality in Practice: Dees, Cramton, and Brer Rabbit on a Problem of Applied Ethics" *Business Ethics Quarterly 2*(1):63-73.

Boero, Riccardo, Giangiacomo Bravo, Marco Castellani & Flaminio Squazzoni (2009) "Reputational Cues in Repeated Trust Games" *The Journal of Socio-Economics 38*:871-877.

Bok, Sissela (1999) [1978] *Lying: Moral Choice in Public and Private Life.* 2nd Vintage Books ed. New York: Vintage Books.

Boles, Terry, Rachel Croson & Keith Murnighan (2000) "Deception and Retribution in Repeated Ultimatum Bargaining" *Organizational Behavior and Human Decision Processes 83*(2):235-259.

Bond, Charles & Bella DePaulo (2006) "Accuracy of Deception Judgments" *Personality and Social Psychology Review 10*(3):214-234.

Bowie, Norman (1985) "Should Collective Bargaining and Labor Relations Be Less Adversarial?" *Journal of Business Ethics 4*(4):283-291.

Bowie, Norman (2017) *Business Ethics: A Kantian Perspective.* 2nd Ed. New York: Cambridge University Press.

Bowie, Norman & R. Edward Freeman (1992) "Ethics and Agency Theory: An Introduction" in Norman Bowie & R. Edward Freeman (eds.) *Ethics and Agency Theory: An Introduction,* 3-22. New York: Oxford University Press.

Braithwaite, Richard (1955) *Theory of Games as a tool for the Moral Philosopher.* Cambridge: Cambridge University Press.

Brandts, Jordi & Gary Charness (2003) "Truth or Consequences: An Experiment" *Management Science 49*(1):116-130.

Brenkert, George (1998) "Marketing and the Vulnerable" *Business Ethics Quarterly: Ruffin Series in Business Ethics*:7-20.

Brennan, Jason (2011) *The Ethics of Voting.* Princeton, NJ: Princeton University Press.

Brennan, Jason (2019) "Should Employers Pay a Loving Wage?" *Journal of Business Ethics 157*(1):15-26

Brennan, Jason & Peter Jaworski (2016) *Markets Without Limits: Moral Virtues and Commercial Interests.* New York: Routledge.

Brown, Campbell (2011) "Consequentialize This" *Ethics 121*(4):749-771.

Brown, Campbell (2016) "The Rightest Theory of Degrees of Rightness" *Ethical Theory and Moral Practice 19*(1):21-29.

Brown, Campbell (2020) "Sex Crimes and Misdemeanours" *Philosophical Studies 177*:1363-1379.

Bruni, Luigino & Robert Sugden (2008) "Fraternity: Why the Market Need Not Be a Morally Free Zone" *Economics and Philosophy 24*:35-64.

Bruni, Luigino & Robert Sugden (2013) "Reclaiming Virtue Ethics for Economics" *The Journal of Economic Perspectives 27*(4):141-163.

Buchanan, Allen (1990) "Justice as Reciprocity versus Subject-Centred Justice" *Philosophy & Public Affairs 19*(3):227-252.

Burr, Anne (2001) "Ethics in Negotiations: Does Getting to Yes Require Candor?" *Dispute Resolution Journal* May-July 2001:9-15.

Burri, Susanne, Daniela Lup & Alexander Pepper (2021) "What Do Business Executives Think About Distributive Justice?" *Journal of Business Ethics* 174:15-33.

Carr, Albert (1968) "Is Bluffing in Business Ethical?" *Harvard Business Review* Jan 1968:143-153.

Carson, Thomas (1993) "Second Thoughts on Bluffing" *Business Ethics Quarterly 3*(4):317-341.

Carson, Thomas (2005) "The Morality of Bluffing: A Reply to Allhof" *Journal of Business Ethics 56*(4):399-403.

Carson, Thomas (2006) "The Definition of Lying" *Nous 40*(2):284-306.

Carson, Thomas (2010a) "Deception and Information Disclosure in Business and Professional Ethics" in George Brenkert & Tom Beauchamp (eds.) *The Oxford Handbook of Business Ethics*, 335-365. New York: Oxford University Press.

Carson, Thomas (2010b) *Lying and Deception: Theory and Practice.* Oxford: Oxford University Press.

Carson, Thomas (2013) "Free Exchange for Mutual Benefit: Sweatshops and Maitland's 'Classical Liberal Standard'" *Journal of Business Ethics 112*:127-135.

Carson, Thomas, Richard Wokutch & Kent Murrmann (1982) "Bluffing in Labor Negotiations: Legal and Ethical Issues" *Journal of Business Ethics 1*(1):13-22.

Chisholm, Roderick & Thomas Feehan (1977) "The Intent to Deceive" *The Journal of Philosophy 74*(3):143-159.

Coakley, Matthew & Michael Kates (2013) "The Ethical and Economic Case for Sweatshop Regulation" *Journal of Business Ethics 117*(3):553-558.

Cohen, Gerald (1995) *Self-Ownership, Freedom and Equality.* Cambridge: Cambridge University Press.

Cohen, Gerald (2000) *If You're an Egalitarian, How Come You're So Rich?* Cambridge, MA: Harvard University Press.

Cohen, Gerald (2008) *Rescuing Justice and Equality.* Cambridge, MA: Harvard University Press.

Cohen, Gerald (2009) *Why Not Socialism?* Princeton, NJ: Princeton University Press.

Cohen, Gerald (2011) *On the Currency of Egalitarian Justice and Other Essays in Political Philosophy.* Princeton, NJ: Princeton University Press.

Cohen, Jonathan (2002) "The Ethics of Respect in Negotiation" *Negotiation Journal* April 2002:115-120.

Cohen, Tara, Erik Helzer & Robert Creo (2022) "Honesty Among Lawyers: Moral Character, Game Framing, and Honest Disclosures in Negotiations" *Negotiation Journal: Research Report*:1-36.

Cooper, Ricky, Michael Davis & Ben Van Vliet (2016) "The Mysterious Ethics of High-Frequency Trading" *Business Ethics Quarterly 26*(1):1-22.

Croson, Rachel, Terry Boles & Keith Murnighan (2003) "Cheap Talk in Bargaining Experiments: Lying and Threats in Ultimatum Games" *Journal of Economic Behavior & Organization 51*:143-159.

Curtis, Charles (1951) "The Ethics of Advocacy" *Stanford Law Review 4*(1):3-23.

Dahan, Yossi, Hanna Lerner & Faina Milman-Sivan (2023) "Shared Responsibility and Labor Rights in Global Supply Chains" *Journal of Business Ethics 182*:1025-1040.

Darwall, Stephen (1977) "Two Kinds of Respect" *Ethics 88*(1):36-49.

Davis, Nathaniel (2016) "If War Can Have Ethics, Wall Street Can, Too" *New York Times,* October 3.

Dees, Gregory & Peter Cramton (1991) "Shrewd Bargaining on the Moral Frontier: Toward a Theory of Morality in Practice" *Business Ethics Quarterly 1*(2):135-167.

Dees, Gregory & Peter Cramton (1995) "Deception and Mutual Trust: A Reply to Strudler" *Business Ethics Quarterly 5*(4):823-832.

DeGeorge, Richard (1992) "Agency Theory and the Ethics of Agency" in Norman Bowie & Edward Freeman (eds.) *Ethics and Agency Theory*, 59-72. Oxford: Oxford University Press.

DeMarco, C. Wesley (2001) "Knee Deep in Technique: The Ethics of Monopoly Capital" *Journal of Business Ethics 31*(2):151-164.

Dougherty, Tom (2013) "Sex, Lies, and Consent" *Ethics 123*(4):717-744.

Dunham, Yarrow, Allison Durkin & Tom Tyler (2018) "The Development of a Preference for Procedural Justice for Self and Others" *Scientific Reports 8*(17740):1-8.

Dworkin, Ronald (2000) *Sovereign Virtue: The Theory and Practice of Equality.* Cambridge, MA: Harvard University Press.

Ekman, Paul & Maureen O'Sullivan (1991) "Who Can Catch a Liar?" *American Psychologist 46*(9):913-920.

Elegido, Juan (2009) "The Just Price: Three Insights from the Salamanca School" *Journal of Business Ethics 90*(1):29-46.

Elegido, Juan (2011) "The Effects of Price Discrimination" *Business Ethics Quarterly 21*(4):633-660.

Elegido, Juan (2015) "The Just Price as the Price Obtainable in an Open Market" *Journal of Business Ethics 130*(3):557-572.

Erat, Sanjiv & Uri Gneezy (2012) "White Lies" *Management Science 58*(4):723-733.

Fallis, Don (2009) "What Is Lying?" *The Journal of Philosophy 106*(1):29-56.

Faraci, David (2019) "Wage Exploitation and the Nonworseness Claim: Allowing the Wrong, To Do More Good" *Business Ethics Quarterly 29*(2):169-188.

Fehr, Ernst & Klaus Schmidt (1999) "A Theory of Fairness, Competition, and Cooperation" *The Quarterly Journal of Economics 114*(3):817-868.

Feinberg, Joel (1984) *The Moral Limits of the Criminal Law Volume 1: Harm to Others.* New York: Oxford University Press.

Feinberg, Joel (1990) *The Moral Limits of the Criminal Law Volume 4: Harmless Wrongdoing.* New York: Oxford University Press.

Ferguson, Benjamin (2016a) "Exploitation and Disadvantage" *Economics and Philosophy 32*:485-509.

Ferguson, Benjamin (2016b) “The Paradox of Exploitation” *Erkenntnis 81*:951-972.

Ferguson, Benjamin (2018) “Can Libertarians Get Away with Fraud?” *Economics and Philosophy 34*:165-184.

Ferguson, Benjamin (2021) “Are We All Exploiters?” *Philosophical and Phenomenological Research 103*:535-546.

Ferguson, Benjamin & Florian Ostmann (2018) “Sweatshops and Consumer Choices” *Economics and Philosophy 34*:295-315.

Ferguson, Benjamin & Hillel Steiner (2018) “Exploitation” in Serena Olsaretti (ed.) *The Oxford Handbook of Distributive Justice,* 533-555. Oxford: Oxford University Press.

Ferguson, Benjamin & Chris Thompson (2021) “Why Buy Local?” *Journal of Applied Philosophy 38*(1):104-120.

Ferretti, Thomas (2022) “An Institutionalist Approach to AI Ethics: Justifying the Priority of Government Regulation over Self-Regulation” *Moral Philosophy and Politics 9*(2):239-265.

Fisher, Roger, William Ury & Bruce Patton (2011) [1981] *Getting to Yes: Negotiating an Agreement without Giving In*. 3rd rev. ed. New York: Random House.

Flanigan, Jessica (2018) “Sweatshop Regulation and Workers’ Choices” *Journal of Business Ethics 153*:79-94.

Fletcher, Guy (2016) “Objective List Theories” in Guy Fletcher (ed.) *The Routledge Handbook of Philosophy of Well-Being,* 148-160. Milton Park: Routledge.

Frankfurt, Harry (1987) “Equality as a Moral Ideal” *Ethics 98*(1):21-43.

Frankfurt, Harry (2005) [1986] *On Bullshit.* Princeton, NJ: Princeton University Press.

Freeman, R. Edward (1984) *Strategic Management: A Stakeholder Approach*. Boston: Pitman.

Friedman, Milton (1962) *Capitalism and Freedom*. Chicago: University of Chicago Press.

Friedman, Milton (1970) “The Social Responsibility of Business Is to Increase Profits” *New York Times Magazine,* September 13.

Friedman, Raymond (1992) “From Theory to Practice: Critical Choices for ‘Mutual Gains’ Training” *Negotiation Journal* April 1992:91-98.

Friedman, Raymond (1993) “Bringing Mutual Gains Bargaining to Labor Negotiations: The Role of Trust, Understanding, and Control” *Human Resources Management 32*(4):435-459.

Friedman, Raymond & Debra Shapiro (1995) “Deception and Mutual Gains Bargaining: Are They Mutually Exclusive?” *Negotiation Journal* July 1995:243-253.

Gaspar, Joseph, Emma Levine & Maurice Schweitzer (2015) “Why We Should Lie” *Organizational Dynamics 44*:306-309.

Gaspar, Joseph, Redona Methasani & Maurice Schweitzer (2019) “Fifty Shades of Deception: Characteristics and Consequences of Lying in Negotiations” *Academy of Management Perspectives 33*(1):62-81.

Gaspar, Joseph & Maurice Schweitzer (2013) “The Emotion Deception Model: A Review of Deception in Negotiation and the Role of Emotion in Deception” *Negotiation and Conflict Management Research 6*(3):160-179.

Gauthier, David (1986) *Morals by Agreement.* Oxford: Clarendon Press.

Gerlach, Philipp, Kinneret Teodorescu & Ralph Hertwig (2019) “The Truth About Lies: A Meta-Analysis on Dishonest Behavior” *American Psychological Association Psychological Bulletin 145*(1):1-44.

Gibbard, Allan (1990) *Wise Choices, Apt Feelings: A Theory of Normative Judgment.* Oxford: Clarendon Press.

Gibbard, Allan (1991) “Constructing Justice” *Philosophy & Public Affairs 20*(3):264-279.

Gneezy, Uri (2005) “The Role of Consequences” *The American Economic Review 95*(1):384-394.

Goodpaster, Kenneth (1991) “Business Ethics and Stakeholder Analysis” *Business Ethics Quarterly 1*(1):53-73.

Gorlin, Rena (ed.) (1990) [1986] *Codes of Professional Responsibility.* 2nd ed. Washington, DC: The Bureau of National Affairs.

Guidice, Rebecca, G. Stoney Alder & Steven Phelan (2009) “Competitive Bluffing: An Examination of a Common Practice and Its Relationship with Performance” *Journal of Business Ethics 87*(4):535-553.

Gunia, Brian (2019) “Ethics in Negotiation: Causes and Consequences” *Academy of Management Perspectives 33*(1):3-11.

Gunia, Brian & Emma Levine (2019) “Deception as Competence: The Effect of Occupational Stereotypes on the Perception and Proliferation of Deception” *Organizational Behavior and Human Decision Processes 152*:122-137.

Guth, Werner, Rolf Schmittberger & Bernd Schwarze (1982) “An Experimental Analysis of Ultimatum Bargaining” *Journal of Economic behavior and Organization 3*:367-388.

Hare, R.M. (1992) “One Philosopher’s Approach to Business and Professional Ethics” *Business and Professional Ethics Journal 11*(2):3-19.

Harris, Sam (2013) *Lying.* Opelousas, LA: Four Elephants Press.

Harsanyi, John (1955) “Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility” *Journal of Political Economy 63*(4):309-321.

Harsanyi, John (1956) “Approaches to the Bargaining Problem Before and After the Theory of Games: A Critical Discussion of Zeuthen’s, Hicks’, and Nash’s Theories” *Econometrica 24*(2):144-157.

Hausman, Daniel (1989) “Are Markets Morally Free Zones?” *Philosophy & Public Affairs 18*(4):317-333.

Hausman, Daniel (1995) “The Impossibility of Interpersonal Utility Comparisons” *Mind 104*(415):473-490.

Hayek, Friedrich (1960) *The Constitution of Liberty.* Oxford: Routledge Classics.

Heap, Shaun Hargreaves & Yanis Varoufakis (1995) *Game Theory: A Critical Text.* London: Routledge.

Heath, Joseph (2014) *Morality, Competition, and the Firm: The Market Failures Approach to Business Ethics.* Oxford: Oxford University Press.

Heath, Joseph (2018a) “’But Everyone Else Is Doing It’: Competition and Business Self-Regulation” *Journal of Social Philosophy 49*(4):516-535.

Heath, Joseph (2018b) “On the Very Idea of a Just Wage” *Erasmus Journal for Philosophy and Economics 11*(2):1-33.

Heath, Joseph (2019) “Is the ‘Point’ of the Market Pareto or Kaldor-Hicks Efficiency?” *Business Ethics Journal Review 7*(4):21-26.

Heath, Joseph (2020) *The Machinery of Government: Public Administration and the Liberal State.* New York: Oxford University Press.

Heilmann, Conrad & Stefan Wintein (2017) “How to be Fairer” *Synthese 194*:3475-3499.

Heinrich, Joseph, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis & Richard McElreath (2004) *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies.* Oxford Scholarship Online.

Hersch, Gil (2020) “You Can Bluff but You Should Not Spoof” *Business & Professional Ethics Journal 39*(2):207-224.

Hess, David & Thomas Dunfee (2007) “The Kasky-Nike Threat to Corporate Social Reporting: Implementing a Standard of Optimal Truthful Disclosure as a Solution” *Business Ethics Quarterly 17*(1):5-32.

Hollander-Blumoff, Rebecca (2010) “Just Negotiation” *Washington University Law Review 88*(2):381-432.

Hollander-Blumoff, Rebecca (2017) "Formation of Procedural Justice Judgments in Legal Negotiation" *Group Decision & Negotiation 26*:19-43.

Hollander-Blumoff, Rebecca, & Tom Tyler (2008) "Procedural Justice in Negotiation: Procedural Fairness, Outcome Acceptance, and Integrative Potential" *Law & Society Inquiry 33*(2):473-500.

Holley, David (1998) "Information Disclosure in Sales" *Journal of Business Ethics 17*(6):631-641.

Horton, Joe (2019) "The Exploitation Problem" *The Journal of Political Philosophy 27*(4):469-479.

Hsieh, Nien-hê (2009) "Does Global Business Have a Responsibility to Promote Just Institutions?" *Business Ethics Quarterly 19*(2):251-273.

Hsieh, Nien-hê (2017) "The Responsibilities and Role of Business in Relation to Society: Back to Basics?" *Business Ethics Quarterly 27*(2):293-314.

Huynh, Toan, Marc Rieger & Mei Wang (2022) "Cross-Country Comparison in Dishonest Behaviour: Germany and East Asian Countries" *Economics Letters 215*:1-5.

Hunter, Larry & Robert McKersie (1992) "Can 'Mutual Gains' Training Change Labor-Management Relationships?" *Negotiation Journal October 1992*:319-330.

Jevons, William (1879) [1871] *The Theory of Political Economy.* 2nd ed. London: MacMillan.

John, Leslie (2016) "Managing Yourself: How to Negotiate with a Liar" *Harvard Business Review* July-August 2016:114-117.

Kalai, Ehud & Meir Smorodinsky (1975) "Other Solutions to Nash's Bargaining Problem" *Econometrica 43*(3):513-518.

Kant, Immanuel (1917) [1795] *Perpetual Peace* (tr. M. Campbell Smith). 3rd ed. London: George Allen & Unwin.

Kant, Immanuel (2016) [1797] "On a Supposed Right to Lie because of Philanthropic Concerns" (tr. James Ellington) in Kevin DeLapp & Jeremy Henkel (eds.) *Lying and Truthfulness*, 81-86. Indianapolis, IN: Hackett.

Kates, Michael (2015) "The Ethics of Sweatshops and the Limits of Choice" *Business Ethics Quarterly 25*(2):191-212.

Kates, Michael (2019) "Sweatshops, Exploitation, and the Case for a Fair Wage" *The Journal of Political Philosophy 27*(1):26-47.

Kates, Michael (2022) "Sweatshops, Exploitation, and the Nonworseness Claim" *Business Ethics Quarterly* First View: 1-22.

Kersten, Gregory, Wojtek Michalowski, Stan Matwin & Stan Szpakowicz (1988) “Representing the Negotiation Process with a Rule-Based Formalism” *Theory & Decision 25*:225-257.

Khorasanee, Kasim (2022) “Spoof, Bluff, Go For It: A Defence of Spoofing” *Journal of Business Ethics* online.

Koehn, Daryl (1997) “Business and Game-Playing: The False Analogy” *Journal of Business Ethics 16*(12/13):1447-1452.

Koplin, Julian (2018) “Choice, Pressure and Markets in Kidneys” *Journal of Medical Ethics 44*:310-313.

Korsgaard, Christine (1996) *Creating the Kingdom of Ends.* Cambridge: Cambridge University Press.

Kramer, Roderick (2009) “Rethinking Trust” *Harvard Business Review,* June, 69-77.

Kristof, Nicholas (2009) “Where Sweatshops are a Dream” *New York Times,* January 15.

Kugelberg, Elsa (2021) “Responsibility for Reality: Social Norms and the Value of Constrained Choice” *Politics, Philosophy & Economics 20*(4):357-384.

Kurki, Visa (2018) “Rights, Harming and Wronging: A Restatement of the Interest Theory” *Oxford Journal of Legal Studies 38*(3):430-450.

Lawlor, Rob (2011) “Organ Sales Needn’t Be Exploitative (But It Matters If They Are)” *Bioethics 25*(5):250-259.

Lawlor, Rob (2014) “Organ Sales: Exploitative at Any Price?” *Bioethics 28*(4):194-202.

Levine, Emma & Maurice Schweitzer (2014) “Are Liars Ethical? On the Tension Between Benevolence and Honesty” *Journal of Experimental Social Psychology 53*:107-117.

Levine, Emma & Maurice Schweitzer (2015) “Prosocial Lies: When Deception Breeds Trust” *Organizational Behavior and Human Decision Processes 126*:88-106.

Lewicki, Roy (1983) “Lying and Deception: A Behavioral Model” in Max Bazerman & Roy Lewicki (eds.) *Negotiating in Organisations,* 68-90. Beverly Hills, CA: Sage Publications.

Lewicki, Roy, Bruce Barry & David Saunders (2020) *Negotiation.* 8th International Edition. New York: McGraw Hill Education.

Liberto, Haillie (2014) “Exploitation and the Vulnerability Clause” *Ethical Theory and Moral Practice 17*(4):619-629.

Lipsey, Richard & Kelvin Lancaster (1956-57) “The General Theory of the Second Best” *The Review of Economic Studies 24*(1):11-32.

Lucas, John (1959) “Moralists and Gamesmen” *Philosophy 34*(128):1-11.

Luce, Duncan & Howard Raiffa (1957) *Games and Decisions: Introduction & Critical Survey*. New York: Dover Publications.

Ma, Zhenzhong (2010) "The SINS in Business Negotiation: Explore the Cross-Cultural Differences in Business Ethics Between Canada and China" *Journal of Business Ethics 91*(Supplement 1):123-135.

Ma, Zhenzhong, Dapeng Liang & Honghui Chen (2013) "Negotiating with the Chinese: Are They More Likely to Use Unethical Strategies?" *Group Decision and Negotiation 22*:641-655.

MacIntyre, Alasdair (1994) "Truthfulness, Lies, and Moral Philosophers: What Can We Learn from Mill and Kant?" *Tanner Lectures on Human Values*, April 6 & 7. Princeton NJ: Princeton University.

MacLeod, Alistair (1984) "Distributive Justice, Contract, and Equality" *The Journal of Philosophy 81*(11):709-718.

Mahon, James (2015) [2008] "The Definition of Lying and Deception" *Stanford Encyclopedia of Philosophy*.

Maitland, Ian (2002) "Priceless Goods: How Should Life-Saving Drugs Be Priced? *Business Ethics Quarterly 12*(4):451-480.

Makioka, Ryo (2021) "The Impact of Anti-Sweatshop Activism on Employment" *Review of Development Economics 25*:630-653.

Malhotra, Deepak & Max Bazerman (2007) *Negotiation Genius*. New York: Bantam Books.

Malmqvist, Erik (2014) "Are Bans on Kidney Sales Unjustifiably Paternalistic?" *Bioethics 28*(3):110-118.

Malmqvist, Erik (2015) "Kidney Sales and the Analogy with Dangerous Employment" *Health Care Analysis 23*:107-121.

Malmqvist, Erik (2017) "Better to Exploit than to Neglect? International Clinical Research and the Non-Worseness Claim" *Journal of Applied Psychology 34*(4):474-488.

Manne, Kate (2019) [2018] *Down Girl: The Logic of Misogyny*. London: Penguin.

Marcoux, Alexei (2006) "Much Ado About Price Discrimination" *Journal of Markets & Morality 9*(1):57-69.

McCullough, Colleen (2016) "Unconscionability as a Coherent Legal Concept" *University of Pennsylvania Law Review 164*(3):779-825.

McMahon, Christopher (1981) "Morality and the Invisible Hand" *Philosophy & Public Affairs 10*(3):247-277.

Meckled-Garcia, Saladin (2008) “On the Very Idea of Cosmopolitan Justice: Constructivism and International Agency” *The Journal of Political Philosophy 16*(3):245-271.

Meng, Jude (2005) “What Profits for a Man to Gain: Just (the) Price (of the Soul)” *Journal of Markets & Morality 8*(1):7-26.

Meyer, Marco & Chun Wei Choo (2023) “Harming by Deceit: Epistemic Malevolence and Organizational Wrongdoing” *Journal of Business Ethics* Feb:1-14.

Meyers, Chris (2004) “Wrongful Beneficence: Exploitation and Third World Sweatshops” *Journal of Social Philosophy 35*(3):319-333.

Michel, Christian (1999) “What Is a ‘Just Price’?” *Journal of Markets & Morality 2*(2):182-196.

Michelman, James (1983) “Deception in Commercial Negotiation” *Journal of Business Ethics 2*(4):255-262.

Miklós, András (2011) “The Basic Structure and the Principles of Justice” *Utilitas 23*(2):161-182.

Miklós, András (2019) “Exploiting Injustice in Mutually Beneficial Market Exchange: The Case of Sweatshop Labor” *Journal of Business Ethics 156*:59-69.

Mitchell, Vincent-Wayne & Joseph Chan (2002) “Investigating UK Consumers’ Unethical Attitudes and Behaviours” *Journal of Marketing Management 18*:5-26.

Moehler, Michael (2018) *Minimal Morality: A Multilevel Social Contract Theory.* Oxford: Oxford University Press.

Moriarty, Jeffrey (2009) “How Much Compensation Can CEOs Permissibly Accept?” *Business Ethics Quarterly 19*(2):235-250.

Morse, Lily & Taya Cohen (2019) “Moral Character in Negotiation” *Academy of Management Perspectives 33*(1):12-25.

Murnighan, Keith, Linda Babcock, Leigh Thompson & Madan Pillutla (1999) “The Information Dilemma in Negotiations: Effects of Experience, Incentives, and Integrative Potential” *The International Journal of Conflict Management 10*(4):313-339.

Nagel, Thomas (1981) [1975] “Libertarianism Without Foundations” in Jeffrey Paul (ed.) *Reading Nozick: Essays on Anarchy, State, and Utopia*, 191-205. Oxford: Blackwell.

Nagel, Thomas (1986) *The View from Nowhere.* New York: Oxford University Press.

Nash, John (1950) “The Bargaining Problem” *Econometrica 18*(2):155-162.

Nicholson, Walter & Christopher Snyder (2017) [2012] *Microeconomic Theory: Basic Principles and Extensions.* Andover: Cengage.

Nieto, Enrique Cáceres (2021) “The Foundations of Legal Constructivism” in Jorge Luis Fabra-Zamora & Gonzalo Villa Rosas (eds.) *Conceptual Jurisprudence,* 295-319. Law and Philosophy Library 137. Switzerland: Springer.

Norman, Wayne (2011) “Business Ethics as Self-Regulation: Why Principles that Ground Regulations Should Be Used to Ground Beyond-Compliance Norms as Well” *Journal of Business Ethics 102*(Supplement 1):43-57.

Nozick, Robert (1974) *Anarchy, State, and Utopia.* Oxford: Blackwell.

Nussbaum, Martha (2011) *Creating Capabilities: The Human Development Approach.* Cambridge, MA, and London: The Belknap Press of Harvard University Press.

Nyerges, Janos (1987) “Ten Commandments for a Negotiator” *Negotiation Journal* January 1987:21-27.

O’Connor, Kathleen & Peter Carnevale (1997) “A Nasty but Effective Negotiation Strategy: Misrepresentation of a Common-Value Issue” *Personality and Social Psychology Bulletin 23*:504-515.

Ó Laoghaire, Tadhg (2022) “Why (Some) Corporations Have Positive Duties to (Some of) the Global Poor” *Journal of Business Ethics* June 7.

Olekalns, Mara, Carol Kulik & Lin Chew (2014) “Sweet Little Lies: Social Context and the Use of Deception in Negotiation” *Journal of Business Ethics 120*(1):13-26.

Otsuka, Michael (2003) *Libertarianism Without Inequality.* Oxford: Clarendon Press.

Pakula, Tom, dir. (1982) *Sophie’s Choice.* Universal City CA: Universal Pictures.

Palmer, David & Trevor Hedberg (2013) “The Ethics of Marketing to Vulnerable Populations” *Journal of Business Ethics 116*(2):403-413.

Parfit, Derek (2011a) *On What Matters: Volume One.* Oxford: Oxford University Press.

Parfit, Derek (2011b) *On What Matters: Volume Two.* Oxford: Oxford University Press.

Pepper, Alexander (2022) *If You’re So Ethical, Why Are You So Highly Paid?* London: LSE Press.

Peppet, Scott (2002) “Can Saints Negotiate? A Brief Introduction to the Problems of Perfect Ethics in Bargaining” *7 Harvard Negotiation Law Review 83*:83-96.

Perrone, Andrea (2014) “The Just Price Doctrine and Contemporary Contract Law: Some Introductory Remarks” *Rivista Internazionale di Scienze Sociali 122*(2):217-236.

Pettit, Philip (1995) “The Virtual Reality of *Homo Economicus*” *The Monist 78*(3):308-329.

Piker, Andrew (2002) “Ethical Immunity in Business: A Response to Two Arguments” *Journal of Business Ethics 36*:337-346.

Post, Frederick (1990) “Collaborative Collective Bargaining: Toward an Ethically Defensible Approach to Labor Negotiations” *Journal of Business Ethics 9*(6):495-508.

Powell, Benjamin & Matt Zwolinski (2012) “The Ethical and Economic Case Against Sweatshop Labor: A Critical Assessment” *Journal of Business Ethics 107*(4):449-472.

Preiss, Joshua (2019) “Freedom, Autonomy, and Harm in Global Supply Chains” *Journal of Business Ethics 160*:881-891.

Preston, Ivan (1994) *The Tangled Web They Weave: Truth, Falsity & Advertisers.* Madison: University of Wisconsin Press.

Provis, Chris (2000a) “Ethics, Deception and Labor Negotiation” *Journal of Business Ethics 28*(2):145-158.

Provis, Chris (2000b) “Honesty in Negotiation” *Business Ethics: A European Review 9*(1):3-12.

Pruitt, Dean (1983) “Achieving Integrated Agreements” in Max Bazerman & Roy Lewicki (eds.) *Negotiation in Organizations,* 35-50. Beverly Hills, CA: Sage Publications.

Quinlan, Michael (1993) “Ethics in the Public Service” *Governance: An International Journal of Policy and Administration 6*(4):538-558.

Quinn, Dennis & Thomas Jones (1995) “An Agent Morality View of Business Policy” *The Academy of Management Review 20*(1):22-42.

Raiffa, Howard (1953) “Arbitration Schemes for Generalized Two-Person Games” in H. Kuhn & A.W. Tucker (eds.) *Contributions to the Theory of Games II,* 361-387. Princeton NJ: Princeton University Press.

Rao, Asha & Stuart Schmidt (1998) “A Behavioral Perspective on Negotiating International Alliances” *Journal of International Business Studies 29*(4):665-693.

Rawls, John (1971) *A Theory of Justice.* Oxford: Oxford University Press.

Rawls, John (1993) *Political Liberalism.* Exp ed. New York: Columbia University Press.

Rawls, John (2001) *Justice as Fairness: A Restatement.* Cambridge, MA: Harvard University Press.

Rees, McKenzie, Ann Tenbrunsel & Max Bazerman (2019) “Bounded Ethicality and Ethical Fading in Negotiations: Understanding Unintended Unethical Behavior” *Academy of Management Perspectives 33*(1):26-42.

Reid, Leopold, Stephanie Eckerd, Lutz Kaufmann & Craig Carter (2023) “From Target to Actor: Contagion of Honesty and Deception Across Buyer-Supplier Negotiations” *Journal of Operations Management 69*:261-283.

Reitz, Joseph, James Wall & Mary Sue Love (1998) "Ethics in Negotiation: Oil and Water or Good Lubrication?" *Business Horizons* May-June 1998:5-14.

Richard, Stanislas (2021) "Wage Exploitation as Disequilibrium Price" *Business Ethics Quarterly*:1-25 (first view: 23 Dec).

Rippon, Simon (2014) "Imposing Options on People in Poverty: The Harm of a Live Donor Organ Market" *Journal of Medical Ethics 40*(3):145-150.

Rivers, Cheryl & Anne Louise Lytle (2007) "Lying, Cheating Foreigners!! Negotiation Ethics Across Cultures" *International Negotiation 12*:1-28.

Rivers, Cheryl & Roger Volkema (2013) "East-West Differences in 'Tricky' Tactics: A Comparison of the Tactical Preferences of Chinese and Australian Negotiators" *Journal of Business Ethics 115*(1):17-31.

Robeyns, Ingrid & Morten Byskov (2020) [2011] "The Capability Approach" *Stanford Encyclopedia of Philosophy.*

Rogers, Todd, Richard Zeckhauser, Francesco Gino & Michael Norton (2017) "Artful Paltering: The Risks and Rewards of Using Truthful Statements to Mislead Others" *Journal of Personality and Social Psychology 112*(3):456-473.

Rutledge, John (1996) "The Portrait on My Office Wall" *Forbes*, December 30.

Ryan, Cheyney (1981) [1977] "Yours, Mine, and Ours: Property Rights and Individual Liberty" in Jeffrey Paul (ed.) *Reading Nozick: Essays on Anarchy, State, and Utopia,* 323-343. Oxford: Blackwell.

Satz, Debra (2010) *Why Some Things Should Not Be for Sale: The Moral Limits of Markets.* New York: Oxford University Press.

Saul, Jennifer (2012) *Lying, Misleading and What Is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford: Oxford University Press.

Scanlon, Thomas (1991) "The Moral Basis of Interpersonal Comparisons" in Jon Elster & John Roemer (eds.) *Interpersonal Comparisons of Well-Being,* 17-44. Cambridge: Cambridge University Press.

Scheffler, Samuel (2005) "The Division of Moral Labour" *Proceedings of the Aristotelian Society 79*:229-253.

Schelling, Thomas (1980) [1960] *The Strategy of Conflict.* 2nd ed. Cambridge, MA: Harvard University Press.

Schneider, Andrea (2000) "Perception, Reputation and Reality: An Empirical Study of Negotiation Skills" *Dispute Resolution Magazine* Summer 2000:24-28.

Schumpeter, Joseph (1976) [1942] *Capitalism, Socialism and Democracy.* 5th ed. London: George Allen & Unwin.

Schweitzer, Maurice & Rachel Croson (1999) “Curtailing Deception: The Impact of Direct Questions on Lies and Omissions” *The International Journal of Conflict Management 10*(3):225-248.

Schweitzer, Maurice, John Hershey & Eric Bradlow (2006) “Promises and Lies: Restoring Violated Trust” *Organizational Behavior and Human Decision Processes 101*:1-19.

Sen, Amartya (1970) *Collective Choice and Social Welfare.* Amsterdam: North-Holland Publishing Company.

Sen, Amartya (1977) “Rational Fools: A Critique of the Behavioral Foundations of Economic Theory” *Philosophy & Public Affairs 6*(4):317-344.

Sen, Amartya (1982) *Choice, Welfare and Measurement.* Cambridge, MA: Harvard University Press.

Serota, Kim, Timothy Levine & Tony Docan-Morgan (2022) “Unpacking Variation in Lie Prevalence: Prolific Liars, Bad Lie Days, or Both?” *Communication Monographs 89*(3):307-331.

Sethi-Iyengar, Sheena, Gur Huberman & Gur Jiang (2004) “How Much Choice is Too Much? Contributions to 401(k) Retirement Plans” in Olivia Mitchell & Stephen Utkus (eds.) *Pension Design and Structure: New Lessons from Behavioral Finance*, 83-96. Oxford: Oxford University Press.

Shadyac, Tom, dir. (1997) *Liar, Liar.* United States: Universal Studios.

Shapiro, Debra & Robert Bies (1994) “Threats, Bluffs, and Disclaimers in Negotiations” *Organizational Behavior and Human Decision Processes 60*:14-35.

Shell, Richard (1991) “When Is It Legal to Lie in Negotiations?” *Sloan Management Review 32*(3):93-101.

Shell, Richard (2018) [1999] *Bargaining for Advantage: Negotiation Strategies for Reasonable People.* 3rd ed. New York: Penguin.

Sherwood, Charles (2022) “A Lie Is a Lie: The Ethics of Lying in Business Negotiations” *Business Ethics Quarterly 32*(4):604-634.

Singer, Abraham (2018) “Justice Failure: Efficiency and Equality in Business Ethics” *Journal of Business Ethics 149*:97-115.

Singer, Peter (1972) “Famine, Affluence, and Morality” *Philosophy and Public Affairs 1*(3):229-243.

Singer, Peter (2009) *The Life You Can Save*. London: Picador.

Skyrms, Brian (2014) [1996] *Evolution of the Social Contract.* 2nd ed. Cambridge: Cambridge University Press.

Smith, Adam (2009) [1759] *The Theory of Moral Sentiments.* London: Penguin.

Smith, Adam (2012) [1776] *Wealth of Nations.* Ware, Hertfordshire: Wordsworth.

Smith, Greg (2012) “Why I Am Leaving Goldman Sachs” *New York Times*, March 14.

Snyder, Jeremy (2008) “Needs Exploitation” *Ethical Theory and Moral Practice 11*(4):389-405.

Snyder, Jeremy (2009a) “What’s the Matter with Price Gouging?” *Business Ethics Quarterly 19*(2):275-293.

Snyder, Jeremy (2009b) “Efficiency, Equity, and Price Gouging: A Response to Zwolinski” *Business Ethics Quarterly 19*(2):303-306.

Snyder, Jeremy (2010) “Exploitation and Sweatshop Labor: Perspectives and Issues” *Business Ethics Quarterly 20*(2):187-213.

Snyder, Jeremy (2013) “Exploitation and Demeaning Choices” *Politics, Philosophy & Economics 12*(4):345-360.

Sobral, Filipe & Gazi Islam (2013) “Ethically Questionable Negotiating: The Interactive Effects of Trust, Competitiveness, and Situation Favorability on Ethical Decision Making” *Journal of Business Ethics 117*(2):281-296.

Sollars, Gordon & Fred Englander (2018) “Sweatshops: Economic Analysis and Exploitation as Unfairness” *Journal of Business Ethics 149*(1):15-29.

Sorensen, Roy (2007) “Bald-Faced Lies! Lying Without the Intent to Deceive” *Pacific Philosophical Quarterly 88*:251-264.

Steiner, Hillel (2017) “Exploitation, Intentionality and Injustice” *Economics and Philosophy 34*:369-379.

Street, Sharon (2006) “A Darwinian Dilemma for Realist Theories of Value” *Philosophical Studies 127*:109-166.

Strudler, Alan (1995) “On the Ethics of Deception in Negotiation” *Business Ethics Quarterly 5*(4):805-822.

Strudler, Alan (2005) “Deception Unraveled” *The Journal of Philosophy 102*(9):458-473

Strudler, Alan (2010) “The Distinctive Wrong in Lying” *Ethical Theory and Moral Practice 13*(2):171-179.

Strudler, Alan (2023) “Lying about Reservation Prices in Business Negotiation: A Qualified Defense” *Business Ethics Quarterly 33*(4):763-776.

Styron, William (2004) [1979] *Sophie’s Choice.* London: Vintage.

Sugden, Robert (2018) *The Community of Advantage: A Behavioural Economist's Defence of the Market*. Oxford: Oxford University Press.

Sullivan, Roger (1984) "A Response to 'Is Business Bluffing Ethical?'" *Business & Professional Ethics Journal 3*(2):1-21,23.

Sunstein, Cass (2021) *Liars: Falsehoods and Free Speech in an Age of Deception.* New York: Oxford University Press.

Sunstein, Cass & Richard Thaler (2005) "Libertarian Paternalism" in Cass Sunstein (ed.) *Laws of Fear: Beyond the Precautionary Principle*, 175-203. Cambridge; Cambridge University Press.

Tenbrunsel, Ann (1998) "Misrepresentation and Expectations of Misrepresentation in an Ethical Dilemma: The Role of Incentives and Temptation" *The Academy of Management Journal 41*(3):330-339.

Tenbrunsel, Ann & David Messick (2004) "Ethical Fading: The Role of Self-Deception in Unethical Behavior" *Social Justice Research 17*(2):223-236.

Thaler, Richard & Cass Sunstein (2008) *Nudge: Improving Decisions about Health, Wealth and Happiness.* London: Penguin.

Thoma, Johanna (2015) "Bargaining and the Impartiality of the Social Contract" *Philosophical Studies 172*(12):3335-3355.

Triandis, Harry, Peter Carnevale, Michele Gelfand, Christopher Robert, Arzu Wasti, Tahira Probst, Emiko Kashima, Thalia Dragonas, Darius Chan, Xiao Ping Chen, Uichol Kim, Carsten de Dreu, Evert van de Vliert, Sumiko Iwao, Ken-Ichi Ohbuchi & Paul Schmitz (2001) "Culture and Deception in Business Negotiations: A Multilevel Analysis" *International Journal of Cross Cultural Management 1*(1):73-90.

Tyler, James, Robert Feldman & Andreas Reichert (2006) "The Price of Deceptive Behavior: Disliking and Lying to People Who Lie to Us" *Journal of Experimental Social Psychology 42*:69-77.

US Holocaust Memorial Museum (2021) "Nuremberg Trials" *Holocaust Encyclopedia* https://encyclopedia.ushmm.org/content/en/article/the-nuremberg-trails.

Valdman, Mikhail (2009) "A Theory of Wrongful Exploitation" *Philosophers' Imprint 9*(6):1-14.

Velleman, David (2015) *Beyond Price: Essays on Birth and Death.* Cambridge: Open Book Publishers.

Vogel, David (2005) *The Market for Virtue: The Potential and Limits of Corporate Social Responsibility.* Washington D.C.: Brookings Institution Press.

Volkema, Roger (2004) “Demographic, Cultural, and Economic Predictors of Perceived Ethicality of Negotiation Behavior: A Nine-Country Analysis” *Journal of Business Research 57*:69-78.

Volkema, Roger, Denise Fleck & Agnes Hofmeister-Toth (2004) “Ethicality in Negotiation: An Analysis of Attitudes, Intentions, and Outcomes” *International Negotiation 9*:315-339.

von Neumann, John & Oskar Morgenstern (1944) *Theory of Games and Economic Behavior*. Princeton NJ: Princeton University Press.

Voorhoeve, Alex (2009) *Conversations on Ethics.* Oxford: Oxford University Press.

Voss, Chris (2022) [2016] *Never Split the Difference: Negotiating as if Life Depended on It.* London: Penguin.

Vrousalis, Nicholas (2013) “Exploitation, Vulnerability, and Social Domination” *Philosophy & Public Affairs 41*(2):131-157.

Vrousalis, Nicholas (2021) “How Exploiters Dominate” *Review of Social Economy 79*(1):103-130.

Walsh, Adrian (2014) “Commentary on Simon Rippon, ‘Imposing Options on People in Poverty: The Harm of a Live Donor Organ Market’” *Journal of Medical Ethics 40*(3):153-154.

Walzer, Michael (1983) *Spheres of Justice: A Defense of Pluralism and Equality.* New York: Basic Books.

Weber, Max (1992) [1904-5] *The Protestant Ethic and the Spirit of Capitalism*. London & New York: Routledge.

Welsh, Nancy (2004) “Perceptions of Fairness in Negotiation” *Marquette Law Review 87*(4):753-767.

Wertheimer, Alan (1992) “Two Questions About Surrogacy and Exploitation” *Philosophy & Public Affairs 21*(3):211-239.

Wertheimer, Alan (1996) *Exploitation.* Princeton, NJ: Princeton University Press.

Wesley, John (1831) [1785] “Sermon 90: An Israelite Indeed” in John Emory (ed.) *The Works of the Reverend John Wesley, A.M.* New York: J. Emory and B. Waugh, for the Methodist Episcopal Church.

Wilkinson, Stephen (2003) *Bodies for Sale.* Abingdon and New York: Routledge.

Williams, Bernard (1981) [1975] “The Minimal State” in Jeffrey Paul (ed.) *Reading Nozick: Essays on Anarchy, State, and Utopia,* 27-36. Oxford: Blackwell.

Williams, Bernard (2002) *Truth & Truthfulness: An Essay in Genealogy.* Princeton, NJ: Princeton University Press.

Wolff, Jonathan (1991) *Robert Nozick: Property, Justice and the Minimal State.* Cambridge: Polity Press.

Wood, Allen (1995) "Exploitation" *Social Philosophy and Policy 12*:136-158.

Yaari, Menahem & Maya Bar-Hillel (1984) "On Dividing Justly" *Social Choice and Welfare 1*:1-24.

Young, Iris (2004) "Responsibility and Global Justice" *The Journal of Political Philosophy 12*(4):365-388.

Young, Iris (2006) "Responsibility and Global Justice: A Social Connection Model" *Social Philosophy & Policy 23*(1):102-130.

Zhang, Jian-Dong, Leigh Anne Liu & Wu Liu (2015) "Trust and Deception in Negotiation: Culturally Divergent Effects" *Management and Organization Review 11*(1):123-144.

Zwolinski, Matt (2007) "Sweatshops, Choice, and Exploitation" *Business Ethics Quarterly 17*(4):689-727.

Zwolinski, Matt (2008) "The Ethics of Price Gouging" *Business Ethics Quarterly 18*(3):347-378.

Zwolinski, Matt, Ben Ferguson & Alan Wertheimer (2022) "Exploitation" *Stanford Encyclopedia of Philosophy.*