

Coercion and Consent

by

James Konow*

Most moral justifications for coercion have been based on one of two arguments: the consent of the coerced, usually understood as univariate and discrete, or the beneficial consequences of coercion; but many cases do not fit these categories. This paper proposes that consent be understood as our inferences about the agreement of agents, which vary in fine degrees with multiple underlying factors, including agent discretion, the choice set, information, and competence. Moreover, consent interacts with other moral values, including consequences, in a pluralistic system of morals that depends on the context. Examples suggest this framework can be reconciled with moral intuitions better than rival systems. (JEL: B3, D6, D7, H1, K4)

1 Introduction

Consent has been characterized as “moral magic” (e.g., Alexander, 1996). An act of sex, a transfer of resources, a blow to the face, or a set of laws is viewed very differently depending on the consent of the person acted on. It is the difference between rape and lovemaking, a theft and a gift, battery and a boxing match, or an autocracy and a democracy. One of the starkest expressions of the strength of this preference is the sacrifice millions are willing to make to provide input to the constraints imposed on them by the state. Voters sometimes wait in line for hours, and millions, across the globe, have died in movements to be able to cast a ballot that is almost certain not to be decisive. The primacy of consent is so compelling, it seems difficult to reconcile violations of it with our moral intuition.

And yet there is abundant evidence that we, in both the private and public spheres of our lives, not only tolerate but routinely endorse coercive measures that forgo consent, and these examples seem equally persuasive. Drivers are expected to stop at every red light, taxpayers are not exempt from funding public programs of which they do not approve, and the taking of private property by the state is often sanctioned,

* Chair of Economics and Ethics, University of Kiel and Kiel Institute for the World Economy, and Professor of Economics, Loyola Marymount University. I am very grateful for the helpful comments and suggestions of Christoph Engel, Dirk Engelmann, Leo Katz, Michael Kurschilgen, Urs Schweizer, Christoph Vanberg, and Eyal Zamir, which resulted in many important changes to the paper. Of course, they are absolved of all responsibility for any remaining shortcomings.

e.g., by eminent domain laws in the United States. Sometimes there are narrow but deep divides about coercion and consent: the individual mandate of the Obama administration's Affordable Care Act closely split both the American public and the United States Supreme Court. What is one to make of these conflicting, and seemingly irreconcilable, intuitions and practices?

This paper presents a framework to account for, and provide guidance on, the justifiable use of coercion and the obligation to secure consent. I propose that consent is constructed out of multiple underlying factors and varies in fine degrees, but that one must also consider other moral principles. The paper takes, as its point of departure, the seminal contributions by John Locke (1689/2003) and David Hume (1752/2007) on the source of moral obligation to obey a government, and, like those authors, I seek an account that is consistent with broadly shared moral sentiments.¹ Locke grounded this on the consent of the governed, whereas Hume made a consequences-based, or *consequentialist*, appeal to the general interests and needs of society. Many, if not most, of the arguments on this topic that have followed over the ages have similarly revolved around consent or consequences. Additional moral considerations (e.g., justice) are, of course, relevant, but I focus here on consequences and consent for the sake of brevity, continuity with previous work, and their importance for a wide range of contemporaneous issues.

I begin by clarifying briefly certain terms and concepts. Coercion is defined as intentionally compelling someone to a choice at the threat of imposing a penalty for noncompliance. There are several points that warrant explanation here. First, coercion is intentional: the coercer seeks to compel, so a driver who accidentally runs a red light and forces another driver to swerve to avoid collision is not guilty of coercion (although the offender might be guilty of many other things, including a misdemeanor). Second, for coercion, the cost that is threatened must involve something to which the victim is entitled. For instance, you are entitled to shop without my interference, so it is coercion if I prohibit you from that. But you are not entitled to a certain price for a car, so the car salesman who declines your suggested price has not coerced. Third, coercion must be effective, i.e., it must compel someone to a choice that, in the absence of the threat, the person would not have taken. This also suggests that the penalty involves something of value to the victim that outweighs the cost of accommodating the demand.²

¹ On this very point with respect to this very topic, Hume wrote “[t]hough an appeal to general opinion may justly, in the speculative sciences of metaphysics, natural philosophy, or astronomy, be deemed unfair and inconclusive, yet in all questions with regard to morals, as well as criticism, there is really no other standard, by which any controversy can ever be decided. And nothing is clearer proof, that a theory of this kind is erroneous, than to find, that it leads to paradoxes repugnant to the common sentiments of mankind, and to the practice and opinion of all nations and all ages” (p. 516).

² But the impact on choice is distinct from the threat itself, which may also be judged in moral terms even when it is ineffective and, therefore, not coercion: a gangster's threat of violence to businesses if they do not pay “protection money” is not coercion if businesses choose to ignore the threat, but the threat itself is still not nice.

At issue are situations in which coercion is a potential or actual force, e.g., military service, sexual relations, and taxation, but not (typically) eating lunch, surfing the Internet, or visiting a friend. Consent involves the waiving of rights or entitlements of some kind, e.g., to one's body, freedom, or property. Specifically, the focus here is on consensual acts that grant others permission to constrain one's own actions. But this sounds like an oxymoron: how can one be willing to be forced to act against one's will? To make sense of this, think of coercion that applies beyond one individual at one point in time. Coercion can generate individual benefits, for example, if it constrains not only oneself but also a larger group, as with mandated cooperation to provide a public good. Coercion can also be desired, even when applied individually, if it solves intertemporal choice problems, e.g., tax withholding at a higher rate or Christmas-club accounts that constrain current consumption and increase savings. These issues are not, however, the chief concern of this paper. Rather, I focus on trying to improve our understanding of consent as a moral value as well as the moral grounds for compromising it.

Intervention signifies interference, that is, coming between two things, such as people and their goals or desires; so state intervention, which compels choices, is a form of coercion. Many of the arguments that follow are fleshed out in the context of state coercion, but analogous reasoning is applicable to private coercion, that is, the kind exercised by nonstate entities (independent of whether it is also supported or prohibited by the state). Such coercion may be viewed as justified, as with most parental authority over children, or unjustified, as with murder or theft. Nevertheless, the domain of morally justified private coercion, or at least the set of interesting and contestable cases involving private coercion, is considerably more circumscribed than that for public coercion, so, in fact, greater attention will be paid here to the latter.

A related question is that of the legitimacy of the state, and Copp (1999) points to two meanings: descriptive and normative. In the descriptive sense, it can refer either to the acknowledgment of governments, as in Hume (1752/2007), or to *psychological legitimacy*, i.e., popular acceptance of the government's right to exercise power, as in Tyler (2006). In the normative (or prescriptive) sense, it refers to moral authority and inquires into morally justifiable criteria for coercion. This paper concerns itself with this second, philosophical question. It is critical, in my view, that a prescriptive theory be able to contribute to a descriptive ethics, but there are distinct goals in the two enterprises. The former concerns the values people *should* have, that is, those that satisfy certain conditions. The dominant traditions in moral philosophy involve assumptions about some minimum level of information and impose impartiality conditions such as the universifiability of rules (as with Kant's categorical imperative), consistency with the sentiments of an impartial spectator (as in Adam Smith's moral theory), or compliance with some criterion that does not discriminate among persons (as with utilitarianism). An implication of these conditions is that the moral act or rule does not depend on the identity of individuals – a requirement similar to what social-choice theorists call *anonymity* (i.e., invariance under permutation of agents). Descriptive analysis of morals, on the other hand, concerns words or

actions about morals, which often reflect not only such impartially justified values, but also the influence of personal interests and limited (or false) information. I will elaborate this point in the next section.

Building on the major philosophical traditions on this topic, I make no grand claims of originality; indeed, most of the principal elements are familiar. But I hope the proposed framework will provide new insights for answering puzzles about and possibly guiding policy on matters of coercion and consent. Section 2 presents a framework for consent as finely graded inferences about individual agreement, which depend on multiple factors. Section 3 addresses challenges to the claim that consent occurs in fine degrees. Section 4 argues that justifiable coercion requires consideration of both consent and consequences, and section 5 concludes.

2 Consent: Complex and Continuous

In the second of his *Two Treatises of Government* (1689/2003), Locke proposed a theory of consent built on a distinction that can be traced to Socrates (see Plato's *Crito*, 360 B.C.). Lockean consent is a binary variable, viz., *express*, which “[N]obody doubts,” or *tacit*, which is implied by other volitional choices. Locke was concerned with the basis of allegiance to a government, and he tied tacit consent to ownership or use of land, or even residing temporarily within the boundaries of a state (see Russell, 1986, for a detailed discussion). The discreteness of Lockean consent can be contrasted with utilitarianism. Although utilitarians are often uncompromising in their devotion to the maximization of utility, utility itself is usually conceptualized as admitting of fine degrees of difference.

Some previous authors have voiced misgivings about Lockean consent and have argued for more differentiated treatments. For example, Bell (2010) adds to Locke's express and implied consent a third category of *hypothetical consent*, a still weaker form that “differs from its express and implied counterparts in that it ignores facts about what any given party does or does not want, instead relying on a counterfactual supposition about what the party *would have wanted*” (p. 18). An example concerns emergency medical aid administered to an unconscious man, and the argument is that anyone would have wanted treatment under such conditions. Bell creates further subtypes of these three, eventually identifying 14 ranked categories of giving or withholding consent, which he calls *graduated consent*, as well as defining *nonconsent*, i.e., actions in which consent plays no role. Although Alexander (2008) describes consent in dichotomous terms, viz., valid or not, he also points to continuity in the multiple considerations (which he calls “scalar properties”) that affect binary moral judgments in general, as well as in the specific case of consent. For instance, information affects our understanding of consent and can vary in degree: surely, a patient has not consented to a hysterectomy if she has no idea what that surgery is, but it is also impractical to enumerate all possible consequences of surgery. How much information is needed for her consent to be considered valid?

By recognizing finer degrees of consent, Bell's graduated consent represents an improvement over Lockean consent, in my view, but it remains a categorical approach. It excludes still finer degrees of consent and does not identify the general properties that govern such degrees. Alexander addresses continuity in underlying properties, but his focus is on identifying the philosophical challenges in going from continuous properties to binary moral judgments and not on arguing for continuous moral judgments themselves, including those about consent.

By contrast, I argue that consent varies in small degrees, which can effectively be treated as continuous. Moreover, consent is also complex: it is constructed by inference out of different and connected factors that strengthen or weaken the degree of consent. Specifically, I describe consent in relation to four factors. These considerations about the giving of consent carry over to the negative range of denying consent, or *dissent*. For conciseness, however, I will focus on situations involving consent rather than dissent. I claim that such an approach is more accurately descriptively accurate of impartial moral views, making it better for explaining numerous puzzles, and is also more useful for the prescriptive analysis of consent, thereby holding promise for resolving real dilemmas.

Let me be clearer about what I mean when an agent consents: he voluntarily approves or gives permission for some choice in which he has an entitlement or right. Thus, I side with Locke by requiring consent to be voluntary and exclude, therefore, "coerced consent" as being true consent. Coerced consent is not consent, it is just coercion: a rape victim did not consent, regardless of whether she fought the attacker, remained silent, or even expressed verbal agreement, if only to avoid some worse fate. Instead, the focus here is on voluntary consent, specifically consent to be coerced.

The agent's consent, therefore, is a type of agreement, but as with all agreements, it is incomplete: the realized state of the world in which coercion is relevant will always contain features that, had they been anticipated, would, at least under the right conditions (e.g., low negotiation costs), have been addressed explicitly in the consent phase. Given, however, the incompleteness of any agreement, this permission itself is never entirely complete. To illustrate this point, imagine the most complete agreement one might actually encounter in the real world, perhaps a detailed contract involving the transaction of goods between two rational, well-informed, and equally powerful parties who, after lengthy consideration, have both signed the contract before witnesses; moreover, there are detailed statutes and voluminous case law to bear on the agreement. The agreement severely penalizes the seller for failure to deliver by a certain date and time, but an extremely unlikely, unforeseen, and unavoidable event – in other words, something that was not provided for in the contract (perhaps an accident that befalls the seller's spouse) – delays that delivery by minutes. The courts might well uphold the penalty, which the seller is forced to pay. But my question is not a legal one but rather a moral one, and I venture to say that our moral intuition tells us that this circumstance does not signify the same degree of seller assent to the penalty as the case in which a provision for the unfortunate event had been explicitly included in the contract.

Even if the unanticipated event is considerably less dramatic and did not trigger an explicit penalty (e.g., transportation costs rose, decreasing the seller's profit), the fundamental point remains that a contingency was not provided for that, under certain circumstances, would have been. Even the most basic theory of consent that only distinguishes express and implied consent supports this claim. The point is that even the most impeccable agreements are incomplete, so that all consent is, in some ways, implied.

In the case of interest here, consent precedes coercion. Since the future is uncertain, one can think of consent in two senses: *ex ante*, i.e., with the information available at the time consent is given (or implied), and *ex post*, i.e., in the realized state of the world in which it becomes relevant for coercion. The above example is *ex post*, since it involves a single realized state of the world. Even a single state is potentially associated with many permissions (or denials of permission). At the time of consent, however, there are an infinite number of possible states of the world, so the incompleteness of consent increases exponentially with *ex ante* consent. For example, a patient's permission to a physician (e.g., signed medical release, general arbitration agreement, implicit understandings about standard medical practices) never fully addresses even a single realized state, and there are many possible states (cure, allergic reaction, uneventful treatment, etc.). It is impossible, therefore, to obtain express consent for every contingency.

Let me push this point further: there is an aspect of hypothetical consent with any action – a question of what agents would have wanted. In evaluating consent, whether in the personal, legal, political, or constitutional realm, we are typically not only quantifying observables but also constructing counterfactuals. Indeed, oftentimes the observables, such as past agreements or social customs, merely serve as proxies for estimating the degree of agent consent, when consent is not express. We can press the argument to the extreme, to situations in which consent ostensibly plays no role. For example, Bell gives, as an example of nonconsent, the case of someone rambling at will through a public forest. But I would argue, even in such cases where the agent intends no act of consent, consent is implicated: let us adapt another of Bell's examples, and suppose our wanderer were attacked by a wild animal and lay unconscious. Then, he might well be held, both by courts and shared intuition, to have consented to medical aid by virtue of his walking in a community where such is the custom, even if that were the furthest thing from his mind, when he set off on his hike. In fact, it is difficult to think of a situation in which consensual issues might not arise: a variation of this example involving medical aid to someone unable to express consent serves nicely to illustrate how far consent extends (perhaps it only fails to reach locations, like deserted islands, where coercion also does not extend?). Certainly, consent is usually a trivial consideration; indeed, most human affairs are carried out without resorting to standards of consent. Nevertheless, the point is that consent is potentially implicated in almost every situation.

The foregoing passage highlights a subtle, but important, difference in common usage of the word consent, which I will now make explicit. On the one hand, consent

refers to an action of the agent, but, on the other hand, it connotes our interpretation of such an action. Indeed, interpretation is introduced the moment we consider implied consent. As explained previously, this paper concerns concepts that are consonant with our impartial moral intuitions, e.g., the views of uninterested observers. So, here consent represents *our impartial inferences about agent consent*.

Thus far, I have made claims about what we cannot usually do, namely, treat consent either as perfect (or complete) or as entirely absent. Nevertheless, I argue that there is wide variation in the degree of consent. Consent in the inferential sense signifies the judgments we make about the extent to which agents have waived their rights – specifically, in this paper, about the permission they grant others to coerce them. It is constructed out of and varies nearly continuously with multiple and connected factors. I will outline four factors that I believe are important for evaluating the degree of consent.³

First, the most important distinction in most discussions of consent is express versus tacit consent, which, in turn, pivots on the *discretion* exercised by the agent. Express consent is so compelling precisely because of the strength of the discretionary act: the express consenter explicitly affirms approval. But even express consent comes in many forms and corresponding degrees. In the legal context, Bell argues that express consent is stronger with a negotiated exchange than a standardized one (e.g., a negotiated contract versus a standard-form contract). In the political context, one might add, in descending order of consent, direct democracy, representative democracy, autocracy, and the many variations on each. But discretion is also central to implied consent: tacit consent carries moral weight precisely because it is tied to other, related discretionary acts from which consent may be inferred. The strength of this inference, and therefore of the implied consent, also varies, from stronger signals of agent discretion (e.g., a conspicuous sign warning pregnant and infirm people not to go on an amusement-park ride) to community expectations (e.g., about the maintenance of one's house). Thus, consent varies in degree with how closely the agent's discretionary act or acts can be associated with permission to coerce, and this differs widely, e.g., according to choices to sign a contract, vote in an election, make a promise on a handshake, or defer to shared expectations and forgo explicit agreement.

The degree of inferred consent also depends on the conditions under which the permission is given. Commonly accepted grounds for not honoring consent are when (1) the consent is coerced, (2) deception is involved, or (3) the decision-maker is not deemed competent to give consent, e.g., as with children or the insane.⁴ I believe these provide clues as to three more general conditions that affect the degree of consent, which I will frame in positive terms: choice set, information set, and agent competence. Thus, consent relates to four factors: discretion plus these three conditions. It is constructed by inference from the discretionary act(s) of

³ I do not claim, however, that this list is exhaustive, since that would be premature at this stage.

⁴ For a related account of what vitiates consent, see Alexander (1996).

the agent subject to the set of possible choices, information, and competence. We never observe *ideal consent*, i.e., an explicit and complete set of permissions for all states of the world selected under perfect conditions: an unrestricted choice set, full information, and the highest levels of competence. Thus, we estimate the consent implied by the agent's discretionary act(s) given the actual conditions regarding choices, information, and competence. I outline these factors below.

As previously stated, I consider coerced consent as coercion rather than consent, since it is not voluntary. But it does share with all versions of coercion the fact that it affects the *choice set*; specifically, it alters the payoffs associated with different choices, which is important for understanding consent. If the coercer desires the agent to choose action A , the former threatens to impose a penalty p on the latter for failing to choose A . This is what we mean by coercion. Let us employ these terms to clarify further what is meant by consent. An agent consents to A if he agrees to A even in the absence of the threat of p . Similarly, if an agent agrees to $\{A, p\}$, that is, to the imposition of p should he not choose A , then that also constitutes consent to A . But if the agent agrees to A only because of the threat of p , then that agreement is not voluntary and is, therefore, not consent.

I have defined consent as our inferences about agent consent, and those depend in a rather intricate way on the discretion exercised by the agent given the choice set (with its associated payoffs). I will illustrate one way in which such inferences depend on this interdependence with a numerical example.

Many countries have some form of compulsory national service, usually military service but sometimes also including civilian service (e.g., Brazil, Finland, Greece, South Korea, Switzerland, and Turkey). The penalty for noncompliance can be incarceration, although monetary fines are sometimes imposed, either officially or de facto. As an example of the latter, Mexican citizens often avoid service by paying a bribe. Consider this stylized example: suppose one can complete national service denoted A and earn $a = 40$ or pay a penalty p (say, a fine or bribe) to avoid service and work at a different job denoted B and earn b . In addition, suppose that these are the only available choices and that people in this situation have different opportunity costs at the alternative job. Specifically, for simplicity, assume a continuum of agents whose earnings, b , are uniformly distributed over $[30, 70]$. Let these numbers represent all relevant values associated with each variable, converted into some common metric such as a currency or utils. If there is no penalty, i.e., $p = 0$, then no coercion occurs. Nevertheless, note that some people will voluntarily choose A , namely, those whose $b \in [30, 40]$. All of these people consent to national service, because their opportunity costs are below the compensation from completing national service.

Now suppose the penalty is 10. Those whose $b \in [30, 40]$ still choose A , but, in addition, the next quartile also choose A , since their payoffs from B net of penalty are $[40 - 10, 50 - 10]$, making A more attractive for them. The remaining agents, whose $b \in [50, 70]$, choose B . Thus, those whose gross earnings $b \in [30, 50]$ choose A , and, of these, one half consent ($b \in [30, 40]$), since they voluntarily choose A , and the other half are coerced ($b \in [40, 50]$), since they are compelled by p to choose A .

Note what happens as the penalty rises. It is straightforward to show that, if $p = 20$, two-thirds of those choosing A are coerced and one-third consent, and if $p = 30$, three-quarters are coerced and one-quarter consent. As the penalty rises, effective coercion increases, and the fraction of agents voluntarily choosing A decreases. Remember, we are interested in the inferences drawn about consent. If we do not know the b -values of individual agents, we cannot determine whether a particular situation involves coercion or consent. Then, the degree of consent reflects the estimated probability that an agent agreed voluntarily. We might have even less information and, at best, a rough idea of the distribution of types. But for inferred consent to vary in this manner, we need only be familiar with this very simple inverse relationship between coercion and consent.

This dependence of inferences about agent consent on the discretion exercised within the choice set is a more general phenomenon. In the example above, there are consent and coercion but no consent to coerce, which is the main concern of this paper. Nevertheless, one can extend the reasoning to such cases, although the analysis becomes slightly more complicated. As previously discussed, people might consent to be coerced when coercion provides some benefit, e.g., funding a public good or increasing personal savings. That is, coercion, in these cases, alters the payoffs not only by adding penalties for some choices, but also by increasing payoffs for the coerced choice. Consider a public good. Without coercion, few, if any, would expect to secure a net benefit from contributing, but with coercion there are net aggregate gains. Typically, though, contributions and benefits are not equally distributed, and some individuals will benefit by less than their contributions. Thus, we can expect a distribution of types consisting of some who consent to coercion and some who do not. Varying the penalties will shift the fraction who are coerced, the fraction who consent, and, therefore, inferences about the degree of consent similar to the previous example.

The choice set in the examples above contained only two elements, but choice sets are actually much richer. Thus, the degree of consent varies inversely not only with coercion but also with constraints on relevant alternative choices. For example, the consent we (not necessarily courts) infer from a patient's signature on an arbitration agreement in a physician's office is greater when it is not required for treatment or alternative caregivers who do not demand such an agreement are locally abundant than when it is required for treatment by the only physician in town. Other constraints result from bundling of choices: bundling is high with standard-form contracts or a two-party political system, and consent is therefore usually lower in those cases than with negotiated contracts or plebiscites. These examples underscore the incompleteness of consent, even estimated at the individual level: a choice actually represents a vector of choices, and the agent might approve of some elements but not others, making the consent to the bundle only partial. Of course, not all constraints affect our inferences about consent, e.g., people often prefer different types of health insurance coverage. Indeed, for some people, certain types of coverage (e.g., for contraception) might be negatively valued. Interestingly, though, there is evidence from economics experiments that people sometimes value

the availability even of options they do not wish to exercise (Falk and Kosfeld, 2006), so we cannot rule out the potential relevance of such choices. Moreover, the effect on consent can be ambiguous, when both discretion and the choice set vary: greater explicit specificity increases express consent, but the possible bundling of conditions not desired by one party or the other can decrease implied consent and create more points of disagreement, as with, perhaps, the 1000-page Affordable Care Act.

Consent obtained by deception is not credible, because, by definition, the consentor has been misinformed. The deliberate nature of deception probably adds to moral objections, but I argue that the more general point for questions of consent is about the agent's *information set*, which can vary in fine degrees. High information (that is, more correct and relevant information) is preferable, *ceteris paribus*. In democracies, the electorate is encouraged to familiarize itself with issues and candidates, whereas under autocratic governments information is often suppressed by the state. Although, in practice, both types of efforts often fall short of their respective and opposing goals, these facts provide indications of the intuition relating moral goodness of information to consent. Similarly, *informed* consent is the cornerstone of many basic protections against abuse, including those in connection with medical procedures and human experimentation. As with constraints on the choice set, inferences about the consent from discretionary acts given the informational constraints on them can, at times, involve some nuances, and these go beyond the scope of this paper (for further discussion, see Alexander, 1996). But the central point is that the strength of the consent varies in direct relation with the extent to which the agent knows to what he or she is consenting.

Children and the insane are typically viewed as incapable of giving consent to many commitments and constraints, on grounds related to competence. Consent varies quite generally, I believe, with the perceived *competence of the agent* to give it. At issue are the agent's capacity to perform the discretionary act of approval, and his abilities (affective and/or cognitive) to process the relevant information. These qualities can fail in various ways, which affect the degree of consent. Regarding the performance of the discretionary act of giving permission, consent is, of course, weakest when agents literally cannot express themselves – due, say, to a physical condition or to being temporarily unconscious or comatose. *Affective* states also relate to the degree of consent; for example, emotional maturity is generally seen as increasing with age and is surely an important reason children may not consent to many decisions. *Cognitive* limitations weaken consent in many ways; e.g., low intelligence, intoxication, and mental illness can invalidate or preclude claims of consent. But there is a range of cognitive limitations, some common to large numbers of people, including cognitive biases such as anchoring, confirmation bias, endowment effects, hyperbolic discounting, and recency effects. This has led to the concept of *libertarian paternalism*, i.e., the notion that one should intervene, without consent, to “nudge” people in certain directions (e.g., Thaler and Sunstein, 2008). Libertarian paternalism is criticized by strong advocates of consent, but this and the other examples are merely intended to illustrate how com-

petence may affect common views of the degree of consent and justification for intervention.

As previously stated, we never observe ideal consent; but, if we keep this fact in mind, that ideal might serve like the North Star: it is not attainable, but it can provide direction – in this case, as a useful means of estimating degrees of consent as distances from the ideal. These degrees are constructed out of the four factors just summarized: the discretionary act of the agent related to the coercion, the choice set itself, the information conditions, and agent competence.

But how, exactly, do we construct consent out of its factors, and how do we resolve conflicts between the factors? As defined here, consent reflects moral preferences, so – as with preferences for anything, including material allocations – we usually formulate these in general terms and leave the precise form to empirical investigation. Nevertheless, the method of reasoning employed otherwise in this paper might provide some guidance on this question. As treated here, consent is constructed out of multiple factors, which exhibit trade-offs, and is defined on an open interval bounded by zero consent and ideal consent. A potential problem is that the ordering might not be smooth; e.g., perhaps only a person with a minimal level of competence can properly process information or exercise discretion, and below this level no increase in any other factor can increase consent, so that trade-offs only occur above certain threshold levels for factors. This seems to imply that consent can be zero, an idea against which I argued previously: consent can be inferred even in the case of hypothetical consent. The tension between these two claims can be resolved, I believe, by noting that consent requires that the agent have had minimal levels of information, competence, and choices in some way related to the coercion at hand at some point in time, while recognizing that not every variation in every factor at every point in time is relevant to every coercion (e.g., giving additional information or choices to a comatose person does not change their inferred consent).

I believe that the multidimensional characterization of consent based on fine degrees of difference in each of the factors of consent allows ranking of cases in a manner more consistent with our moral intuition than a unidimensional system, such as Lockean consent or Bell's graduated consent.

Locke's statements suggest that express consent is stronger than implied consent, which accords with our intuition. The conception of consent proposed here also ascribes a higher degree of consent to the former, *ceteris paribus*, but it follows from the factors of consent: Express consent involves a discretionary act that more directly aligns with the granting of permission. Nevertheless, implied consent sometimes outranks express consent, and the current proposal can also handle such cases. For instance, even written consent to be frisked would not be valid if secured under false pretenses (an impostor posing as a police officer), whereas boarding a flight strongly implies such consent. Other examples of strong implied consent include leaving one's car keys in the ignition, which can make the owner liable for any theft. Similarly, obtaining a driver's license is considered implied consent to submit to a sobriety test, a position that has been widely upheld in the

courts. On the other hand, some forms of express consent have been ruled invalid. A cannibal, whose victim willingly responded to an Internet solicitation and who videotaped his consent to be killed and eaten, was, nevertheless, convicted of murder (http://en.wikipedia.org/wiki/Armin_Meiwes). “Clickwrap agreements” require users to click “I agree” to download software on the Internet. Although courts have, more often than not, enforced such agreements (and the related “shrinkwrap licenses”), it is notable that they have failed to uphold some versions (e.g., *Brower v. Gateway 2000, Inc.*, and *Williams v. America Online*), calling one such agreement “unconscionable.”

When implied consent is stronger than express consent, it is due to the fact that the greater discretion associated with the latter is more than offset by the countervailing influence of other factors of consent, such as information or competence. The typical passenger is well informed of the requirements for air travel, whereas the police impostor employed deception. Drivers are held competent to lock their cars and avoid intoxication, whereas the victim of cannibalism had mental problems, and the software user might find the language of a license too technical or time-consuming. Thus, express and tacit consent both vary continuously, and the latter sometimes trumps the former.

Nevertheless, when factors other than discretion are held constant, express consent is stronger than tacit consent. Indeed – unlike other differences in discretion, choices, information, and competence – an explicit statement of consent typically causes a discrete, rather than continuous, shift in our inferences about the agent’s consent. This claim not only aligns with our intuition, but is also consistent with the fact that express consent is not usually overturned with ease, i.e., by a small opposing variation in other factors. For instance, suppose I would hide my car keys if I knew my roommate had recently received a traffic ticket, to avoid any impression of implied consent to borrow my car. But if I did not know that and granted him express permission to borrow my car, that express consent would not be undone by the revelation of his traffic ticket. That said, however, there are many cases in which the marginal influence of expressing consent is so minor as to have a negligible effect on consent. For example, in most countries, the understanding that one is obliged to secure a license in order to be authorized to drive is so strong that it is difficult to imagine that signing an affidavit consenting to that requirement would alter the perception of consent in any practical way.

Turning now to Bell, his graduated consent involves various ranked categories, and in that system, a *negotiated* contract represents a higher level of consent than a *standardized* one. But suppose that, in arriving at a negotiated contract, a parts supplier deceives the buyer about conditions, and this will permit later increases in the price of the parts. Further, consider a standardized contract between a landlord and tenant that involves no deception and corresponds to the wishes of both parties. Bell’s consent does not provide for distinctions based on such factors and therefore ranks the negotiated parts contract above the standardized rental agreement in terms of consent. According to the conception of consent proposed here, however, this ranking is transposed, since the rental agreement incorporates more favorable

information conditions and does not impose a binding constraint on the choice set any more than a negotiated agreement does.

3 *Continuity of Consent: Challenges and Responses*

The proposal above rests on two claims: consent is inferred from multiple factors, and both consent and its factors vary continuously (except, at times, express consent). The second claim has been the subject of some debate in the literature, so this section presents and responds to some challenges to this claim.

The continuity assumption seemingly conflicts with many real-world situations involving consent in which policies (e.g., collective choice procedures, legislation, or court rulings) are binary or otherwise discrete. For instance, voters often face the choice of only two parties, the age of adulthood with its many consensual rights and responsibilities occurs in an instant, and judges and juries must often rule either guilty or not guilty. One might counter that, on closer inspection, many policies are more continuous than they at first appear. Many political systems accommodate parties representing a wide ideological range, and even two-party systems allow wider expression through primary elections, propositions, and the ability to vote on many individual candidates, rather than only a party. In some jurisdictions, the age of consent is introduced gradually: sexual consent is possible at 16, voting at 18, and the drinking of alcohol at 21 (and some European countries set limits on alcohol content that increase with age). Legal guilt is determined relative to punishments that are already scaled to be proportionate to the offense (e.g., manslaughter versus second-degree murder versus first-degree murder) and is further smoothed through sentencing.

But even if such scrutiny reveals finer grades, there clearly remains much discreteness in policies. Only a rare voter, if any, can point to a politician or party that reflects his or her views on every issue. Even when the age of consent is implemented in stages, the stages are few and still take effect abruptly rather than gradually: one's vote does not count as a fraction that increases with age. Although there exist categories of offenses, they remain categories, and legal judgments remain largely discrete: an accused rapist is not deemed 90% guilty if the evidence suggests a 90% probability he committed the crime.

The mere fact that policies are often discrete, however, does not necessarily imply that the underlying *moral* space in which we construct consent is similarly discrete. Policies can be lumpy even if moral concepts are continuous. Since consent is defined, in this paper, in terms of morality and inferences, the most relevant evidence on it is our moral intuition. Nevertheless, I have, at times, marshaled evidence from policies to support my arguments, so it is only fair to confront this proposal with policy evidence in apparent contradiction to it. The goal is not an exhaustive defense – a presumably impossible task and one that would go beyond the scope of the paper. Rather, I attempt to demonstrate several plausible explanations for discreteness in the policy space that do not require discreteness in the moral space of consent, while

addressing some important challenges to my position. Issues of consent and policy continuity have been debated mostly in the legal domain, so my discussion focuses on that context.

Imagine we could perfectly determine the degree of consent to all coercions at no cost. For instance, suppose we could identify the exact discretion, information, competence, and choice set of people and construct their consent for each item among the rights and responsibilities associated with adulthood. This is a very bold assumption, but even if it held, it would certainly not be costless to calibrate each right or responsibility. Presumably, one would have to create and maintain a data bank for each person and each item. Moreover, the factors of consent vary over time, which would require monitoring and adjusting each item with respect to each person's (hypothetically known) consent, such as the extent to which they may be compelled to serve in the military or to pay taxes. Often factors do not even vary monotonically; e.g., although competence might initially increase with age, it can subsequently decline with senility, and there are also short-run variations due to fatigue or intoxication that would necessitate adjustments in policies. These costs are, in most cases, blatantly prohibitive. Since consent varies by person, item, and time, any broadly applied discrete policy will always be, in some sense, arbitrary, but the high costs of a finely calibrated system of policy creation and enforcement make such a recourse impractical.

Apart from these concerns, though, there are limits, of course, to the precision with which consent can be empirically specified. Even perfect conceptual clarity about consent does not imply the existence of an exact measure that can be straightforwardly applied to policy. Some factors, such as the information and competence of the agent, are largely latent variables that might, at best, be approximated through proxies, and even a mostly explicit variable, like the choice set, cannot be precisely catalogued. Although people have strong and converging intuitions about many matters of consent, in the end, these involve subjective values and variables subject to measurement error. Differences across individuals and contexts present challenges to resolving conflicts and to taking collective action. But this does not imply that we must abandon any attempt to base policies on consent: its primacy in many contexts is compelling, and often we are sufficiently confident about our estimates of it to craft simple policies. Such uncertainty, though, makes it difficult to fine-tune policies. It also creates a space for "moral wiggle room" where agents can make misrepresentations, which I now discuss.

On the one hand, agents might misrepresent facts to others. For example, allowing partial rather than full consent to alcohol testing might have unfortunate incentive effects: when the factors of consent are not (entirely) verifiable, agents could strategically claim ignorance of drinking laws or constrain their choice set by failing to take along cab fare on a night of drinking. Thus, the best approach might be a simple, discrete policy: grant certain choices to adults while holding them fully responsible, and withhold those options from entire groups that are considered, on average, not to be sufficiently competent. On the other hand, agents might misrepresent matters to themselves, i.e., engage in self-deception, which can obstruct agreements involv-

ing consent. Relevant results on this point can be found in a series of studies by Babcock, Loewenstein, and their coauthors (e.g., Loewenstein et al., 1993; Babcock and Loewenstein, 1997). They report evidence that self-serving moral beliefs are an important reason for bargaining failures, which result in labor strikes and in court cases proceeding to costly trials. Although their studies focused on fairness biases, the straightforward implication for consent is that adding finer distinctions to a contentious legal context is likely to be counterproductive and to increase the incidence of disagreements.

There are additional reasons discrete policies can be superior to continuous ones, even if consent is continuous. Kaye (1982) addresses legal cases in which one knows a defendant is liable only with a certain probability and argues that it is preferable, under most conditions, to hold the most likely defendant accountable for full damages rather than to allocate damages in proportion to the probability of liability. Let me illustrate this argument using a criminal case mentioned briefly earlier: assume it is known that the probability that an accused rapist is guilty is 90%. Suppose the uncertainty concerns whether or not the act was consensual, and the sentence for this crime is 20 years. For every innocent defendant found guilty, then, the judicial error is 20 years. Put differently, out of ten defendants under these circumstances, one is misjudged, and the judicial error is 20 years. Now suppose ten defendants are sentenced according to the probability of guilt. For the innocent defendant, this reduces the error to 18 years. But the sentences are also reduced by two years for the nine guilty defendants, increasing the judicial error for them by 18. Thus, the proportional system increases judicial error from 20 to 36 years. Even if inferred consent lies at an intermediate point on the interval, the policy the produces the least error is discrete.

The mapping of consent onto policies, then, is sometimes a step function, producing discrete changes in policies, as when judgments include serious charges, lesser counts, or innocence. Indeed, there is a single step with all-or-nothing policies, such as guilty or not-guilty verdicts. Such discreteness in policies related to consent can create the impression that consent itself is discrete. For example, informed patient consent is the only thing, from a legal standpoint, separating surgery from criminal battery. A court ruling on whether consent was given is unlikely to be affected by marginal changes, such as the addition of a few sentences to express consent, or differences in implied consent based on whether the patient had voluntarily undergone other surgeries previously. Even when a policy is discrete, though, it does often respond to differences in the degree of consent, when they occur close to the discontinuity in policy. There are qualifying conditions that courts have ruled to invalidate consent to surgery, even after the fact – in particular, those affecting competence, such as mental illness, the influence of alcohol, drugs, or medications, or being in labor – and in such cases, the degree of impairment of competence matters. And the absence of express consent does not preclude surgical procedures: emergency medical assistance without prior consent has often been upheld by the courts. Thus, the factors of consent can influence even discrete policies, and understanding these underlying forces can help us identify where these policy switches should occur.

Leo Katz (2011) makes a clever and inventive case for the all-or-nothing (which he calls “either/or”) character of legal doctrines. Here I wish to address his more specific argument that not only legal policy, but also consent itself, is all-or-nothing and not a matter of degree. He illustrates his claims with a hypothetical scenario involving a defendant who demands \$1000 from a victim, who then complies, but the threat employed varies across five versions of the scenario: (1) killing the victim, (2) destroying the victim’s treasured rose bushes, (3) driving through the victim’s neighborhood with negligence, putting the victim and his/her family at risk, (4) publicizing the victim’s past relationship with the defendant, which the victim would find highly embarrassing, and (5) ending their friendship. According to Katz’s characterization of the degrees-of-consent argument, coercion decreases and consent increases progressively as we proceed from Case 1 to Case 5. Nevertheless, he argues that this story is actually incorrect and that the validity of consent does not, in fact, vary in this fashion. Instead, consent is totally and equally invalid in Cases 1 to 4 and then totally valid in Case 5. There are two parts to his argument. One is the positive argument in favor of all-or-nothing. The other, and more extensively worked-out, part involves first developing and then critiquing a degrees-of-consent argument while employing this critique to bolster the case for all-or-nothing consent. Since my own argument for continuous consent differs from his version for degrees of consent, and since his critique of that version is very intricate, I will focus on his affirmative case for all-or-nothing.

Quite simply, Katz’s argument for all-or-nothing consent is based on blameworthiness. The validity of consent depends on the blameworthiness of the defendant, and “Blameworthiness stays the same from Case 1 to Case 4 and then suddenly plummets to zero when we reach Case 5,” which “is just another way of saying that the defendant behaved equally coercively in the first four and not at all coercively in the last” (Katz, 2011, p. 168). Underlying this reasoning is an insightful categorical distinction about entitlements. Katz formulates this with respect to the defendant: he is entitled to end the friendship (Case 5), but, in the other cases, he “is not entitled to do these things either as a matter of criminal law (Cases 1 and 2) or as a matter of civil law (Case 3) or simply as a matter of morality (Case 4)” (p. 166).

The definition of coercion introduced at the start of this paper also pivots on entitlements. It was formulated, though, with respect to the victim, the person being coerced rather than the coercer: the penalty with which the victim is being threatened involves something to which the victim is entitled. Thus, the victim is entitled to his life, his rose bushes, the safety of himself and his family in their neighborhood, and, presumably, his privacy. It seems, however, that he is not entitled to a friendship. Thus, I agree wholeheartedly with Katz that entitlements provide a critical distinction for identifying the presence of coercion. And, in his example, it does not matter whether coercion is defined in terms of what the victim is entitled to or of what the defendant is not entitled to do to the victim.

Nevertheless, consider why, even if we accept that coercion is total and equal in Cases 1 to 4 and that there is a discontinuity with Case 5, this does not imply that consent is all-or-nothing. Remember that, to be consent, it must be voluntary;

but this example suggests little basis for freely choosing the action desired by the robber: in the absence of the threat, why would the victim hand over \$1000? This is unlike the earlier example of national service, in which some people can be expected to participate voluntarily. If coercion is total in Cases 1 to 4, then consent is not invalid – it is simply not present at all. Consider now Case 5, in which the threat is to end a friendship, and let us accept that this case does not involve an entitlement, in contrast to Cases 1 to 4. But if it does not involve an entitlement, then it also does not involve consent. Consent refers to agreements or permissions relating to one's entitlements or rights – both my and Katz's usage – but then consent is not implicated in this context. On this reading, therefore, it is correct that consent does not vary continuously across these versions of the scenario, but only because it is not a consideration in any of them.

Moreover, I believe there are good reasons to doubt the claims of equal coercion across Cases 1 to 4 and of zero coercion in Case 5 and their implications for inferences about consent. Although this scenario is not well suited to illustrating the consent proposed here, due to the lack of explicit variation in the factors constituting that consent, people might well impute their own assumptions about variation in these factors as well as in entitlements and penalties. As a result, I expect a survey with a response format that permits measurement in degrees (e.g., a Likert scale running from “strongly agree” to “strongly disagree”) would reveal some differences across the five versions. For example, with respect to the interpretation of entitlements, few would dispute that people are entitled to their lives (Case 1), but more would likely doubt an entitlement to privacy about relationships (Case 4), especially in this day of social media and reality shows. Regarding penalties, surely most people would, in fact, consider a threat to the victim's life (Case 1) as more coercive than the threat to destroy the victim's rose-bushes (Case 2). And, although friendship does not imply strong entitlements (Case 5), personal relationships commonly evoke expectations of certain rights and responsibilities, and some respondents would likely see threats to withdraw friendship as imposing penalties and, therefore, involving some level of coercion. Respondents might even interject assumptions about the factors of consent. For example, the victim's discretionary act of paying the \$1000, when the threat is low (e.g., destroying rose bushes or ending a friendship), might be seen less as a reaction to the threat than as a voluntary response to other considerations (e.g., possible pity for the defendant), which introduces some role for consent. Any such variation, though, suggests degrees that can be reconciled with the conception of consent proposed here.

Katz is careful not to base his claim of discreteness in consent solely on discreteness in the law: the coercion in Case 4 is not illegal, but he considers the consent invalid nonetheless. Many examples in favor of discreteness in consent, however, are developed in a purely legal context and run the risk, therefore, of erroneously attributing properties of legal doctrines to consent. We have already discussed one basis for doubting that logic: there are reasons for discrete policies, whether legal or otherwise, that have nothing to do with consent. This section concludes with some

caveats specifically about arguing from legal context, which relate to fundamental differences between the legal and nonlegal domains.

The law and morality are routinely viewed as closely connected, e.g., in natural-law theory. Thus, rape is both immoral and illegal. But the boyfriend who threatens to end the relationship if his girlfriend does not have sex with him is not guilty of a crime, although many people would surely consider his conduct immoral and involving a measure of coercion. Indeed, much immoral behavior is not illegal, e.g., most lies, or failure to aid others whose plight is not our fault. Most of us would find an existence with no space unregulated by the law oppressive. So, not all immoral behavior is illegal, but perhaps the most immoral acts are illegal, and the more immoral the act, the more strongly it is sanctioned. For example, an innocuous lie violates no law, but slander (or oral defamation) allows limited sanctions, and libel (written or broadcast defamation) allows stronger sanctions. But the sanctions for illegal acts do not always seem to vary according to their immorality. For the law, stealing the last \$100 from a destitute person constitutes the same act of theft as stealing \$100 from a rich person who does not even notice the loss, although we perceive the former to be more immoral than the latter. But surely the law at least penalizes more-immoral acts as a whole, even if the sanctions do not always vary according to the immorality of the acts? Even this statement does not seem to hold: most people would likely consider it more immoral to cheat on one's spouse than to take a few office supplies from the workplace, although infidelity is no longer illegal in most jurisdictions, whereas taking property from the workplace is theft. Of course, there are connections between the law and morality, but the general point is that one cannot evaluate morality directly from legal doctrine. In particular, for the topic at hand, one needs to exercise caution about attributing to consent discrete changes in legal doctrine that occur because of switches between the legal and nonlegal domains. This would be a danger, for example, in Katz's scenarios, if we were to compare only Cases 1 to 3, which involve illegal behaviors, with Case 5, which does not.

In contrast to the legal context, the goal of the current paper is to shed light on impartial moral views of consent and coercion. Discreteness in policies does not imply the need for a discrete moral concept – to the contrary, I argue that it is precisely by keeping in mind the near continuity of the moral considerations that we might better recognize the discrete switches in policies that are sometimes necessitated in light of the influence of factors within the context as well as situations when the policy switches occur for reasons distinct from consent.

4 *Consent and Consequences*

In his essay “Of the Original Contract” (1752/2007), Hume argued for obedience to government based on social utility rather than on consent. He briefly acknowledged the appeal of consent (p. 503), but declared it insufficient and let the full weight of the justification for state coercion rest on its consequences. Like Locke, therefore, Hume also assumed a *monist* position, i.e., one based on a unitary moral justifica-

tion. Indeed, the subsequent debates on this issue have been dominated by monist arguments, mostly of the consensualist or the consequentialist variety. Of course, both schools of thought allow exceptions. For instance, even the most committed libertarian, who fiercely defends individual consent, allows state coercion to prevent the private violation of freedoms, e.g., individuals may not use force or deception against others. Nevertheless, in these cases, the monist position is still maintained, since these qualifications do not constitute a new and competing rule but rather remain idiosyncratic exceptions that preserve the rule.

Thus, there are conflicting intuitions and evidence on consent and consequences. If individual consent is required for coercion, then that inextricably implies it must be unanimous. Short of unanimity, one cannot claim to have obtained individual consent. But even allowing for commonly accepted exceptions to consent (viz., coerced consent, deception, and incompetence), it is clear that the state could not discharge most of the functions it currently does – indeed, it seems unlikely it could carry out any activities whatsoever, even the protection of person and property – if unanimity were required.

One might hold that unanimity sets the bar too high, but the alternative abandons individual consent and raises the sticky question of where the marker should be set. Simple majority rule is common in many institutions, but history is rife with examples of majorities suppressing minorities and squashing their ability to give or withhold consent, as in some recently democratized North African countries. Similarly, the exclusive regard for the interests of society as the criterion for the implementation of state coercion has been behind many dark episodes in history. If we set aside the intentions of state actors and take consequences as the sole justification for coercion, we are left with cases such as the actions of the Chinese government, which its critics accuse of employing capital punishment against political opponents and forcing the migration of hundreds of thousands (1.2 million due to the Three Gorges Dam alone), but which has achieved sustained economic growth near the double digits for extended periods. As another example, could the utilitarian view survive if torture successfully protected the United States and produced the greatest social utility?

But perhaps these examples do not actually contradict common morality. Of course, the appeal of governments to moral justifications for their policies might be a mere façade, but some evidence suggests popular support for their measures. It seems that (at least until recently) most Chinese have endorsed their government's policies: the Pew Research Center (2008) reported that more than eight in ten Chinese were satisfied with their country and its economy, the highest ranking of 24 countries surveyed. This fact seems to bolster the consequentialist case, not only by adding utilitarian-like evidence of subjective satisfaction to the objective evidence from economic variables, but also by suggesting a high level of popular consent. In addition, we have recently witnessed torture employed on utilitarian grounds even by the U.S., a country with a strong tradition of consent. Moreover, a national survey found that, at one point, 50% of Americans favored (and 46% opposed) the harsh interrogation techniques of their government, including waterboarding, on

suspected terrorists (Steinhauser, 2009). Is our moral sense genuinely troubled by these examples?

As stated previously, I wish here to distinguish the sentiments that can be reconciled with some impartiality condition from the claims that people frequently profess (and surely often believe) as interested parties and/or under unfavorable informational conditions. Thus, the support of many, even most, Chinese for the authoritarian policies of their government does not satisfy any impartiality condition, if it is associated with personal advantage. For example, indifference to coercion is not credible if it dissipates with the loss of the benefits associated with the rulers (as recently appears to be the case with slowing economic growth and rising environmental problems in China). As noted, the major traditions in moral philosophy require some minimum level of information as well as assume or imply anonymity, i.e., that the rightness of an act or rule should not depend on a person's place in society. And yet much of the rest of the world, which is not as dependent on the Chinese government as its citizens are and which has more balanced access to information, views something profoundly wrong about many uses of coercion against segments of the Chinese population. Similarly, popular support among many Americans for harsh interrogation techniques against suspected foreign terrorists does not constitute a normatively valid moral position, since it does not satisfy anonymity: such a rule is not invariant with respect to the nationality of torturer and victim. Moreover, both the rarity of support for torture as manifested in international agreements prohibiting its use and the circumstances under which torture occasionally gains in acceptance suggest that it reflects passions and fears of stakeholders rather than impartial moral intuitions.

In the previous two sections I argued that the conception of consent that distinguishes fine degrees of consent based on a set of four factors provides a better fit to our moral intuitions about consent and the use of coercion than unidimensional and discrete concepts such as Lockean consent and Bell's graduated consent. But can the addition of complexity and continuity resolve all challenges to the consensualist position? Consider the following two puzzles. Incarceration represents one of the strongest forms of coercion, and yet the U.S. Supreme Court has repeatedly ruled to place few limits on the use of deception by the police to secure confessions (including lying about evidence and eyewitness statements), the most extreme violation of the information condition associated with consent. In section 2, we also discussed restrictions on the choice space due to bundling of conditions of agreements. For instance, the functions of government represent an extreme example of bundling – the electorate can usually choose only among highly aggregated packages of policies. Voters agree to a panoply of public services, regulations, and income and wealth transfers, wherein any individual voter surely objects to many of the forms of coercion that are placed on him or her. But the proposal here does not explain why there is bundling to begin with. Why, instead, is consent not obtained from each person for each policy?

As an attempt to address such paradoxes, I propose a broader moral framework that includes both consent and consequences. Although most of the dominant theories

of ethics are monist, there are also *pluralistic* theories of ethics, which maintain that there are multiple moral principles. Some pluralistic approaches treat ethics as *context-specific*, e.g., moral relativists claim that moral values differ across contexts such as culture. Defenders of this position often cite differences in moral norms across cultures and ages, such as those regarding polygamy, infanticide, and honor killings. In the area of justice, I have argued, by contrast, that popular moral views are *context-dependent* (Konow, 2001, 2003). This position states that people operate from a common set of moral values, but the interpretation and application of those values is sensitive to the context, that is, to the salient set of people and variables being judged. Thus, moral judgments of behavior or norms can differ because of differences in the context, or in interpretations of even the same context, despite agreement about general moral principles. An example in the context of justice explains why Americans, on average, tolerate higher levels of economic inequality than Europeans: the former are more inclined than the latter to attribute differences in actual outcomes to effort than to luck, although both groups share the view that fairness requires effort but not luck to matter (see Alesina and Angeletos, 2005). I believe it is fruitful to apply such a context dependent approach to questions involving consent. Let me elaborate and illustrate this using the two puzzles above.

Regarding the admission of deceptive police interrogation, the U.S. Supreme Court has never stated its precise rationale, but there are indications it did not entirely disregard consent.⁵ Its early use of the so-called “voluntariness requirement” suggested a justification for deception on the basis of suspects having been informed of their rights to remain silent and to legal representation. That is, the idea is that subjects gave implied consent by being made aware of but failing to invoke these protective measures. From a less legally and more morally oriented angle, however, it seems difficult to defend the position that merely being informed of these rights conveys a suspect’s permission to be misled and lied to. Moreover, critics have pointed to instances of false confessions. Indeed, the Court has more recently moved away from this standard, and legal arguments for the use of deception increasingly emphasize the social benefits: “the substantial value of deception in obtaining confessions is based on long experience. Given the limited proof of the false confession problem, there is little question that the benefit of deception outweighs its costs” (Magid, 2001, p. 1206). Regardless of one’s position with respect to this practice, we recognize in these arguments a balancing of consensual and consequentialist claims. The shift in justification toward the latter reflects, I suspect, an increasing acknowledgment that the stronger case for compromising consent deemphasizes consent and argues from the consequences.

The consequence with which economists are traditionally most concerned is efficiency, and this also undergirds the classic economics argument for the state: it provides public goods that are otherwise unfunded or underfunded due to the free-rider problem. In economics, the gold standard for efficiency is Pareto efficiency,

⁵ Many of the details presented here draw heavily on Magid (2001), although my interpretations often differ from hers.

which is obtained when no further improvements are possible according to the Pareto principle. This principle calls for changes that produce benefits without making anyone worse off, that is, changes that produce no losers. In fact, Buchanan and Tullock (1962) base consent itself on this principle: they require *conceptual unanimity*. That is, any state constraints on individuals must be Pareto-improving, as only such constraints are presumed to secure unanimous consent (a weaker version only requires constraints to satisfy a supermajority rule). The case in favor of implied consent with this criterion seems strong: would not competent, well-informed agents consent to such policies?

In the real world, though, few, if any, policies produce no losers, rendering the Pareto principle of little practical relevance. Indeed, economists, in practice, actually tend to employ the *compensation principle*, which only requires that policies create more gains than losses, irrespective of the distribution of those gains and losses across individuals. Nevertheless, all this begs our original question of why we observe bundling of public policies, including public-good provision. Lacking the unanimous implied consent of the Pareto principle, there should not only be no bundling of public goods, but no public provision of public goods at all. So, returning to our original question, is there no way to secure consent individually?

In fact, mechanisms have been proposed to overcome the free-rider problem that avoid coercion by eliciting individual willingness to pay, e.g., Clarke (1971) and Groves (1973). Nevertheless, empirical tests of such mechanisms have been less promising (e.g., Attiyeh, Franciosi, and Isaac, 2000), and suspicion is often cast on the cognitive demands of these devices. These facts mean that the use of such mechanisms is associated with lower consent, due to problems of agent competence. The fact that they do not deliver the promised efficiency further weakens their case. Even if a mechanism could perfectly reveal preferences, however, there would still be the costs of eliciting valuations from every person for every public good provided, collecting a complicated schedule of fees, and, to the extent the public good is local or impure, identifying the beneficiaries toward whom it should be directed. This would obviously entail very considerable costs that would have to be balanced against the gains in consent; indeed, for many individuals, the increased cost might exceed the premium they pay in the absence of such a mechanism. Issues of cost are also relevant in consideration of the many other public policies, such as regulations, property-right laws, and transfer programs, that are in force without high levels of consent, let alone unanimous consent.

Thus, people balance the costs of consent with its benefits, which include consent's intrinsic value as well as its potential instrumental value. Regarding this last point, some claim high-consent societies are, on average, better-functioning and happier ones (Frey and Stutzer, 2000). According to the framework presented here, we should find this most true, not based on a univariate notion of consent, such as the mere presence of free elections, but when the underlying conditions for high consent apply, viz., a sufficiently unconstrained choice set, high information, and high competence. When there are significant deficiencies in one or more of those factors, consent itself is not high, and, if consent has instrumental value, the consequences

should also not be as favorable. Such deficiencies seem likely in some recently democratized societies, and critics of democratization in these societies also point to heavy economic and social costs. Thus, the interaction between consent and consequences can be negative or positive. The two are sometimes mutually reinforcing, as when prior voter consent enhances implementation of public policies, or when a successful parliamentary government is more inclined to call for elections.

The general point is that neither consent nor consequences alone are sufficient to justify morally the use of coercion, in general.⁶ This does not rule out, though, one or the other dominating in particular contexts. We saw with consent alone that context can make one factor more salient than another; e.g., when competence is low, improvement on that margin is critical to the validity of discretion. Similarly, in a multicriterion system, consent or consequences can dominate judgments; e.g., requiring express consent for organ donation presumably reflects the primacy of consent in this context, which can be obtained at a comparatively low cost (i.e., with minimum negative consequences). Thus, there may be corner solutions, emphasizing either consent or consequences alone.

Despite such cases, neither value alone provides a general explanation of moral intuitions or general direction in choosing actions and policies. There are often interactions, and the particular context at hand will shift the weight to the one or the other. The relative importance of consent or coercion depends on many contextual factors, including the magnitude of benefits from coercion, the costs of obtaining express consent, the costs and benefits of unbundling issues, the similarity to issues that have received higher forms of consent, etc. In evaluating the justification for coercion, the framework proposed here incorporates the four factors underlying consent and also stresses the importance of a pluralistic moral system. Other arguments employed in the discussion of consent alone carry over to the treatment of the multicriterion system: the nearly continuous and multiple moral factors, and the usefulness of this perspective even when the choice set is limited. Discrete jumps in policies and rationales for policies do not imply equally discrete moral factors. As before, we should not confuse the continuity of policies with that of the moral space or fail to appreciate the multivariate set of factors in the moral space, which include not only the four factors underlying consent but also potentially numerous consequences. Categorical and univariate approaches can lead to an artificial ordering of principles, e.g., a lexicographic one where one principle takes absolute priority over the other, and systems that produce inconsistencies with common moral intuitions.

5 *Conclusions*

This paper has proposed a framework to account for our moral intuitions about coercion and consent. Consent is defined in terms of our inferences about agent consent,

⁶ For an extensively developed argument about the moral importance of both consequentialist and nonconsequentialist concerns in the legal context, see Zamir and Medina (2010).

which are based on multiple moral factors and vary in fine degrees. The framework also rests on moral pluralism; in particular, I focused on consequences as the primary rival to consent, but not because consequences are the only moral principle that might conflict with consent. Rather, this was in the interests of brevity and because of the historical precedence and contemporaneous importance of consequences in debates over the use of coercion.

These arguments about consent and moral pluralism seek to reconcile *our* moral intuitions in these matters, but what if *we* have differing intuitions? I took the position of context dependence that, at a deeper level, we share moral values, even though differences are to be expected in moral judgments in specific contexts. Although I have often made reference to the views of “many” or “most” people, I have not presented a rigorous argument about what differences, qualitative or quantitative, are acceptable. That is an admittedly important but also very challenging question, and one I do not attempt to answer here. I would point out, however, that virtually every philosophical argument about morality (even moral relativism) supports its case with claims about moral intuitions. Moreover, I would argue that the disparity in moral views is frequently overstated, as least if one restricts attention to the informed and impartial views that are the focus of this paper. Experimental and vignette studies have found significant reduction in differences in moral views with increases in impartiality and information (Konow, 2009a,b). Lewinsohn-Zamir (2012) reports results from a vignette study that is specifically supportive on the topic of this paper: coercion reduces the perceived goodness of a given outcome, and this finding is similar across very different populations consisting of students, laypeople, and business people.

In addition, we should not, in my view, demand of a moral theory about consent that it generate an uncontroversial ranking of policies or an unambiguous “best” policy. Rather, it should fit actual moral intuitions, even when we are conflicted and opinions are sharply divided – indeed, *especially* in such cases, because those are the situations most in need of clear and accurate guidance to the underlying moral values. Policy debates on knife-edge issues such as the Affordable Care Act, abortion, and same-sex marriage are often dominated by strong commitments to monist positions, either consensualist or consequentialist, that create impediments to a deeper understanding of the issues and, therefore, to agreement. Applying a pluralistic system and employing the more balanced calculus that usually governs our moral judgments on less controversial matters might promote greater clarity and agreement on more contentious questions.

Further work could examine interactions between consent and other moral principles, including justice, which might alter some conclusions. For example, if people value only consent and the Pareto principle, then forgoing express consent seems innocuous if a policy generates no losers, but this conclusion is no longer obvious if people also value equity and the Pareto-improving policy also produces a distribution of gains that is viewed as unfair. In addition, consent might be incorporated into a virtue ethics framework, where the emphasis is on the development of character virtues. Such a framework might consider the effects of exercising consent on

character and generate conclusions about the desirability of such opportunities. For example, if the discretionary action of expressing consent bolsters virtuous traits such as generosity, this would give added moral weight to broadening consent. Finally, there are many possible applications of an approach that highlights consent in a pluralistic system that could be investigated, such as a consensualist motivation for voting and how it interacts with consequentialist considerations like the cost of voting.

References

- Alesina, A., and G.-M. Angeletos (2005), "Fairness and Redistribution," *The American Economic Review*, 95(4), 960–980.
- Alexander, L. (1996), "The Moral Magic of Consent (II)," *Legal Theory*, 2(3), 165–174.
- (2008), "Scalar Properties, Binary Judgments," *Journal of Applied Philosophy*, 25(2), 85–104.
- Attiyeh, G., R. Franciosi, and R. M. Isaac (2000), "Experiments with the Pivot Process for Providing Public Goods," *Public Choice*, 102(1–2), 95–114.
- Babcock, L., and G. Loewenstein (1997), "Explaining Bargaining Impasse: The Role of Self-Serving Biases," *The Journal of Economic Perspectives*, 11(1), 109–126.
- Bell, T. W. (2010), "Graduated Consent in Contract and Tort Law: Toward a Theory of Justification," *Case Western Reserve Law Review*, 61(1), 1–68.
- Buchanan, J. M., and G. Tullock (1962), *The Calculus of Consent: Logical Foundations of Constitutional Democracy*, University of Michigan Press, Ann Arbor.
- Clarke, E. H. (1971), "Multipart Pricing of Public Goods," *Public Choice*, 11(1), 17–33.
- Copp, D. (1999), "The Idea of a Legitimate State," *Philosophy & Public Affairs*, 28(1), 3–45.
- Falk, A., and M. Kosfeld (2006), "The Hidden Costs of Control," *The American Economic Review*, 96(5), 1611–1630.
- Frey, B. S., and A. Stutzer (2000), "Happiness, Economy and Institutions," *The Economic Journal*, 110(466), 918–938.
- Groves, T. (1973), "Incentives in Teams," *Econometrica*, 41(4), 617–631.
- Hume, D. (1752/2007), "Of the Original Contract," in: *The Philosophical Works of David Hume, Volume III*, Elibron Classics, Adamant Media Corporation, New York, pp. 494–517.
- Katz, L. (2011), *Why the Law Is So Perverse*, The University of Chicago Press, Chicago (IL).
- Kaye, D. (1982), "The Limits of the Preponderance of the Evidence Standard: Justifiably Naked Statistical Evidence and Multiple Causation," *American Bar Foundation Research Journal*, 7(2), 487–516.
- Konow, J. (2001), "Fair and Square: The Four Sides of Distributive Justice," *Journal of Economic Behavior & Organization*, 46(2), 137–164.
- (2003), "Which Is the Fairest One of All? A Positive Analysis of Justice Theories," *Journal of Economic Literature*, 41(4), 1188–1239.
- (2009a), "Is Fairness in the Eye of the Beholder? An Impartial Spectator Analysis of Justice," *Social Choice and Welfare*, 33(1), 101–127.
- (2009b), "The Moral High Ground: An Experimental Study of Spectator Impartiality," MPRA Paper No. 18558, University Library of Munich, <http://EconPapers.repec.org/RePEc:pra:mprapa:18558>.
- Lewinsohn-Zamir, D. (2012), "Taking Outcomes Seriously," *Utah Law Review*, 12(2), 861–902.
- Locke, J. (1689/2003), *Two Treatises of Government and a Letter Concerning Toleration*, edited by I. Shapiro, Yale University Press, New Haven (CT).
- Loewenstein, G., S. Issacharoff., C. Camerer, and L. Babcock (1993), "Self-Serving Assessments of Fairness and Pretrial Bargaining," *The Journal of Legal Studies*, 22(1), 135–159.

- Magid, L. (2001), "Deceptive Police Interrogation Practices: How Far Is Too Far?" *Michigan Law Review*, 99(5), 1168–1210.
- Pew Research Center (2008), *The 2008 Pew Global Attitudes Survey in China: The Chinese Celebrate their Roaring Economy, as they Struggle with its Costs: Near Universal Optimism about Beijing Olympics*, Washington (DC), <http://www.pewglobal.org/files/2008/07/2008-Pew-Global-Attitudes-Report-2-July-22-2pm.pdf>.
- Plato (360 B.C.), *Crito*, available at <http://classics.mit.edu/Plato/crito.html>.
- Russell, P. (1986), "Locke on Express and Tacit Consent: Misinterpretations and Inconsistencies," *Political Theory*, 14(2), 291–306.
- Steinhauser, P. (2009), "Poll Finds Lack of Support for 'Torture' Investigations," CNN/Opinion Research Corporation, <http://www.cnn.com/2009/POLITICS/05/06/bush.torture/>.
- Thaler, R. H., and C. R. Sunstein (2008), *Nudge: Improving Decisions about Health, Wealth, and Happiness*, Yale University Press, New Haven (CT).
- Tyler, T. R. (2006), "Psychological Perspectives on Legitimacy and Legitimation," *Annual Review of Psychology*, 57, 375–400.
- Zamir, E., and B. Medina (2010), *Law, Economics, and Morality*, Oxford University Press, Oxford.

James Konow
Chair of Economics and Ethics
Kiel University and
Kiel Institute for the World Economy
Olshausenstrasse 40
24098 Kiel
Germany
and
Department of Economics
Loyola Marymount University
Los Angeles, CA 90045-2659
U.S.A.
jkonow@lmu.edu