Moral Learning in the Open Society:
The Theory and Practice of Natural Liberty

Gerald Gaus and Shaun Nichols

I. LIBERTY AS A DEFAULT?

Many liberal political philosophers have claimed that moral and political life rests on a principle of "natural liberty" — a general presumption in favor of freedom of action. As Joel Feinberg puts it, "liberty should be the norm, coercion always needs some special justification."[1] John Rawls advances a wider specification, identifying "restrictions" (rather than only "coercion") as requiring justification: "there is a general presumption against imposing legal and other restrictions on conduct without sufficient reason."[2] Some of Mill's specifications are even less focused on the law: "in practical matters, the burthen of proof is supposed to be with those who are against liberty" — liberty is the default; its limitation requires justification.[3] Perhaps the widest specification is Locke's: "all men are naturally in … a State of perfect Freedom to order their actions ... as they think fit ... without asking leave, or depending upon the Will of any other Man."[4]

A number of philosophers — both from within and without the liberal tradition — have rejected this principle. Some suspect that it is covertly libertarian, hostile to claims of distributive justice (this, despite the fact that Rawls, Feinberg, and Mill and other non-libertarians endorse it). To others, it is simply another piece of liberal ideology, seeking to insinuate into the very foundation of moral thinking a claim

that (negative) liberty has a special status, and is somehow more basic than positive liberty, welfare, equality or justice. Indeed, the principle seems to suppose that law and even morality itself are somehow inherently liberal; once one understands the structure of legal and moral reasoning and justification, one is committed to the view that liberty is the default. The opponents of liberalism and negative liberty seem to lose out before we ever get to substantive argument, simply by appeal to the nature of legal and moral thinking.

We focus on this last claim: is a commitment to liberty as the default somehow implicit in our moral thinking? We know that philosophers disagree. Our concern in this paper is how normal human reasoners answer this question. When people reason on the basis of moral rules, do they suppose that, in the absence of a prohibitory rule they are free to act, or do they suppose that morality always requires a justification establishing a permission to act? Our answer is mixed: we shall show that rule followers sometimes suppose a principle of natural liberty, sometimes its denial. Yet, we shall argue, there is reason to believe that rule followers are inclined — in some cases strongly inclined — to suppose that natural liberty is indeed the default. Both systems seem possible: is there a reason to adopt one or the other? We argue that systems of social morality based on a principle of natural liberty have a decisive advantage over their competitors: they are well adapted to effectively exploring the constant novel circumstances that arise in open, dynamic, societies.

Our analysis is divided into three parts. In section II we clarify the principle of natural liberty, distinguishing several versions, and sketch some criticisms that have been raised against it. Section III considers how a principle of natural liberty enters into moral learning and rule following behavior. Here we present empirical research concerning the conditions under which moral learners conclude that their system is, or is not, one based on a default of liberty. Finally, in section IV we argue that systems of moral rules premised on a principle of natural liberty have decisive advantages in helping us to navigate, and to expand, the constant novelty that confronts diverse, open, societies.

II. LIBERTY: NATURAL OR JUSTIFIED?

A. The Minimal Principle of Natural Liberty

The intuitive attraction of a principle of natural liberty was perhaps best articulated by S. I. Benn, who asked us to

> [i]magine Alan sitting on a public beach, a pebble in each hand, splitting one pebble by striking it with another. Betty, a casual observer, asks him what he is doing. She can see, of course, that he is splitting pebbles; what she is asking him to do is to explain it, to redescribe it as an activity with an intelligible point, something he could have a reason for doing. There is nothing untoward about her question, but Alan is not bound to answer it unless he likes. Suppose, however, that Betty had asked Alan to justify what he was doing or to give an excuse for doing it. Unlike explanations, justifications and excuses presume at least prima facie fault, a charge to be rebutted, and what can be wrong with splitting pebbles on a public

beach? Besides, so far as we can tell, Alan is not obliged to account to Betty for his actions....

Suppose Betty were to prevent Alan from splitting pebbles by handcuffing him or removing all the pebbles within reach. Alan could now quite properly demand a justification from Betty, and a *tu quoque* reply from her that he, on his side, had not offered her a justification for splitting pebbles, would not meet the case, for Alan's pebble splitting had done nothing to interfere with Betty's actions. The burden of justification falls on the interferer, not on the person interfered with. So while Alan might properly resent Betty's interference, Betty has no ground for complaint against Alan.[5]

Let us call this the Full-fledged Liberal, or the

Strong Principle of Natural Liberty:

(1) (*a*) Alan does not need a permission granted by any moral rule *R*, in order to perform an action, ϕ, without moral fault.

(*b*) So he is not always morally required to cite a permission to ϕ in the face of a moral challenge to his ϕ–ing by Betty.

(*c*) In this sense Alan is under no standing obligation to justify his ϕ-ing.

(2) (*a*) Interference with, or restriction of, Alan's ϕ-ing requires justification.

(*b*) Unless Betty, the interferer, can cite a moral permission to interfere, she is normally morally blameworthy for her interference.

According to the Strong Principle of Natural Liberty, in Benn's pebble-splitting case Alan properly resents Betty's interference (because of clause 2), while she has no ground of complaint against him — there is no presumptive fault, so no call for Alan

to justify himself (because of clause 1).

Note that the statement of the Strong Principle of Natural Liberty appeals to the idea of moral rules. The focus is on what a number of philosophers have called "social morality": a system of social rules that guides behavior and forms a basis for shared normative expectations about what one may or may not do, and what one can demand that others refrain from, or must do, and shared empirical expectations as to whether people will conform to these rules.[6] There are many sorts of social rules: descriptive rules (what side of the sidewalk do we walk on here?), mere social norms (say, norms of deference in class cultures) and moral rules. When agents see social rules as moral, they not only understand them as the basis of coordinated expectations, but their normative attitudes and personal moral beliefs support these rules, and they internalize them: they see the rules as internal demands and will experience guilt when they fail to conform, and resentment and indignation when others fail.[7] As Kurt Baier stressed, these rules must be teachable, and indeed they must be taught. "People are neglecting their duties if they do not teach these moral rules to their children."[8] Equipped with these, a child can navigate social life, being aware of the normative expectations of others (what others think the child ought to do), and when the child can expect others will be sensitive to her normative demands. When these conditions are met within a group or social network, a moral rule R is the basis for mutually acknowledged moral demands, and expectations about behavior. Rules of social morality are quintessentially about the sorts of social

interaction that is crystalized in Benn's example: one person is acting and another challenges his action, demanding that he desist.[9] Here, unless there are shared expectations about what is appropriately demanded, and what sort of action is permitted or prohibited, the parties will fail to successfully coordinate their actions and demands.

To be sure, some of the rules that we might identify as "moral" are perhaps better understood as personal rules or norms.[10] Whereas a rule R is a rule of social morality only if R is a rule that some group or social network accepts and generally acts upon, sometimes a moral rule is conceived of as an unconditional personal commitment: if Alan endorses an unconditional rule against lying, Alan will not lie, regardless of what others do or what they expect of him. Such personal rules are also taught from one generation to the next, and form a part of a person's moral code. The Strong Principle of Natural Liberty focuses too much on interpersonal claims and challenges to capture this personal aspect of morality. However in teaching such a code to one's child, clause 1a of the principle is obviously relevant, i.e., what we will call

> *The Minimal Principle of Natural Liberty*: Alan does not need a permission granted
> by any moral rule R, in order to ϕ without moral fault.

This rule certainly can be supposed in teaching a personal moral code to one's child, even if no one else in the society accepts it. The Minimal Principle of Natural Liberty

is the crux of the idea that liberty is the default: when Alan φs, it is not the case that he *must* be permitted by some moral rule R to φ if he is to φ without moral fault. Thus, he can be morally free to φ, even though he could cite no rule (either social or personal) justifying his φ-ing, indeed even if there *is* no rule granting permission. Given this apparently modest principle, we can see why Betty cannot assume that, as a matter of course, she can advance a well-grounded moral demand that Alan justify his action, for he need not always possess a justification.

B. Rejecting the Minimal Principle of Natural Liberty

Somewhat surprisingly, objections to the Full-fledged Liberal Principle of Natural Liberty often seem focused on the Minimal Principle. In rejecting the principle of natural liberty, Ranier Forst insists:

> "[a] In a *moral* context, a person can be called autonomous only if he or she acts on the basis of reasons that take every other person into account, so that these reasons are mutually justifiable. [b] Wherever the actions of a person affect others in a morally relevant way, they must be justifiable on the basis of reciprocally and generally binding norms, and therefore all those affected — individually — can demand that the agent justify his action on the basis of reasons that are 'not reasonable to reject'...."[11]

A Kantian Criticism. In part [a] of the above quote Forst claims that if *M* is a moral context, an autonomous Alan must only act on reasons that are justifiable to Betty. Thus it cannot be the case that there is any moral context *M,* such that a person in *M*

need not justify himself to others. So it seems Forst's claim must be that the Minimal Principle of Natural Liberty cannot be part of the moral sphere (it cannot operate within the moral domain); if the Principle of Natural Liberty operated within a moral context *M* then it would be false that in all moral contexts Alan must be able to justify his action to Betty. Thus because on Forst's account morality always requires justification and the Minimal Principle of Natural Liberty denies that justification is always called for, both cannot be correct about morality.

On closer examination it is not clear that Forst is successful in rejecting Minimal Natural Liberty. Forst's specification of the moral context in [b] may be read as advancing:

> *The General Prohibition Principle*: There is a general prohibition on acting without justification *when one's actions affect others*.

But then it must be the case that Alan does not need a permission granted by any moral rule *R*, in order to ϕ without moral fault (Minimal Natural Liberty), unless ϕ-ing falls under the General Prohibition Principle. More generally, any time a system specifies a *triggering condition* for moral justification or, more generally, morality (e.g., affecting others, affecting their welfare), it is implied that when the condition is not met, the person is free to act without having a moral justification, and that simply *is* the Minimal Principle of Natural Liberty.

Perhaps, it might be thought, a divide between the moral and the nonmoral

context can, after all, be reconciled with the rejection of the Minimal Principle of Natural Liberty. If there is a divide between the moral and non-moral contexts, and if the Minimal Principle of Natural Liberty is on the "non-moral side of this divide" (where the trigger for moral justification has not been pulled), the Principle of Natural Liberty is not, by definition, part of morality, so it cannot be a "moral liberty." And this may be what Forst has in mind, as he insists that all moral claims to liberty require justification: the only plausible "presumption" in favor of liberty, he tells us, would be one of "equally *justifiable* liberty." [12] But justification is constitutive of the moral context M, so no liberty "outside" of M can have a moral dimension. But this is to misconceive a moral liberty — perhaps to confuse it with a moral claim to freedom. On a fairly straightforward Hofeldian-inspired analysis, a moral liberty to ϕ simply is equivalent to a "no moral duty not to ϕ."[13] So a sphere in which no moral duty can apply is, by definition, a sphere of moral liberty. To say that, given a principle of natural liberty, morality permits one to proceed in that sphere is not to claim that morality justifies an explicit permission, much less a claim on others to allow one to proceed — only that morality is silent.

In the end, there are two ways to revise Forst's complaint so that it truly rejects the Minimal Principle of Natural Liberty— both drop the moral context condition. (*i*) What we might call a "one sphere" solution simply requires that all actions — including Alan's pebble splitting — must be justified to others. This is the straightforward interpretation of rejecting "natural" liberty: there is only "justified"

liberty, because all liberty is within morality. (*ii*) Alternatively a "two sphere" solution would identify a moral sphere in which the General Prohibition Principle is justified and another (basically self-regarding sphere) in which a principle of minimal liberty is explicitly morally justified. With both solutions there is no line between the moral and what lies outside of it. In the one sphere solution there is no division. In the two sphere solution there is a line *within* morality between self-regarding freedom, where *further* justification is not required, and the General Prohibition Principle, where it is. On either proposal there is no room for natural liberty as a default, for all possible actions are sorted into a moral category; either a single one or two mutually exclusive and exhaustive parts of the fully justified moral domain.

A Welfarist Criticism. We have focused on what may be broadly called Forst's "Kantian" rejection of natural liberty.[14] The other common route to rejecting natural liberty is via a welfarist or utilitarian standard of evaluation. Thus Claude Hillinger and Victoria Lapham insist, "[w]henever the choices of one individual impinge on the welfare of others, there is no general presumption in favor of freedom of individual choice."[15] On the face of it we are simply back to the "other-regarding" triggering condition, but suppose a certain welfarist view where (*i*) the aim is always the maximization of individual welfare, and (*ii*) we suppose that individual rational action always seeks to maximize personal welfare. Then, if no one else's welfare is

relevant, and so the person goes ahead and maximizes her own welfare, that just *is* the best way to maximize welfare. So here we may say that there is after all a justification of the liberty: if the person cares about her own welfare, giving her the liberty to pursue it is justified in the self-regarding case. In any event, in moral philosophy classic act-utilitarians are clear that for all possible acts of an individual, he is obligated to do that which best maximizes overall utility, so there can be no question of a "natural liberty" that is not justified by the utilitarian standard.[16]

C. Two Desiderata for Social Morality

Recall that our primary interest here is a system of social-moral rules that provides the basis for sustained network of shared expectations and commonly recognized claims. As Rawls would say, our concern is a social practice that provides for the public adjudication of our claims in a system of social cooperation.[17] Kurt Baier is certainly correct that one desideratum for such a moral practice is that it can be transmitted from one generation to the next. "Morality is meant to be taught to all members of the group in such a way that everyone can and ought always to act in accordance with these rules."[18] If we wish to understand the place of natural liberty in a system of social morality, we should consider its place, not in a scheme of vague and abstract categorizations or principles of morality, but in a teachable, and so learnable, system of moral rules. We need to inquire what learnable systems of social morality which include it — and reject it — look like. And we should examine how

moral learners come to the conclusion that the Minimal Principle of Natural Liberty

characterizes morality, and under what conditions they will agree with critics like

Forst and act utilitarians, that morality denies any such liberty. And if one system is

learnable and the other is not, or if one is more easily learned than the other, we may

find that one or the other system is better suited to play morality's expected role in

human life.[19]

This leads to the second desideratum for a social morality. In the above quote

Baier argues that it ought to be the case that the taught rules, at least in principle, can

always be complied with. Now without accepting all that this might imply, it does

seem that an acceptable set of moral rules should generally be *decisive*, in the sense

that for any action ϕ, the system of morality allows, requires or prohibits ϕ (with the

"or" being exclusive). Formally, this is a consistency condition — for any action ϕ,

the moral system is single-functioned, providing one and only one answer.[20] If the

moral system is not single-functioned in this way, it is in principle impossible that a

person can always comply with it. This is obvious if the system fails to be single

functioned because, for some action ϕ, it yields both the outputs in the same

circumstances that ϕ is prohibited and that it is required. Of course the system may

yield "prima facie" or "pro tanto" opposing judgments, which need to be weighed

and, as a practical matter, it may be very difficult to do the weighing, but these

complexities are all consistent with the requirement that the system is not, in

principle and all-thing-considered, "over-complete,"[21] i.e., inherently inconsistent.

Another failure of decisiveness, under-completeness, may seem less worrying. Here, in a context in which φ is a feasible option, the system does not generate any of the judgments: φ is required, φ is prohibited, or φ is permitted. In this context under-completeness also makes it impossible to comply with the system: if the system is under-complete among "φ is permitted, φ is required, φ is prohibited" one cannot "always to act in accordance with these rules," for the rules do not entail any instructions. One cannot go ahead and φ in this case, for that would be to suppose that φ is permitted, and that is what the system is incomplete about. And one cannot simply refrain from φ-ing, for the system is incomplete about whether φ is required.

Decisiveness is typically assumed in systems of deontic logic: if an action is not prohibited or required, then it is permitted, and it cannot be both permitted and prohibited. Now to be sure, real-world systems of social rules do not possess such logical clarity. Social rules are sometimes ambiguous, and a person can satisfy normative expectations in different ways. [22] Nevertheless, for social morality to generally provide a system of shared expectations, ambiguity must be held firmly in check. As Rawls insisted, a fundamental function of a system of social ethics is to provide a decisive resolution of disputes.[23] To achieve coordination in cooperative contexts the rules must guide our expectations, so that one can anticipate whether a person must, must not, or may φ, and there is also a shared understanding when one may demand that another φs (or refrains) and when one appropriately blames others for what they have done. A social morality that lacks decisiveness will thus be

flawed, given the very job that we wish our shared morality to perform.[24] While some ambiguity may well be not only unavoidable but useful, systems of high ambiguity simply cannot fulfill their coordinating role.[25] Thus decisiveness is a desideratum, though complete decisiveness is not a requirement. Like cricket???, the rules of morality may sometimes be unclear, but unless the rules typically provide unambiguous guidance, the game cannot be played.

If we focus on social morality — a system of shared empirical and normative expectations within a group that allows coordination of behavior — it is clear the great merits of a generally decisive system that can be taught and learned. However, even if we take a simple personal view of morality — according to which moral rules are a person's unconditional rules of behavior — decisiveness is a great virtue. When the rules are not decisive a person who seeks to always conform to morality can experience decisional paralysis: she confronts some feasible set of actions and cannot select from it. Here one not only confronts incompatible options that one cannot rank, but the selection of any option implies that one may have done wrong because one failed to select another.[26]

III. NATURAL LIBERTY AND MORAL LEARNING

A. Teachable Decisive Morality: Closure Rules

In order for a system of morality to be teachable and decisive, it must contain rules that (*i*) are sufficiently fine-grained to generally secure coordination of expectations

as to what others will do and (*ii*) close the system so that it is always decisive.[27] To

be sure, (*ii*), a closure principle, is not formally required for a system of rules. If all

action types were specified, then theoretically each type could be binned into the

prohibited or the permitted. Depending on the number of action types and the

capacity of the learner, this list could be internalized by rote. In the case of human

rule systems, however, it seems that the list of action types is vast and the cognitive

resources limited.[28] Real human reasoners can only learn a relatively modest set of

rules, and they do not possess the cognitive resources to infer all the possible

applications of the set of rules. Moreover, because moral learning is gradual, it is

important for the human learner to have some way of making a reasonable

prediction about whether an action of a new type falls into the prohibited or the

permitted class. Hence, it is likely that people do make some kind of inferences that

bear on the rest of the rule system. The sort of inference in which we are interested

are closure rules: if in circumstances *C* the set of moral rules do not instruct whether

ϕ is prohibited, required or permitted, then a default rule is invoked which implies

one may ϕ, one may not ϕ or one must ϕ. Drawing on this idea, John Mikhail

identifies one such closure rule,

> *The Closure Principle of Natural Liberty*: Whatever is not prohibited (and this
>
> includes the non-performance of specific acts) is permitted.[29]

On this closure rule, an agent consults the system of rules and determines whether

her ϕ-ing is prohibited by some rule in the system; if it is not, then she is free to ϕ. Thus the Closure Principle of Natural Liberty implies the Minimal Principle of Natural Liberty, i.e., Alan does not, as a matter of course, need a permission granted by any moral rule *R,* in order to ϕ without moral fault. As Mikhail notes, we can contrast this principle of natural liberty to another closure rule:

Residual Prohibition Principle: Whatever is not permitted is prohibited.

Under this closure rule Alan is prohibited from ϕ-ing unless the system has a rule permitting ϕ–ing. The Residual Prohibition Principle is inconsistent with the Minimal Principle of Natural Liberty: if there is no permission rule *R* that allows Alan to ϕ, he violates a prohibition and so is at moral fault. If, accepting the need for a closure rule, a moral theory or conception of rules rejects the Closure Principle of Natural Liberty, it is committed to the Residual Prohibition Principle. If an agent does not have a liberty to ϕ (i.e., one does not possess a "no duty not to ϕ") when the rules are silent, then it must be the case that when the rules are silent she *does* have a duty not to ϕ, that is, ϕ is prohibited.

We now can reanalyze the debate about the Minimal Principle of Natural Liberty as a comparative analysis of two different types of moral rule systems, those based on the Closure Principle of Natural Liberty and the Residual Prohibition Principle. Moralities, we have been insisting, must be teachable and learnable, and so must appeal to one or the other closure rule. Thus we may wonder, are both the types of

moral systems learnable?

B. The Experiments

People have knowledge of numerous rule systems in numerous domains, and different closure principles apply in different domains. This is easy to see by considering games. Nowhere in the rules of chess does it forbid replacing a rook with a wizard who can kill the king with his gaze. Yet this is obviously prohibited in chess. Whatever is not expressly permitted by the rules of chess is prohibited. Of course this still leaves a great deal of flexibility, otherwise chess would not be very interesting. But here we have a game that is closed under residual prohibition – any move that is not expressly permitted is prohibited. Consider now the childhood game of not letting a balloon touch the floor. Apart from the fact that you lose if the balloon touches the ground, there is but one rule: no holding the balloon. Anything else goes. It is okay to hit the balloon with the shoulder, knee, head, elbow, etc. It is okay to stand on one's head to hit the balloon. While Mom might get mad if you stand on the kitchen table to hit the balloon, that too is allowed by the lights of the game. Except for holding the balloon, anything goes. In the balloon game, the rules are closed under natural liberty (in the game).

There are a number of ways to learn closure rules. Often, rather than being explicitly taught, learners infer them. We will suggest that at least in some cases,

these inferences are rational. Our focus here is one way in which learners make such inferences, through *pedagogical sampling*. Sometimes a learner has reason to think that the training examples she receives come from a teacher who is trying to help her.[30] This procedure can be part of a normatively rational procedure. If Betty knows that Alan is trying to teach her a rule (say), and she knows that Alan assumes that she is rational, then Betty should expect Alan to give examples that are maximally useful for a rational agent trying to acquire the rule.

Recent evidence indicates that people do in fact engage in learning based on an assumption of pedagogical sampling. In one study, teachers were told to help learners determine the dimensions of a rectangle by indicating with a green circle a positive example (falling inside the rectangle) and by indicating with a red X a negative example (falling outside the rectangle).[31] A good teacher who assumes that the learner is rational will pick near misses. The positive example should be just inside one corner of the rectangle, and the negative example should be just outside the opposite corner. If a learner assumes that the teacher has done this, she will choose a rectangle that barely includes the positive example and barely excludes the negative example. And in fact, in these tasks, learners do choose rectangles that stretched from the positive example in one corner to the edge of negative example in the other corner. This suggests that the learners are, rightly as it happens, assuming that the examples were chosen with the learner in mind.

Turning to rules, suppose a teacher has a limited amount of time to teach

someone a rule system. There are, of course, many factors to consider. For instance, if it is much worse to commit a violation than to forgo an option, the teacher might focus on the violations. But in general, one prevailing factor in good teaching will be efficiency. Typically, when the closure principle for a rule system is Residual Prohibition, there will be many more action types prohibited than permitted; when it is the Closure Principle Natural Liberty, many more action types will be permitted than prohibited. Thus, if a teacher wants the learner to infer a closure principle, it will be efficient for the teacher to provide examples from the smaller set and anticipate that the learner will infer that the closure rule applies to the larger remainder. If the teacher knows that the relevant principle is the Closure Principle of Natural Liberty, the teacher should provide examples of prohibitions, since the learner will infer that there is a larger remainder of permissions. If the relevant principle is the Residual Prohibition Principle, the teacher should provide examples of permissions. Accordingly, if the learner assumes that the teacher is trying to provide evidence that would be most efficient for a rational learner, then the learner should infer the Residual Prohibition Principle when the teacher provides permission rules and the Closure Principle of Natural Liberty when the teacher provides prohibition rules.

To explore this issue empirically we ran a series of small learning studies. In the first two studies participants were told that they would be learning about a game from a teacher and that they would have to use this information to determine other

rules of the game. In both initial studies, we presented participants with either two permission rules or two prohibition rules. Then we had them determine whether a new action would be prohibited or permitted.

In the first study, they were told about a ball game. The instructions in the prohibition-training condition went as follows (in the permission-training condition, the "not" was omitted.):

Before you play Game X, the teacher tells you two rules.

1. You are not allowed to bounce the ball on the ground.

2. You are not allowed to hold onto the ball.

The teacher says that while these aren't the only rules, it's enough to get you started.

Participants were then asked whether they thought "you are or are not allowed to toss the ball in the air," using a 7-point scale from "You are not allowed to toss a ball into the air" to "You are allowed to toss a ball into the air." We predicted that participants in the prohibition-training condition would be more likely to say that you *are* allowed to engage in the new action type. That is exactly what we found. Participants in the prohibition condition tended to say that it is permissible to toss the ball in the air (*M*=5.17), which was significantly different from the responses in the permission-training condition (*M*=3.94, *t*(112)=2.9, *p* < .01).[32]

To reinforce this result, we ran a second study in which we used a board game

for the rule system. This also afforded a very closely matched set of training examples. As before, participants were told that they were being taught a game and they needed to infer other rules from limited information. In the prohibition-training condition, they received the following (again, the "not" was omitted in the permission-training condition):

1. You are not allowed to put a red peg on a black square.

2. You are not allowed to put a green peg on a black square.

Participants were asked to indicate whether you are or are not allowed to put a blue peg on a black square. As in study 1, we found a clear difference in the predicted direction. Participants in the prohibition condition tended to say that it is permissible to put a blue peg on a black square (*M*=4.8), and this was significantly different from the responses in the permission-training condition (*M*=3.3, *t*(47)=2.82, *p*<.001).[33]

Thus the results confirm our hypothesis that participants will be more likely to infer the Closure Principle of Natural Liberty when given only prohibition rules. Our prediction was based on the idea that participants would exploit the assumption of pedagogical sampling — that the teacher was trying to communicate the rule system in an efficient manner. In keeping with this idea, we found that several participants, in both conditions, explicitly appealed to what the teacher would have done to optimize teaching, e.g., "The teacher would have told me in the

instructions" (prohibition-training condition); "I would think the teacher would have told us if blue was allowed to go on a black square" (permission-training condition).

The previous studies focus on rules of games, but we also wanted to explore something closer to a normative rule system.[34] In a third study, participants were told the following:

> There is a Farm with squeaky mice, and all the mice are supposed to follow The Rules of Mice, written in a book. The Farm has four barns: Red, Blue, Yellow and Green.

In the prohibition-training condition, they received the following:

> The Rules of Mice Book says
>
> 1. Squeaky mice are not allowed to be in the Red Barn
>
> 2. Squeaky mice are not allowed to be in the Yellow Barn.

In the permission-training, of course, the "not"s were omitted. We had a third, mixed, condition with these rules (counterbalanced for order):

> 1. Squeaky mice are not allowed to be in the Red Barn
>
> 2. Squeaky mice are allowed to be in the Yellow Barn.

Participants were asked to indicate (on a 7-point scale) whether or not mice are allowed to be in the Green Barn. As in studies 1 & 2, we found a clear difference in the predicted direction. Participants in the prohibition condition tended to say that it

is permissible for the mice to be in the Green Barn (*M*=5.73), and this was

significantly different from the responses in the permission-training condition

(*M*=3.16, *t*(110)=7.8, *p*<.0001). The mixed condition produced results in the middle

(*M*=4.33), showing significant difference from both the prohibition-training

condition (*t*(125)=4.6, *p*<.0001) and the permission training condition (*t*(138)=3.76,

p=.0002).[35]

Responses differed from chance in the prohibition-training condition (one

sample t-test *t*(48)=6.92, *p*<.0001) and in the permission training condition

(*t*(61)=3.194, *p*=.0022). This shows that, in simple cases, people can easily learn either

a residual prohibition rule (when trained on permissions) or a principle of natural

liberty (when trained on prohibitions). Interestingly, responses in the mixed-

training condition showed a trend in the direction of thinking the new act would be

allowed — indicating an inclination towards inferring the Closure Principle of

Natural Liberty, but this did not reach significance (one sample *t*(77)=1.8, *p*=.07 n.s.).

In light of the trend in the mixed case, we ran one final experiment. In the

previous experiments, we used a 7-point scale, which gives participants the

opportunity to choose the midpoint under uncertainty. In the mixed case, this was a

very common response. In real life, the learner often cannot afford to stay on the

fence. So we ran a study with a 6-point scale, forcing participants to pick a side.[36] In

addition we made explicit that there were only two stated rules, by saying "The

Rules of Mice Book has only two rules." Note in this case subjects cannot be

understood as guessing what a missing rule might be; they know that this is the complete set of rules, so they are being induced to directly determine the closure rule. It is also important to stress that they are not being taught a closure rule, or even instructed that there is one. In this experiment subjects knew that they had received the entire set of rules, so they cannot be taken as guessing what the missing rule might be. And critically, it cannot be the case that a closure rule is part of the rule system — if they employ a closure rule it must be the case that it comes from outside the system. Neither Natural Liberty nor Residue Prohibition can be a justified rule in the set. Yet, when induced to take sides by the six-point scale, subjects show a strong tendency to close the system. Responses differed from chance in the prohibition-training condition (M=5.5; one sample t-test $t(27)$=13.33, p<.0001) and in the permission training condition (M=1.8; $t(26)$=5.05, p<.0001). And as before, there were significant differences between all conditions.[37] Most importantly, in this task, responses in the mixed condition also differed from chance (M=4.79, $t(41)$=4.84, p<.0001). Participants were markedly more likely to think that the new action was allowed when given an example of one prohibition rule and one permission rule.[38]

These results indicate that people can infer closure principles in rationally appropriate ways, based on limited evidence. In addition, it is clear that people can learn different closure principles for different domains. There is not a single closure principle that applies to all rule systems. Social (or personal) morality itself could well be divided into different domains with different closure rules; we have said

nothing to indicate that all social (or personal) morality is of one piece. That is a topic for further investigation. In what follows we will speak of "morality" as if it was only one domain, but it should be stressed that the experiments do not speak to that issue.

IV MORAL RULES, UNCERTAIN ENVIRONMENTS, AND INNOVATION

A. Bodo Ethics

Our experiments suggest that systems of moral rules based on either closure rule are in principle teachable and decisive, though it does seem that in the face of ambiguity learners lean toward natural liberty. Those who primarily learn prohibition rules, our data indicates, tend to treat their morality as being one of natural liberty rather than residual prohibition. Thus a morality primarily of prohibitions will tend to be one in which participants infer Minimal Natural Liberty, while those who have learned moralities focused on permission rules will tend to infer that what morality does permit, it prohibits — i.e., unless there is an explicit permission, morality prohibits an action. Thus those who endorse the two desiderata for social morality (§II, *C*) and would follow Forst in his rejection of natural liberty, have reasons to teach a social morality of moral permissions, and so encourage moral learners to accept the alternative closure rule, the Residual Prohibition Principle. We have shown that they may well be able to do this. Should they wish to?

It is hard to see why. Prohibitory systems of moral rules grounded on the

Closure Principle of Natural Liberty (and so which include the Minimal Principle of Natural Liberty) have a decisive advantage over the permission-based systems that lead leaners to reject natural liberty: they allow individuals and groups to adjust to quickly changing and complex environments, and provide a hospitable morality for what Mill called "experiments of living."[39] To see — in perhaps an overly-stark way — the dynamic contrast between moralities of natural liberty and residual prohibition, consider Russell Hardin's charge that rule-based ethics are versions of "Bodo ethics;"

> Axel Leijonhufvud… characterizes the village society of eleventh century France in which the villager Bodo lived. We have detailed knowledge of that society from the parish records of the church of St. Germaine. Today one would say that that church is in the center of Paris, but in Bodo's time it was a rural parish distant enough from Paris that many of its inhabitants may never have seen Paris. Virtually everything Bodo consumed was produced by about eighty people, all of whom he knew well. Indeed, most of what he consumed was most likely produced by his own family. If anyone other than these eighty people touched anything he consumed, it was salt, which would have come from the ocean and would have passed through many hands on the way to St. Germaine, or it was spices, which would have traveled enormous distances and passed through even more hands.[40]

Hardin advances a variety of criticisms of Bodo ethics; one important claim is that in Bodo's world a rule-based morality made perfect sense, as his life confronted "a relatively limited range of kinds of actions."[41] Given this, a set of rules that identified

all the available action-types, and then sorted them into the categories of the prohibited, required or permitted, was a plausible basis for social order. Indeed, perhaps a system of rules in such a very simple social world could do without a closure rule: if the available action-types are few and very stable, perhaps the system could be closed without any default rule. But even Bodo ethics probably could not identify all the available act-types in the village, and even in this very stable world novel events occasionally occurred. For Bodo ethics, the Residual Prohibition Principle is entirely appropriate. All the useful or appropriate act-types could be identified in the course of village history; given its stable environment the justifiability of these act-type also would be stable. The Residual Prohibition Principle would then close the system and better secure social coordination, by sorting any act-type not identified in the system of rules as prohibited.

Hardin insists that in our social world, with its immensely "richer range of activities," technologies and scale, and where new act-types are constantly arising (e.g., blogging, cyber-terrorism, intellectual theft), a rule-based morality cannot be the primary mode to secure social order and coordinate our expectations.[42] Indeed, Hardin believes that Kantian ethics — understood as action based on rules — is sort of latter-day Bodo ethics.[43] Hardin's critique is insightful concerning Residual Prohibition Systems. As new act-types arise (i.e., those that are not categories in the original set of rules), they will be sorted into prohibitions: act-types that are not on the list of the permitted are prohibited. Thus new ways of acting are morally

prohibited. To be sure, eventually the system might be revised so that these new act-types are specifically sorted into, say, the permitted category, but that process of explicit revision will be relatively slow, and will never comprehensively categorize all the new act-types that arise in dynamic societies. A residual prohibition system is thus be highly conservative, and will have great difficult adapting to new environments, in which the social or individual value of engaging in certain action types fluctuates.

It might be thought that a residual prohibition system could cope with the emergence of new act-types through analogy or similarity. This is, indeed, a way in which change occurs in rule systems: a new type of action is reinterpreted as akin to some familiar one that is covered by an existing rule. In 2015, for example, the United States Federal Communications Commission categorized some activities of cable companies as public utilities,[44] thus drawing cable companies' provision of internet services under rules originally designed for telephone providers. In similar ways, "blogging" can be categorized as publishing, as "cyber-terrorism" as violence. And therein lies the problem. In all these cases the extensions of current categorizations are uncertain and controversial. There is real disagreement about these matters, and that is why centralized authoritative bodies such as the Federal Communications Commission make authoritative (and often highly controversial) rulings. In the informal system of social regulation we call social morality, when the Residual Prohibition Principle is operative, new act-types are thus either simply

prohibited (and so the morality is conservative) or are reinterpreted by individuals, thus rendering scope of the rules highly uncertain, and so undermining their teachability and decisiveness.

In contrast, a system of natural liberty will sort new act-types as permissible; one is free to engage in a new type of action that is not covered by existing prohibitions. Now we do not wish to suggest that systems of natural liberty are entirely immune to disputes about how to categorize new actions (which is why, presumably, Hardin thinks all rule-based moralities are infected with the limits of Bodo ethics). Consider a rule that was clear in, say, 1980: one has the right to control information about (legal) activities that occur in one's own home. Others are thus prohibited from obtaining and using this information for commercial purposes without one's consent.[45] As the new action type of web browsing arose, and Google can use information about one's browsing history (that in one way is) "on one's home computer" to select ads for future viewing, some argue that this new act-type should be included in the older prohibition against "commercial spying," while others resist this analogy.

There is, however, a critical difference between the two systems, even in the face of uncertainty and dispute. In a natural liberty system, if a moral innovator does not conclude that the analogy holds, he will conclude that morality allows his innovative activity; in a residual prohibition system, unless the innovator concludes that a relevant analogy holds — that the new action type is analogous to a permitted type

— she will desist. It is in this sense that a natural liberty system encourages experimentation and discovery. The moral experimenter — one is who doing new things in new ways — need not first convince himself that the new action type falls under a previous permission; he proceeds as long as he does not conclude that the new type falls under a current prohibition. Given that in many circumstances any analogy to previous types of action will be very imperfect and so uncertain, this asymmetry is of great significance, freeing the innovator from proving (to himself) his freedom to discover.

Thus far we have been considering only the point of the view of the innovator, who consults his own understanding of morality, and checks whether, as he understands it, morality gives his innovative activity a green light. However, as Mill stressed, "[g]enius can only breathe freely in an atmosphere of freedom."[46] If one expects that others who, say do uphold the disputed analogy, will believe one's action to be prohibited and act on their belief, then the innovative activity is apt, after all, to be checked. Worry about this sort of interpersonal case drives us back to the Strong, or Full-fledged Liberal, Principle of Natural Liberty, which we discussed at the outset, viz.

> (1) (*a*) Alan does not need a permission granted by any moral rule *R,* in order to ϕ without moral fault.
>
> (*b*) So he is not always morally required to cite a permission in the face of a moral challenge to his ϕ–ing by Betty.

(*c*) In this sense Alan is under no standing obligation to justify his ɸ-ing.

(2) (*a*) Interference with, or restriction of, Alan's ɸ-ing requires justification.

(*b*) Unless Betty, the interferer, can cite a moral permission to interfere, she is morally blameworthy for her interference.

The Strong Principle of Natural Liberty is an interpersonal principle of justification, putting the interpersonal onus on Betty, who would claim that Alan's innovative activity is prohibited: *she* must provide the justification, showing that Alan's new act-type falls under a previous prohibition. The Minimal Principle of Natural Liberty is important but, of course, it is insufficient to provide that atmosphere of freedom that Mill so valued.

B. Residual Prohibition without Conservatism?

It may seem that a critic of natural liberty such as Forst could avoid the conservatism implicit in a residual prohibition system by appealing to a version of what we deemed the General Prohibition Principle (§II, *B*): while there is a general prohibition on acting without justification, if one can justify one may proceed. So:

> *The Proceed with Justification Principle*: If one is engaging in some new action type γ for which there is no current permission in the system of moral rules, one may γ if and only if one can justify γ-ing to others.

For Forst, this may be a Kantian public reason principle, in which one justifies γ-ing

to all free and equal persons; to a utilitarian it could be a principle that allows new activity to be justified by appeal to the general welfare. In environmental thinking, this could be a version of the "precautionary principle," which allows new action when a clear case can be made that its benefits exceed the costs. All such views might insist that, by including this general principle within the system of rules, a residual prohibition system can be dynamic, allowing those new act-types that are justifiable, useful, and so on.

The Proceed with Justification Principle clearly lacks a sort of epistemic decisiveness; it is often extraordinarily unclear how these calculations are to be made. While, perhaps, in theory Proceed with Justification might have a definitive answer, it is most unlikely that people will be able to coordinate on it. At least from the perspective of social morality, and its aim to coordinate normative and empirical expectations, inclusion of such a principle scores badly on an important desideratum. And in some cases, given our current information there may be no determinate answer we can reach, and so an even deeper idea of determinacy is violated.

Still, it might be thought that the costs in determinacy (and perhaps teachability) entailed by accepting The Proceed with Justification Principle are exceeded by the benefits. The benefits, however, are meager — even with the addition of this principle, moralities of residual prohibition remain deeply conservative. Innovators can seldom justify their innovative activity, not only because they have little idea of

its consequences, but often very little idea of just what they are doing. A classic example is Alexander Fleming's discovery of penicillin; when he noticed a "blob of mold" when cleaning out his petri dishes, he did not know that he was about to discover penicillin.[47] As Fleming remarked "One sometimes finds what one is not looking for." Fleming may well have been unable to justify his experiments that led to discovery of penicillin, because he did not know that is what he was doing, and no one could have known. And that is why Benn focuses on a senseless activity such as pebble-splitting: recall that Betty's demand that Alan justify his pebble splitting requires him "to explain it, to redescribe it as an activity with an intelligible point." But it is precisely what the innovator may well be unable to do. There is no logic of discovery; some innovators seek a result, and achieve something in the neighborhood, others find something entirely different, and others are not quite sure what they are doing or why they are doing it. If, before proceeding with his innovative activity the innovators must justify it, very often the justification will not be forthcoming, as he has no clear idea of what it is that he is trying to justify. And, if so, the Residual Prohibition Principle will once again come into play, with its conservative implications.

V. CONCLUSION: TEACHING MORALITY FOR AN OPEN SOCIETY

A system of moral rules that is conducive to innovation and experimentation, we have argued, is one that is grounded on the Minimal Principle of Natural Liberty

and its associated closure rule. Such a system is far from Hardin's Bodo ethics. In addition, our findings point to the conclusion that the teaching of such a morality should focus on prohibition rules. Those who primarily learn prohibition rules, we have stressed, tend to treat their morality as being one of natural liberty rather than residual prohibition. This, then, provides grounds for claiming that the morality of an open society — one that is friendly to innovation and discovery — will be one focused on prohibitions. The point, of course, applies with equal force to the law. As a recent editorial of *The Economist* pointed out in relation to the liberalization of Cuba, "In place of a 'positive list' of permitted private activities, the government should publish a negative one that reserves just a few for the state. All others would then be open to private initiative, including professions such as architecture, medicine, education and the law."[48]

Natural liberty, then, seems fundamental to the moral rules of an open society, in which not only the outcomes, but the point of our action is often unknown. Consequently, in an open society we will very often be unable to justify a claim to liberty. As Hayek observed, "Freedom granted only when it known beforehand that its effects will be beneficial is not freedom. If we knew how freedom would be used, the need for it would largely disappear."[49]

Acknowledgements

---

[1] Joel Feinberg, *Harm to Others* (New York: Oxford University Press, 1984), 9.

[2] Rawls, *Justice as Fairness: A Restatement*, Erin Kelly, ed. (Cambridge, MA: Harvard University Press, 2001), 44.

[3] John Stuart Mill, The *Subjection of Women* in *The Collected Works of John Stuart Mill*, J.M. Robson, ed. (Toronto: University of Toronto Press, 1965), vol. XXI, 262. See also *On Liberty* in ibid., vol. XVIII, 299. In *The Principles of Political Economy* (ibid., vol. III, 938) Mill presents a formulation focused on law.

[4] John Locke, *Second Treatise of Government* in *Two Treatises of Government*, Peter Laslett, ed. Cambridge: Cambridge University Press, 1960), 287 (§4).

[5] Stanley I. Benn, *A Theory of Freedom* (Cambridge: Cambridge University Press, 1988), 87.

[6] See, for example, Kurt Baier, *The Moral Point of View: A Rational Basis of Ethics* (Ithaca: Cornell University Press, 1965), chap. 10; Baier, *The Rational and the Moral Order* (La Salle, IL: Open Court, 1995), 157; Peter Strawson, "Social Morality and Individual Ideal," *Philosophy*, vol. 36 (1961): 1–17; David Gauthier, *Morals By Agreement* (Oxford: Clarendon Press, 1986), 269; Gerald Gaus, *The Order of Public Reason* (Cambridge: Cambridge University Press, 2011), chap. I; David Gauthier, *Morals by Agreement* (Oxford: Oxford University Press, 1986), 269; Peter Strawson, "Social Morality and Individual Ideal," *Philosophy*, vol. 36 (1961): 1–17; Gaus, *The Order of Public Reason* (Cambridge: Cambridge University Press, 2011).

[7] See Shaun Nichols, *Sentimental Rules* (Oxford: Oxford University Press, 2004). For other ways to distinguish moral from other social norms, see Bicchieri *et al.*, *Norms in the Wild,* chap. 3; Geoffrey Brennan, Geoffrey, Lina Eriksson, Robert E. Goodin and Nicholas Southwood, *Explaining Norms*. Oxford: Oxford University Press, 2013), Part I; Gaus, *The Order of Public Reason,* chap. IV.

[8] Baier, *The Moral Point of View*, 195-96.

[9] On the relation of challenge to justification, see the enlightening analysis of Carl Wellman, *Challenge and Response: Justification in Ethics* (Carbondale and Edwardsville: Southern Illinois University Press, 1971).

[10] See Cristina Bicchieri, *The Grammar of Society: the Nature and Dynamics of Social Norms* (Cambridge University Press, 2006), 20-1; Shaun Nichols, "Emotions, Norms, and the

Genealogy of Fairness," *Politics, Philosophy & Economics,* vol. 9 (August 2010): 275–96.

[11] Ranier Forst, "Political Liberty" in *Autonomy and the Challenges to Liberalism,* John Christman and Joel Anderson, eds. (Cambridge: Cambridge University Press, 2005): 226-42 at 230. Reference deleted; emphasis in original.

[12] Forst, "Political Liberty," p. 242. Emphasis in original. Stephen Macedo appears to have similar qualms, maintaining that "[a]n account of basic liberties is itself a product of the justificatory enterprise…. No general presumption of liberty as non-interference forms a prior baseline."("Why Public Reason? Citizens' Reasons and the Constitution of the Public Sphere," 13. Available at: http://ssrn.com/abstract=1664085). Macedo's reference to "basic" liberties makes it more difficult to interpret his complaint. A system of morality that accepts a principle of natural liberty (including the Full-fledged Liberal Principle) can insist that important, *basic,* liberties such as freedom of speech are specially justified, and that is why they are morally protected (a view that, indeed, Rawls seems to take). In relation to, say freedom of speech, one not only possesses a blameless liberty to speak in public, but a *claim right* to do so. Macedo's deep worry, like Forst's, seems to be that the Minimal Principle of Natural Liberty attributes to Alan a blameless liberty to ϕ that is itself exempt from the need to be morally justified, and that unacceptably biases the moral system toward liberty rather than, say, equality or claims of justice, which do face the burdens of justification. Jonathan Quong has similar worries. See his "Three Disputes about Public Reason: Commentary on Gaus and Vallier," www. publicreason.net/wp-content/PPPS/Fall2008/JQuong1.pdf.

[13] See Wesley N. Hohfeld, "Fundamental Legal Conceptions as Applied in Judicial Reasoning," (1917). Yale University, Faculty Scholarship Series, Paper 4378.

[14] For his overall Kantian-inspired theory see his *The Right to Justification: Elements of a Constructivist Theory of Justice*, translated by Jeffrey Flynn (New York Columbia University Press, 2011).

[15] Claude Hillinger and Victoria Lapham, "The Impossibility of a Paretian Liberal: Comment by Two Who Are Unreconstructed," *Journal of Political Economy*, vol. 79 (Nov. - Dec., 1971): 1403-1405 at 1403-4. See also C. K. Rowley and A.T. Peacock, *Welfare Economics: A Liberal Restatement* (London: Martin Robertson, 1975), 82-3. Both these criticism are based on a type of utilitarianism. Hillinger and Lapham, as well as Rowley and Peacock, are responding to Amartya Sen's Principle of Minimal Liberalism. Sen objects to these responses on the grounds that "Everyone's right to do anything whatsoever is made conditional on non-opposition by others…." "Liberty, Unanimity and Rights," *Economica,* New Series, vol. 43, (Aug., 1976): 217-245 at 227. In our case, we might say "No one can do anything whatsoever unless they first have moral approval."

[16] There are special complications here concerning purely consequentialist theories of blame. The core of the distinction between objective and subjective rightness/wrongness is that between actions that objectively do maximize utility and those which, given the person's beliefs and information, could be rationally expected to promote utility. See J. C. C. Smart, "An Outline of a System of Utilitarian Ethics" in J. J.C. Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973),

3-74 at 30ff.

[17] John Rawls, "Justice as Fairness" in *John Rawls: Collected Papers*, edited by Samuel Freeman (Cambridge, MA: Harvard University Press, 1999 [1958]): 47-72.

[18] Baier, *The Moral Point of View*, p. 195.

[19] See John Rawls, "The Independence of Moral Theory," in *John Rawls: Collected Papers*, edited by Samuel Freeman (Cambridge, MA: Harvard University Press, 1999), 286-302 at 286.

[20] See Kenneth O. May, "A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decision," *Econometrica*, vol. 20 (October, 1952): 680-684 at 681.

[21] For a discussion see Amartya Sen, *On Ethics and Economics* (Oxford: Blackwell, 1986): 66ff.

[22] See Cristina Bicchieri and Alex Chavez. "Norm Manipulation, Norm Evasion: Experimental Evidence," *Economics and Philosophy*, vol. 29 (Special Issue 2, July 2013): 175-98; Bicchieri and and Hugo Mercier, "Self-serving Biases and Public Justifications in Trust Games," *Synthese,* vol. 190 (2013): 909–922.

[23] John Rawls, *A Theory of Justice,* rev. edn (Cambridge, MA: Harvard University Press, 1971), 115-16. In specifying the desiderata for a social morality, Rawls accepts a version of Baier's teachability requirement (113n).

[24] This supposes that it makes sense to talk of the function of morality. For a defense of this claim, See Gaus, "On Dissing Public Reason, *Ethics*, vol. 125 (June 2015).

[25] See Gaus, **The Tyranny of the Ideal: Justice in a Diverse Society** (Princeton: Princeton

University Press, 2016) cap. I.

[26] It is not generally the case that a rational person will be paralyzed by under-complete preferences, as Sen's analysis of Buridan's ass illustrates: it is often rational to select from a maximal set of undominated options. Even if the ass cannot rank the relative attractions of eating the hay to the right or left, it can know that either is better than starving so it ought to select one. But that is not the case when one seeks to conform to a deontic system: here one really would be unable to choose what act is consistent the rules. This case also should be distinguished from a social choice problem where we are trying to choose what set of rules to follow among a maximal set of alternative systems of rules. See Sen "Maximization and the Act of Choice," in his *Rationality and Freedom* (Cambridge, MA: Harvard University Press, 2002): 158-205 at 181ff; Gaus, *The Order of Public Reason*, 303ff.

[27] Rawls, *A Theory of Justice*, rev. edn (Cambridge, MA: Harvard University Press, 1009), pp. 300-301.

[28] As Gilbert Harman argues, the ideal of a belief system that is deductively closed, or complete under logical implication, is neither required by rationality nor realistic for humans., *Reasoning, Meaning and Mind* (Oxford: Clarendon, 1999) 21-3.

[29] John Mikhail, *The Elements of Moral Cognition: Rawls' Linguistic Analogy and the Cognitive Science of Moral and Legal Judgment* (Cambridge: Cambridge University Press, 2010), §6.3.1. We are thus translating a requirement to ϕ as a prohibition on failing to ϕ.

[30] See, for example, Patrick Shafto, Noah Goodman, and Tom Griffiths, "A Rational

Account of Pedagogical Reasoning: Teaching By, and Learning from,

Example," *Cognitive Psychology*, vol. 71 (2014): 55-89.

[31] Ibid., 65.

[32] Responses in the permission-training condition did not differ significantly from

chance, but responses in the prohibition-training condition did (one sample t-test

$t(59)=3.65$, $p<.001$). All subjects were recruited through Amazon's mechanical turk.

There were 114 subjects in the first study, 45 female.

[33] Responses in the permission-training condition did not differ significantly from

chance, but responses in the prohibition-training condition did (one sample t-test

$t(28)=2.53$, $p=.017$). There were 49 subjects, 26 female.

[34] This experiment draws on Denise Dellarosa Cummins, "Evidence for the Innateness of

Deontic Reasoning," *Mind & Language*, vol. 11 (June 1996): 160–90; "Evidence of Deontic

Reasoning in 3- and 4-year-olds." *Memory and Cognition*, vol. 24 (1996): 823–829.

[35]There were 188 subjects, 82 female.

[36] On the tendency of normative life to reconceive continuous categories into

dichotomous ones that induce the agent to make a choice, see S.I. Benn and G.F. Gaus,

"The Liberal Conception of the Public and Private" in their edited collection *Public and*

Private in Social Life (New York: St. Martin's Press, 1983), 31-65.

[37] Permission vs. prohibition: $t(53)=10.2$, $p<.0001$; permission vs. mixed: $t(67)=6.98$,

$p<.0001$; prohibition vs. mixed: $t(68)=2.05$, $p<.05$.   There were 93 subjects, 39 female.

[38] At least some of the participants explained their answers in ways that conform to

pedagogical sampling. Here are a few examples:

"It would seem that if squeaky mice aren't allowed in the Green Barn, it would be explicitly noted, though there is no way to know for certain."

"I would assume that if the Green barn was off limits, there would be a rule specifying that."

"It never said specifically that they were not allowed in the green barn, so until expressed it is allowed."

[39] Mill, *On Liberty*, p. 261. See also Ryan Muldoon, "Expanding the Justificatory Framework of Mill's Experiments in Living," *Utilitas*, available on CJO2015. doi:10.1017/S095382081400034X.

[40] Russell Hardin, "From Bodo Ethics to Distributive Justice," *Ethical Theory and Moral Practice* vol. 2 (1999): 399–413, at 401-02. See also Hardin, "The Priority of Social Order," *Rationality and Society,* vol. 25(2014) 407-421 at 411ff.

[41] Russell Hardin, *Indeterminacy and Society* (Princeton: Princeton University Press, 2003), 98.

[42] Hardin, *Indeterminacy and Society*, 98. In "The Priority of Social Order" Hardin is more hospitable to norm-based ordering.

[43] Hardin, *Indeterminacy and Society*, 98.

[44] Federal Communications Commission, *Report and Order on Remand, Declaratory Ruling, and Order* http://transition.fcc.gov/Daily_Releases/Daily_Business/2015/db0312/FCC-15-

24A1.pdf

[45] See Benn, *A Theory of Freedom*, 289-91.

[46]Mill, *On Liberty*, vol. XVII, 267.

[47] See *The Discovery and Development of Penicillin, 1928-1945* (Royal Society of Chemistry and American Chemical Society, 1999).

[48] "Reforming Cuba," *The Economist*, May 16, 2015.

[49] F. A. Hayek, *The Constitution of Liberty* (London: Routledge, 1960), 31.