

Sustaining rules: a model and application *

John Turri
john.turri@gmail.com

Abstract: I introduce an account of when a rule normatively sustains a practice. My basic proposal is that a rule normatively sustains a practice when the value achieved by following the rule explains why agents continue following that rule, thus establishing and sustaining a pattern of activity. I apply this model to practices of belief management and identify a substantive normative connection between knowledge and belief. More specifically, I propose one special way that knowledge might set the normative standard for belief: knowing is essentially the unique way of normatively sustaining cognition and, thereby, inquiry. In this respect, my proposal can be seen as one way of elaborating a “knowledge-first” normative theory.

Keywords: rules; practices; normativity; belief; knowledge

1. Introduction

We value many things. Some are easily obtained, others aren't. We also disvalue many things. Some are easily avoided, others aren't. To obtain or avoid any of these things, we need ways of doing so, ways that are good enough to be worth the effort. Worthy ways per-

* This is the penultimate version of a paper forthcoming in *Knowledge first* (Oxford University Press), ed. J. A. Carter, E. C. Gordon, & B. Jarvis. Please cite the final, published version if possible.

sist and unworthy ways perish. Worthy ways of the right sort proliferate, thereby instituting and sustaining practices.


These observations can seem uncontroversial, even bordering on platitudinous. But they provide most of the materials needed to develop a robust theory of the rules that normatively sustain a practice. In short, they point toward a theory of what I call *sustaining rules*.

Before proceeding I should clarify what I mean by ‘rule’ and ‘practice’. A *rule* is a principle or standard of conduct. Conduct can be assessed according to a rule. If you conform to the rule, then your conduct is correct according to the rule. If you break the rule, then your conduct is incorrect according to the rule. A *practice* is an ongoing way of doing things. Practices are instituted, can persist, change, and expire. I’m thinking of practices as concrete patterns of actions, held together as patterns by the rules that their agents follow. The rules “hold the practice together,” by analogy to the way that water binds together isolated bits of powdered milk into liquid milk. A practice could also be thought of as an abstract type, a description of a series of actions done according to a rule. For convenience, I’ll continue speaking of practices as concrete patterns rather than abstract types, and of rules sustaining practices thus construed. I’ll have more to say about how I’m thinking of practices as the discussion unfolds.

2. *The model: a simple and intuitive presentation*¹

There are six basic elements of the model: agent, goal, rule, achievement, persistent pattern and explanation.² A rule sustains a practice when the agent follows the rule and achieves what it prescribes, the value of which in turn explains why the agent continues following the rule, thereby establishing a persistent pattern. In a normatively sustained practice, the value achieved through earlier instances of rule-following breeds subsequent rule-following. In a word, there is reproduction via value produced.

At the risk of oversimplification, it might be helpful to visualize the model like so:



The arrows represent explanatory relations. The emergence of a persistent pattern in this way *just is* the emergence of a practice: it *makes* the practice. Generation is the initial, limiting case of sustenance.

Strictly speaking, it is the activity of agents that generates and

¹ The view is partly inspired by Millikan's (2005) 'biological model' of public linguistic meanings as 'stabilizing functions.' It is also broadly related to the theory of 'domain' normativity and performance assessment discussed in Sosa 2007: chs. 2 and 4, and the project of the primitive enquirer ('A') discussed in Williams 1978: pp. 23ff. There are also superficial similarities to Alston's (1991) 'doxastic practice approach' to epistemology, though deep disagreements divide me and Alston. My thinking on this subject was influenced by reading in the summer of 2010 a paper by Benjamin Jarvis, entitled 'Assertion and Advocacy' (unpublished ms).

² Several of these elements could be plural. More on this below.

sustains practices, what we might call *sustaining agency*. Sustaining agency and sustaining rules are but two sides of the same coin.³ I aim to illuminate both sides, though it has proven easier for me to conduct the discussion directly in terms of rules and allow the points about agency enter at the appropriate junctures.

Our powers and abilities equip us to exercise some measure of control over outcomes, such that successful outcomes can *manifest* our powers and abilities.⁴ (For convenience, I'll often just say either 'powers' or 'abilities' instead of mentioning both.) An outcome is our *achievement* if and only if it manifests our abilities. Not every successful outcome is an achievement. For example, some are due to luck.⁵ A *feasible outcome* is one that we can achieve as we are actually constituted, which requires that we actually have some ability to produce it. An *infeasible outcome* is one that we have no ability to produce as we are actually constituted. (An infeasible outcome might still be possible, even if we actually have no ability to produce it.)

It is trivially true that sustaining rules prescribe a feasible outcome. For if we couldn't achieve the prescribed outcome, then subsequent rule-following couldn't be explained by our previously achieving the outcome; and if subsequent rule-following couldn't be explained by our previously achieving the outcome, then we couldn't sustain the practice by following the rule. Of course, the

³ MacIntyre 1999 has influenced my thinking here.

⁴ For a partial theory of this relation, see Turri in press.

⁵ For more on this way of thinking about powers and outcomes, see Sosa 2007, Sosa 2011, Greco 2010, Turri 2011a, Turri 2011b.

rule-following behavior might somehow persist even if we achieve nothing valuable through it. We might continue on in the same hopeless way. And our persistence in such hopeless ways might even incidentally cause something valuable to happen to us, perhaps through blind luck or the intervention of a benevolent and merciful higher power. But this wouldn't amount to us sustaining the practice in the sense I'm interested in.

My view about sustaining rules and agency does not commit me to any particular view about the rules that constitute *the actions that occur within a practice*. It does not even commit me to the view that the actions that occur within a practice are constituted by rules.⁶ From the fact that a practice is sustained by a rule, it does not follow that actions occurring within that practice are made possible by the rule, or indeed by any rule. For example, the practice of horse racing involves mounting and riding a horse. But intuitively the actions *mounting* and *riding* aren't constituted or made pos-

⁶ Searle (1969: ch. 2.5) and Williamson (2000: ch. 11) both take a different approach to constitutive rules; Rawls (1955: 25) presents a related theory, what he calls "the practice conception" of rules, according to which rules are "logically prior" to certain actions. One main reason that I don't offer my view as a theory of "constitutive rules" is that the phrase 'constitutive rules' carries considerable baggage, and I wish to avoid verbal squabbles. More importantly, sustaining rules are not "general" or "conventional" rules (Williamson 2000: 238–9), "regulative rules" (Searle 1969: 33), or "rules of thumb" either (Rawls 1955: 23). Sustaining rules, as I define them, differ from all these widely discussed categories. In particular, I am not offering a theory of rules that necessarily "create the possibility of new forms of behavior" (Searle 1963: 33) or are "logically prior" to any actions that occur within the practice. My present purpose is to identify and explain a theory of sustaining rules and aside from ensuring that what I'm doing is not confused with what Rawls, Searle, Williamson and others have tried to do, I lack the space to compare the approaches.

sible by the rules of horse racing. Rather, they are antecedently given actions that the practice of horse racing incorporates into more complex patterns of activity.

3. *Refinement and elaboration*

Having presented the basic model, in this section I include some finer points and qualifications to further clarify the view. I include these refinements and qualifications at this point in order to forestall common questions and misconceptions. Those uninterested in such refinements can skip to section 4.

Not all ways of sustaining a practice are ways of *normatively* sustaining it. My theory applies to both ways of sustaining a practice, but I'm really interested in the latter. Normatively sustained practices have *genuine reason-giving force*, at least for their practitioners: practitioners have good reason to achieve what the rules prescribe. But what distinguishes normatively sustained practices in this way? Every practice has a point, which can be identified with that goal, achievement of which explains why practitioners persist in following the rule. (I don't intend to mark an important distinction between a point and a goal. They are both just objectives or ends. I often use 'point' for a practice and 'goal' for agents, to help keep straight the different elements of the model.) But not all goals are good. A normatively sustained practice has a good point, which explains its reason-giving force.

The practice's point needn't be intrinsically or necessarily good

in order to imbue it with reason-giving force. It is enough, I think, that the point is usually good, all things considered, in some way or another under normal circumstances. In a word, the point need only be ordinarily good. Consider *preventing tooth decay*, which is the point of the practice of brushing your teeth daily. Preventing tooth decay is instrumentally good in normal circumstances because in those circumstances it tends to promote better health, which tends to promote greater happiness. This explains why we have a reason to follow the rule *brush your teeth daily*. Preventing tooth decay is only contingently instrumentally good. In some strange circumstances, it would be instrumentally bad to prevent tooth decay — e.g. a situation in which all people without rotten teeth are tortured to death, or in which our teeth rot as part of a natural and healthy exfoliation process. Nevertheless, although I think that a liberal approach on this point is advisable, it is consistent with the basic model to impose stricter requirements on the point of a normatively sustained practice.

In order for a rule to sustain a practice, practitioners needn't succeed every time that they engage in the practice or attempt to follow the rule. They needn't even succeed most of the time. They need only succeed enough of the time for the practice to sustain itself through the value it brings. Depending on the point of the practice, very infrequent success could be enough. The practice of diamond hunting and other forms of prospecting are like this. And as examples of such practices show, the involvement of later parties in a practice might be explained by the benefits achieved through the

rule-following of previous parties long ago. For example, a family might engage in the practice of diamond prospecting over many generations and succeed only once in a century. Nevertheless, the benefits earned are enough to sustain the pattern of rule-following of family members in between successes.

Not all rule-following sustains a practice in the sense I'm interested in. In particular, repeated but isolated rule-following that fails to explain further instances of similar behavior doesn't sustain a practice. For example, suppose that we are consistently convinced by a different person with a different reason each day to hang our laundry out to dry, and the benefit of doing this previously doesn't explain why we do it subsequently. Even though we frequently follow the rule *hang the laundry out to dry*, and we benefit from doing so, this is not a normatively sustained laundering practice or even, in the sense I'm interested in, a practice at all. There is no "reproduction value" here. Someone observing our daily actions might reasonably infer that we had instituted a laundering practice, but that is because in this case outward appearances are misleading. By contrast, if we had as our goal *saving energy*, and we followed the rule *hang the laundry out to dry* and thereby achieved energy savings, and this benefit explained why we continued following the rule and reaping the benefit, then it would be a normatively sustained laundering practice. Here we have reproduction through the value produced.

Closely related to this last point, continued rule-following behavior must have an appropriate etiology in order for it to constitute a

practice. Even if a set of activities follows the exact same rule in exactly the same way, it doesn't follow that they are instances of a single practice. Practices are like biological species in this respect. Specimens must share a proper evolutionary etiology to be conspecifics. Even if a group of specimens share all their gross morphological features, it doesn't follow that they form a species. They might still be a cryptic species. Moreover, practices that prescribe outwardly similar behavior in a wide range of circumstances might still be very different due to differences in their rules. Again, the comparison to biological species is helpful: two species might be very different genetically (follow different "genetic rules") despite being outwardly similar due to convergent evolution.

In light of the points made in the last two paragraphs, I find it very plausible that even though practices are uncontroversially *artificial* in an important sense of that term, they nevertheless have an underlying nature that is partly determined by their history. In this respect they are importantly similar to natural kinds.⁷

Some practices are intentionally designed, but they needn't be. Many important practices are not intentionally designed. For example, the practice of assertion was not intentionally designed. Neither were the chimpanzee practices of smashing open nuts with rocks or using straws to extract termites from crevices. Again, we might compare practices to species, which can be intentionally designed through husbandry or genetic engineering, even though

⁷ This is in keeping with the most influential contemporary scientific work on human norms, culture and evolution (e.g. Richerson and Boyd 2005, Tomasello 2008, 2009).

most have evolved without intentional design through natural selection.

Goals can be complex and difficult to identify precisely. Often we don't just want some particular kind of thing, but rather that kind of thing without any admixture of another particular type of disvaluable thing, both in this specific case and also in the long run. The importance of avoiding disvaluable admixture will partly determine the shape of any sustainable practice organized around the goal.

Powers and abilities needn't be infallible. Indeed, virtually none is infallible. A power or ability enables you to achieve a certain result in a certain environment. It might enable you to achieve the result on fifty percent of attempts in a normal environment, on ten percent of attempts in an especially hostile environment, and on ninety percent of attempts in an especially hospitable environment. From the fact that you sometimes fail to, say, flip the switch when you try, it doesn't follow that you lack the ability to flip the switch. Neither does it follow that you have power only over whether you *try* to flip the switch. You still have the power to flip the switch. It's just not an infallible power.

Getting lucky is a way of getting what we want and it is, trivially, good to have good luck. But, equally trivially, getting lucky is not a way of achieving success, and so not a way of sustaining a practice. Achievement is the antithesis of lucky success. Achievement is success manifesting ability, and the extent to which success manifests our ability is inversely proportional to how lucky we get.

Consistent with this, there could be cases where some combination of luck and ability conspire to bring success. Indeed, most actual cases of success are probably like this. Success is often a matter of degree, as can be the extent to which it manifests our abilities. I will abstract away from these complications and speak as if these were dichotomous matters.

Not all rule-following is conscious or explicitly intentional. In fact, I think that most rule-following is instead habitual and automatic, in the sort of way that North American drivers drive on the right side of the road. North Americans habitually, mechanically follow the rule *drive on the right side*. They don't consciously reflect on it or wake up in the morning and form an explicit intention to drive on the right side today.

Closely connected to the previous point, it's worth noting that we typically *practice* ways of doing things so that they become "second nature" — that is, so that we conduct ourselves that way "without thinking," and especially without consciously deliberating or explicitly formulating and executing an intention. 'Practice' in this sense, a synonym of 'rehearse', is clearly closely related to the sense of 'practice' relevant to my view of how rules sustain practices, a synonym of 'customary procedure' or 'routine'. It is through rehearsal that things become routine: by practice is practice established.⁸

⁸ There are potentially interesting points of convergence here with how Aristotle characterizes moral virtue in the *Nicomachean Ethics*. Writes Aristotle, "Moral virtue comes about as a result of habit" (Aristotle 350 BCE, Book 2.1, 16-17).

Practitioners must be implicitly sensitive to the rules of the practice, otherwise they couldn't follow the rule as opposed to merely conform with it. I take no positive stand on what form such sensitivity must take or what makes it true that such sensitivity is present. But it's clear that sensitivity doesn't require being able to explicitly formulate the rules. Practitioners needn't be especially good judges of what the rules are, even if we expect patterns in their behavior to offer helpful clues to what the rules are. (Actions speak louder than words.) Practitioners might not be explicitly aware of the rules that they're following, or that they're engaging in the practice that they're engaged in. They might even sincerely judge themselves to *not* be engaged in a practice that they regularly engage in. Plausibly, at least some people engaged in contemporary racist or sexist practices are like this. It is possible that observers of a practice are better positioned than the participants to discover what the rules are.

I'm not offering a theory of rule-following. I assume that people do follow rules and I take this assumption to be uncontroversial.

4. An example: Robin Hood and his merry men

This section presents, and the following three sections discuss, a toy example illustrating the model.

Robin Hood and his Merry Men:

When Robin was a young lad, he witnessed a curious event that left a lasting impression. Just outside a poor village at the edge of the forest, a bold group of criminals robbed a very wealthy noble-

man who was transporting his treasures. As they made their getaway, out of exuberance and without design, the criminals tossed some of the stolen treasure along the road in the village. The peasants retrieved it and felt very obliged. When the Sheriff arrived to investigate, none of the peasants revealed what they knew, namely, who the criminals were and which direction they fled. Instinctive gratitude led them to remain silent.

The criminals robbed from the rich and “gave” to the poor. But this wasn’t their practice. It was an isolated event and they weren’t in any way aiming at what turned out to be, from their perspective, a very fortunate outcome.

Years later when Robin was a man, he became disillusioned with the unjust distribution of wealth in the English feudal system. He desired greater distributive justice in the local economy. Robin recalled the successful robbery he witnessed as a young lad. The lesson of that robbery stuck with him and inspired his incipient plan to enhance distributive justice: rob from the rich and give to the poor. Robin recruited a group of like-minded merry men and set up in Sherwood Forest. They robbed rich people traveling through the forest and gave the windfall to the poor, retaining enough to pay their expenses and sustain their redistributive enterprise. The local population benefitted greatly from this and adored Robin and his merry men because of it. Consequently, the locals refused to cooperate with the Sheriff of Nottingham’s effort to imprison Robin and his merry men. Things continued this way for a long time.

Here we have a normatively sustained practice: the practice of Robin Hoodery. An initial application of the model to it might go as

follows. The agents are Robin and his gang. Their goal is greater distributive justice, which is a good thing. The relevant abilities include their thieving skills.⁹ The rule is, roughly, *rob from the rich and give to the poor*. They follow this rule and achieve the goal, which in turn explains why they persist in this pattern of behavior, thus instituting and sustaining the practice.

If Robin Hoodery is to be sustained, then its rules — call them *Robin's Rules* — will be more complicated than the initial formulation just suggested. The merry men must also not undermine their support in the community by, for example, also robbing from poor people. If they rob everybody they encounter, rich and poor alike, then their support in the community will be undermined, which will very shortly lead to their demise. Robin's Rules will also instruct them to avoid this disvaluable admixture of results. At the same time, it's implausible that the merry men would be engaging in the practice incorrectly by not constantly robbing, and presumably the rules should reflect this. Here, then, is one plausible suggestion for a set of rules for Robin Hoodery:

Robin's Rules

Robin's Imperfect Duty: Not infrequently, rob a traveler.

Robin's Rule of Permission: You may rob a traveler only if

⁹ I'm eliding some real but minor complications for the sake of clarity and simplicity. The peasants also play a role in the practice, even if their participation is mainly passive, and their behavior is prescribed by rules of the practice pertaining to gratitude and complicity. The number of abilities is actually surprisingly many, as is fitting for a complicated activity like organized thievery.

the traveler is rich.¹⁰

Robin's Perfect Duty: You must distribute the plunder to the local poor, less the necessary and appropriate overhead costs of the redistributive activity.

Prevailing conditions of human existence demand that a successful plan have some flexibility. Lack of precision is the cost of flexibility. Robin's Rules aren't perfectly precise. What counts as *not infrequently* will depend on how much redistribution is needed to keep the local population on Robin's side. What counts as *necessary and appropriate overhead* will likewise vary and be determined by such things as the costs associated with robbing travelers (e.g. the right kinds of weaponry, disguises, and other equipment) and maintaining the merry men's living quarters and health. Attitudes of the local poor will also influence what is appropriate. For example, if Robin and his merry men start retaining enough to live a lavish lifestyle, then it might alienate their allies and lead to their demise. Plausibly there is no unique right way to set values for these variables in the plan. Vigilance will be required to keep the values within appropriate ranges. Appropriate attention to environmental factors will shape how Robin and his merry men adjust the values

¹⁰ Alternatively: Avoid robbing non-rich travelers. Avoidance could be understood as a matter of degree, in which case the rule needn't forbid the robbing of *any* non-rich traveler; instead, it might require only that non-rich travelers not exceed a certain proportion of overall travelers robbed. As will be noted several times in subsequent footnotes, I am sympathetic to rules requiring only that participants' activity be marked by certain *central tendencies*. I include these points in footnotes because it's important to flag that the model is flexible in this way. At the same time, I relegate these points to footnotes because I want to keep the proposals in the main text uncluttered.

as their practice unfolds. Otherwise put, the relevant environmental factors *remotely condition* an effective interpretation of the rules and thereby the practice.¹¹

Robin's Rules prescribe certain outcomes. The merry men have no power to make these outcomes just happen, or happen by luck, or happen by magic. They can achieve these outcomes only by exercising their ability to produce the outcomes. There simply is no other way for them to do it. And it is by achieving these things that the merry men achieve what the rules prescribe and thereby achieve the goal of the practice, namely, greater distributive justice. Greater distributive justice is their achievement, an outcome that manifests their skill, effort and agency. And the value of their achievements breeds further activity of the same sort.

5. *Success and evidence*

I am often asked why Robin's Rules aren't framed in terms of *evidence*, such as the following alternative for a permission rule:

Robin's Evidential Rule of Permission: You may rob a traveler only if you have good evidence to believe that the traveler is rich.

Here *good evidence* does not guarantee that the traveler is rich, but it does make it probable.¹²

¹¹ This idea is inspired by Ben Jarvis's work on "distal shaping" (unpublished ms). Jarvis expressed some concern about my use of 'distal shaping' in the present context, so I switched to 'remote conditioning' to pick out the related phenomenon described in the main text.

¹² In my experience, this sort of alternative gets posed by those who favor a

This is an important suggestion providing an opportunity to explain two other features of sustaining rules, namely, what sort of rules are most eligible to be sustaining rules, and how sustaining rules give rise to derivative rules of a practice. In the remainder of this section, I discuss the former feature; in the following two sections, I discuss the latter.

Evidential rules can have an important derivative place in a practice, but they are inferior to what I call *direct success rules* as candidates for sustaining rules. But before getting to that, I'd like to first point out that although there is a fairly natural evidential substitute for Robin's Rule of Permission, there is no natural evidential substitute for the others. Consider:

Robin's Imperfect Evidential Duty: Not infrequently, do something that you have good evidence to believe is robbing a traveler.

Robin's Perfect Evidential Duty: You must do something that you have good evidence to believe is distributing the plunder to the local poor, less the necessary and appropriate overhead costs of the redistributive activity.

Given how unnatural these are, it seems unlikely that a thoroughgoing commitment to evidential substitutes will be attractive. In light of this, and in order to give the alternative evidential approach the best run for its money, I'll focus on the more natural evidential per-

non-factive conception of good evidence, so I entertain it accordingly. For work investigating how closely this tracks the ordinary way of evaluating evidence, see Turri 2015; Turri in press c.

mission rule.

Robin's Rule of Permission is a better candidate for a sustaining rule because it is a *direct success rule*. Direct success rules directly prescribe the goal. Evidential rules do not. Instead, evidential rules prescribe something that you have evidence to believe will count as satisfying the goal. The difference is critical. Consistently following Robin's Evidential Rule of Permission could directly undermine the practice. Despite following the rule, the merry men might still rob many poor people and thereby alienate their allies, which under prevailing conditions will effectively end the practice of Robin Hoodery. This would happen if the merry men's evidence was consistently enough misleading. However, consistently following a sustaining rule would not undermine a practice. Indeed, just the opposite. In line with this, consistently following Robin's Rule of Permission wouldn't undermine Robin Hoodery, but would instead help sustain it. This is because it is a direct success rule.

When I say that direct success rules directly prescribe the goal, this is an oversimplification due largely to the potential complexities of goals. Strictly speaking, it is only a practice's entire set of sustaining rules that must directly prescribe the goal. That is, the set of sustaining rules, taken as a whole, will have enough content to prescribe the goal. Robin's Rules form such a set for Robin Hoodery. Individual rules within the set might merely prescribe things that are, at least given prevailing conditions, partly constitutive of or instrumental to the goal of the practice. For example, Robin's Rule of Permission by itself isn't enough to specify the goal of Robin

Hoodery , and following it isn't enough to sustain the practice.

6. *Heuristics and evidence*

What role do evidential rules play in a practice, if not as sustaining rules? They function derivatively as *heuristics* and as *standards of criticism*. I discuss heuristics in this section and standards of criticism in the next section.

Heuristics are important and useful. A heuristic is intended as a convenient substitute for its correlative direct success rule. If the direct success rule prescribes *S*, then a correlative heuristic would prescribe *H*, where achieving *H* (1) is at least typically easier than achieving *S*, and (2) correlates well enough with achieving *S*. (For convenience I'll call a heuristic's correlative direct success rule its *superior rule*.) The strength of the correlation determines the quality of the heuristic. Following a good heuristic is a good way of approximating the results of following the superior rule. If the heuristic is good enough, then following it might even be a way of following its superior rule. In light of this, we might call heuristics *indirect success rules* for a practice.

Nevertheless, their usefulness and importance can't conceal the fact that even good heuristics are parasites. A heuristic's place and importance in a practice derives entirely from that of its superior rule. If a once good and popular heuristic becomes bad due to some change in the environment, then we would expect its popularity to diminish over time among practitioners committed to the superior

rule. And if a better heuristic was discovered, over time committed practitioners would abandon the old heuristic and adopt the new one. The change of preference for one heuristic over the other is explained by the primary importance of the underlying superior rule. And it is the enduring commitment to the underlying superior rule that explains why we have continuity of practice across the more superficial change of heuristics. This is similar to the way that, in corporate or military planning, an improvement in tactics is understood by reference to continuity in the underlying strategy that the tactics are intended to implement.

Notice something important but easily overlooked. Heuristics are also direct success rules in their own right. This might sound paradoxical initially, but a moment's reflection reveals that it is trivially true. A heuristic tells you directly to do something too. If you don't do it, then you break the rule and your conduct is incorrect according to the heuristic.

Even if it's easier to follow a heuristic than its superior rule, it still might not be easy to follow the heuristic. There can also be heuristics for heuristics, so there might be a temptation to keep finding heuristics that are easier and easier to follow, until we reach rules that we can infallibly follow whenever we choose. However, the relation *x is a good heuristic for y* isn't transitive, because the relation *x correlates well with y* isn't transitive. From the fact that H^* correlates well with H , and H correlates well with S , it doesn't follow that H^* correlates well with S .

Returning now to Robin's Evidential Rule of Permission, I think

that it functions as a heuristic in the practice of Robin Hoodery. (A better name for it would be ‘Robin’s Permission Heuristic’.) Robin and his men might use evidence of wealth as a proxy for wealth, but the evidential rules are still derivative and parasitic. This fits the pattern of a heuristic in subordinate position to its correlative direct success rule.

To illustrate the point, suppose that the prevailing fashion around Sherwood is as follows: all and only rich people wear purple robes, whereas all and only poor people wear green rags. Further suppose that this is common knowledge. Since the color of someone’s robe is typically much easier to detect than their net worth, this provides Robin and his men with a convenient heuristic for selecting eligible travelers to rob: *rob a traveler only if the traveler wears a purple robe*.

Now suppose that fashions change without Robin and his merry men hearing about it. The local rich people suspect that they’re being marked for robbery because of their purple robes. So they decide to give away their vast quantities of purple robes to the local poor, who are more than happy to wear them; the rich people then begin wearing green rags instead. Now when Robin and his merry men rob purple-robed travelers, they are robbing poor people — and, perhaps, then giving to the rich in green rags! They are breaking Robin’s Rule of Permission, but they are following Robin’s Permission Heuristic. Will this conduct sustain Robin Hoodery?

No, it obviously won’t. Their actions destabilize the practice and even threaten to undermine it entirely. If Robin Hoodery is to per-

sist, then Robin and his merry men will have to alter their behavior, bringing it back in line with the Rule of Permission. In the process, if they are wise, they will also adjust their views about what constitutes evidence of wealth. They will interpret the Permission Heuristic in light of what counts as correct observance of the Rule of Permission. Moreover, if given a choice between having their behavior conform to the Rule of Permission or the Permission Heuristic, committed practitioners will choose the former. In light of all this, I submit that the Rule of Permission helps to sustain Robin Hoodery, whereas the Permission Heuristic is just that, a heuristic.

7. *Criticism and evidence*

Now let's turn to standards of criticism. Here I mean 'criticism' in the negative sense of 'expressing disapproval' or 'fault finding', rather than the more neutral sense of 'evaluation'. Standards of criticism, in the relevant sense, are actually standards of *punishment*, mild though it may be. Telling someone that they did something bad or wrong is, in effect, to impose a mild verbal penalty for their behavior.

Many practices involve cooperation among people in different roles. When people's fates are intertwined, it makes sense to have ways of influencing others to avoid incorrect behavior. Criticism is one way of influencing others. Thus it makes sense that social practices will involve standards of criticism, enabling participants to effectively influence associates by negatively reinforcing correct beha-

violation.

From the fact that someone breaks a rule, it doesn't follow that criticizing them or their action is appropriate. If criticism won't improve behavior — or worse, if it will lead to more rule-breaking, or cause some other serious damage — then it would be dysfunctional for a practice to condone such criticism. A well functioning practice will tend to condone criticism only if, and to the extent that, it is instrumental in promoting the goals of the practice.

What will promote the goals of the practice is typically a contingent matter. Humans tend to resent criticism of incorrect behavior that the agent had good reason to think was correct. Thus we would expect successful human practices to reflect this fact by incorporating standards of criticism that place a premium on the agent's evidence at the time of action. We typically excuse people from criticism when they had good reason to believe that their action was correct, even if it wasn't.¹³

Standards of criminal liability often display just this pattern. For example, common law typically requires an objective violation of a rule (*actus reus*) and an understanding that the conduct would violate a rule (*mens rea*) in order for someone to be legally guilty and thus subject to punishment. Proving that you had good evidence that your conduct was legal is enough to avoid guilt, because it is enough to invalidate a claim of *mens rea*. There are many exceptions to this and a host of other technical legal issues threaten to

¹³ For experimental verification of this empirical claim, see Turri 2013; Turri and Blouw 2014.

crowd in almost immediately. But for my purposes it's enough to note this central tendency in the law as a way of illustrating the phenomenon I'm describing. Breaking the law is considered necessary but (usually) insufficient to warrant a penalty.

Of course, following some rules, such as those forbidding the killing of innocents, might be so important to us that we're willing impose punishment for breaking them at all (for further discussion, see the papers collected in Simester 2005). We hold people strictly accountable for breaking such rules and do not excuse them, regardless of their evidence. But strict accountability is the exception in human practice, not the rule.

Standards of criticism clearly play an important role in human practices. But standards of criticism aren't suitable to hold a practice together, so they're not good candidates for sustaining rules. Escaping criticism can be a good thing but it's not the sort of thing that explains why a practice persists. As an analogy, consider a family consisting of a mother, father and children. The parents have a duty to provide food for the children: the First Family Rule is *parents must give their children daily bread*. If the parents consistently provide bread to the children ('sustenance' in one sense), then they're following the rules and this contributes to the family's continued existence ('sustenance' in another sense). By contrast, consider the evidential substitute, the First Family Evidential Rule: *parents must do something that they have good reason to believe constitutes giving their children daily bread*. If the parents consistently escape criticism for not providing bread to the children — be-

cause, say, most things that appeared to be bread were just a bunch of nutritionless lookalikes made from ash, water and artificial flavoring — then their conduct won't contribute to the family's continued existence, even though they're following the evidential rule. Indeed, before long their conduct will ensure the family's demise. Only objective success in providing sustenance can sustain the family.

Returning now to the evidential substitutes for Robin's Rules, I propose that in addition to functioning as heuristics, they also help set standards of criticism. Following Robin's Permissibility Heuristic (i.e. the Evidential Rule of Permission) exempts the merry men from criticism for their actions, even though it doesn't set the standard for permissible conduct.

Recall the imagined situation in which it was common knowledge that all and only rich people wore purple robes. Robin and his men know this, so they have good evidence that by robbing a purple-robed traveler, they are acting permissibly. But then the rich secretly decided to give away their purple robes to the poor and wear green rags instead. Now if Robin robs a purple-robed traveler, he acts impermissibly. But others would be reluctant to criticize him for this mistake (at least initially). That's because he was following the Permissibility Heuristic, which puts his action beyond appropriate criticism.

In general, where there is a human social practice with sustaining rule R, and H is commonly employed by practitioners as a good heuristic for R, then we would expect there to be a derivative stand-

ard of criticism, C, to the effect that agents are generally exempt from criticism if they follow H but thereby violate R.¹⁴

8. *Belief management*

We have ongoing ways of managing our beliefs. Supposing that belief management is a normatively sustained practice, what can the present model tell us about it?

It is a familiar idea that the point of belief management involves believing truths and not believing falsehoods.¹⁵ Either goal on its own is easy enough to accomplish, but their combination can make belief management challenging. If the only goal were to believe truths, then *believe everything* would be a good rule to follow. Likewise, if the only goal were to avoid falsehoods, then *believe nothing* would be a good rule to follow. But following either of these rules would do us no good; nor would it sustain the practice of belief management. Indeed, it would be a total disaster! We want true beliefs, but without a harmful admixture of falsehood. In the abstract, it is difficult to say precisely how to correctly balance these twin goals, but no one doubts that belief management aims at some mix of acquiring true beliefs and avoiding falsehoods.

Belief management might also aim at acquiring beliefs that serve other purposes. These might not be prototypical cases of beneficial belief management, but not only are they possible, they al-

¹⁴ My views on this are partly inspired by Williamson's (2000: 256) distinction between the impermissible and the reasonable, and DeRose's (2002: 180) related distinction between primary and secondary propriety.

¹⁵ E.g. James 1897, DePaul 2001, and David 2005.

most certainly actually happen. Athletes work themselves into the false belief that they are better than their opponent, so that they perform better; scientists focus intently on the virtues of their inductive theories and ignore the problem of the pessimistic meta-induction from the history of science; religious fanatics find ways to convince themselves that the evidence for evolution is an elaborate hoax, so that they can sleep better; free-market fundamentalists manage to continue believing that no government regulation of commerce has good consequences, so they don't have to admit to themselves the error of their ways; people who overestimate their competence tend to be more confident and better performers as a result; etc.

I accept that belief management could also aim at these outcomes and that such outcomes can benefit us and thereby help to sustain the practice of belief management. Although I remain open to the possibility that such outcomes shouldn't be counted as benefits of belief management *per se* but of some other practice instead, I don't find that suggestion very plausible. More generally, I find it implausible that the sustaining rules of belief management will necessarily disallow or discourage such outcomes. Consequently, I will propose an application of the model to the practice of belief management, assuming that these non-truth-directed outcomes promote the goal of belief management and bring value that helps sustain the practice.

One way to handle all of this is to posit distinct practices of belief management. One central practice is what I call *position man-*

agement. The goal of this practice is to equip us with beliefs needed to accomplish our goals and continue in a suitable way of life — that is, to help position us for success. We might need to accept or reject certain things just to get our projects off the ground or keep them going. If science would stall unless we believed that the pessimistic meta-induction from the history of science is fatally flawed, then we are well served to get ourselves to believe as much. More generally, if cognition would stall unless we believed that we know that our fundamental belief-forming processes are reliable, then we have an interest in getting ourselves to believe that. For those who would fall into despair and ruin without the crutch of supernatural religious beliefs, they are prudent to keep themselves believing such things. Again, the point of believing any of these things would be to help position us for success. Such positioning can pertain to global matters, as with the reliability of our basic belief-forming processes, or to exceedingly local matters, as with a cancer patient who needs to believe, “The odds are good that I’ll survive this illness,” in order to keep things together and have a chance at survival. Position management is a purely strategic belief management practice; whether the beliefs are *true* is beside the point. The point is effective positioning.

Another central belief management practice is what I call *inquiry management*, which itself seems to divide into at least two further practices. One side of inquiry is *attention management*. Which questions should we consider? And which of these deserve more of our time and energy? There is no point in considering ques-

tions useless to us, or in devoting more time and energy to a question than it deserves. We could be the best truth-detectors in the world, but that wouldn't matter if we never considered worthy questions. Acquiring worthless things won't sustain a practice. Attention management is also a strategic belief management practice, concerned with aiming cognition in the right direction. Its point is to set an appropriate cognitive agenda.

The other side of inquiry is *truth-detection*, or *cognition proper*. The goals and rules of position management don't apply here. And once a question has been considered, attention management has done its work, for better or worse. The only concern at this point is whether P is true. The aim of cognition is twofold: get the correct answer and avoid the incorrect answer. Getting the correct answer means believing P, if P is true, and believing not-P, if P is false. Avoiding the incorrect answer means not believing P, if P is false, and not believing not-P, if P is true.

What are the sustaining rules of belief management? The point of belief management in general is to *acquire useful beliefs*. This is a very broad goal, but that is appropriate because belief management is a very broad practice serving our interests across the whole range of life's activities. Call any belief, true or false, that benefits us a *useful belief*. Just as robbing some travelers was required to sustain the practice of Robin Hoodery, the formation of some beliefs will be required to sustain the practice of belief management. So at least one of the rules of belief management will be:

Belief's Imperfect Duty: Not infrequently, form a belief.

Forming useless beliefs won't sustain the practice, so we should also expect a rule of this form:

Belief's Rule of Permission: You may believe P only if believing P is useful.¹⁶

But as we have seen, beliefs can be useful in different ways. They can be useful because they position us to proceed and perform well enough. The practice of position management looks after that. And they can be useful because they are correct answers to questions that matter. The practice of inquiry management looks after that.

What are the rules of position management? Whether the belief is true is beside the point for this practice. Focusing for the moment on the most serious cases where the presence or absence of certain beliefs would immobilize us, two rules of this practice might be:

Positional Rule of Permission 1: You may believe P only if believing P will not immobilize you.

Positional Rule of Permission 2: You may refrain from believing P only if doing so will not immobilize you.

The practice might also incorporate standing recommendations:

Positional Recommendation 1: To the extent that believing P will help you, it is recommended that you believe P.

Positional Recommendation 2: To the extent that believing P will harm you, it is recommended that you not believe

¹⁶ Perhaps strict permission is the wrong deontic category. If so, a natural alternative would be *Belief's Rule of Discouragement:* You ought to believe P only if believing P is useful. Another weaker alternative is *Belief's Rule of Recommendation:* It is recommended that you believe P only if believing P is useful.

P.

What are that rules of inquiry, beginning with attention management? Attention management is presumably characterized by at least a pair of rules:

Attention's Imperfect Duty: Not infrequently, pose a question.

Attention's Rule of Permission: You may pose a question only if its answer matters.¹⁷

Once attention management has played its role, cognition begins.

What are the rules of cognition? It's doubtful that we break the rules of cognition by not answering every question posed. Nevertheless, there is an expectation that questions will be answered most of the time. This suggests that the duty is imperfect but still stringent, such that you're required to usually answer. Add to this that cognition aims at providing true answers and avoiding false ones, and a plausible pair of rules suggests itself:

Cognition's Imperfect Duty: Usually when the question whether P is posed, form a belief that P, or form a belief that not-P.

Cognition's Rule of Permission: You may form a belief only if it is true.

Just as the merry men achieved the goal of Robin Hoodery by achieving what Robin's Rules prescribe, so too will we achieve the

¹⁷ Again, I'm open to the possibility of replacing 'may' with 'ought' or 'should', and also to the suggestion that the rule is better understood as a stringent imperfect duty, placing restrictions only on the overall ratio of useful to non-useful questions asked.

goal of belief management by achieving what the rules prescribe. In the case of cognition, true beliefs are prescribed, so we achieve true beliefs. We are naturally equipped with powers of perception, are disposed to trust the word of others, and are habituated to make certain inferences. These are powers we have for achieving true beliefs: *powers of detection and discovery*. Some truths are harder won than others, but in a wide range of typical cases we can more or less directly achieve true beliefs through the exercise of these powers.

I am offering a view about the rules that normatively sustain practices. This does not commit me to the view that actions that occur within a practice are themselves made possible by the practice's rules, or indeed by any rules at all. Even if the rules of belief management pertain to the activities of questioning and believing, it does not follow that either *questioning* or *believing* is made possible by these rules, or that they could occur only in the context of such a practice. To the contrary, one might argue that inquiry management aims to manage the antecedently given activities of questioning and believing. In any event, here I am not taking a definite stand on whether questioning or believing can or must be metaphysically grounded in a prior practice.

9. *Knowledge and belief management*

If what I have said thus far is on the right track, then in light of an important result from recent epistemology, we learn something im-

portant about the relationship between knowledge and belief management. First I'll explain the result from recent epistemology. Then I'll explain what this entails about knowledge's relationship to belief management.

It has recently been persuasively argued that to know *just is* to achieve true belief.¹⁸ This thesis has been defended on a wide range of grounds, including the following: it provides a simple, elegant and extensionally adequate definition of knowledge, it solves the Gettier problem,¹⁹ it underwrites a compelling account of epistemic value, and it helps explain why knowledge is the norm of assertion. This is not the place to review these arguments. I am convinced that they succeed and for present purposes I will assume that they do.

Now suppose that the rules of cognition prescribe acquiring true beliefs. And suppose that it is *by achieving what the rules prescribe* that we sustain a practice. Thus it is *by achieving true belief* that we achieve the goals of cognition and thereby sustain the practice. So since achieving true belief just is knowing, it is *by knowing* that we sustain the practice of cognition and, in turn, inquiry. Knowing is the only way to do this. It is by knowing, and by knowing alone, that we hold the practice together.

Knowledge is not the goal of belief management. The sustaining rules of belief management don't mention knowledge. But know-

¹⁸ Sosa 2007, Greco 2010, Turri 2011a, Turri in press a. Important historical antecedents include Aristotle, Descartes and Reid.

¹⁹ This includes so-called "fake barn" cases, which many philosophers categorize as Gettier cases. For further discussion and review of relevant literature, see Turri 2011a; Turri, Buckwalter and Blouw 2015; Blouw, Buckwalter and Turri in press; and Turri in press b.

ledge still plays a unique and essential role: it is the constitutive means and sustenance of cognition and inquiry.²⁰

It is widely believed that there is an important normative connection between belief and knowledge.²¹ I have proposed one special way that knowledge sets the normative standard for belief: knowing is essentially the unique way of normatively sustaining cognition and, thereby, inquiry. Knowing is the constitutive means and sustenance of these central practices of belief management. Of course, there are other belief management practices, such as position management, where knowledge does not play this role, and I have said nothing about whether knowledge is a norm of belief *itself*, independently of any practice of managing beliefs. But that is consistent with knowledge playing a special and unique normative role in our practices of cognition and inquiry.

Acknowledgments — For helpful feedback, I thank Matt Benton, Wesley Buckwalter, Tim Kenyon, Rachel McKinnon, Adam Morton, Ernest Sosa, and Angelo Turri. Special thanks go to Ben Jarvis. This research was supported by the Social Sciences and Humanities Research Council of Canada, the Association of Commonwealth Uni-

²⁰ Again, I am open to the possibility of replacing Cognition's Rule of Permission with a Rule of Discouragement involving 'ought' or 'should'. I am even open to the possibility that the rule states only an imperfect duty, such as: For the most part, form a belief only if it is true. If the relevant deontic rule were such an imperfect duty, then we could still say that knowledge is essentially the central sustaining tendency of cognition: in order to normatively sustain the practice, the central tendency of cognition must be knowledge. Finally, I am open to the possibility that knowledge does not, strictly speaking, require truth, but only approximate truth (see Turri 2011c).

²¹ E.g. Williamson 2000, Sutton 2007, Huemer 2011.

versities, the British Academy, The Character Project at Wake Forest University and the John Templeton Foundation, the National Endowment for the Humanities, and an Early Researcher Award from the Ontario Ministry of Economic Development and Innovation.

References

- Alston, William P. 1991. *Perceiving god: the epistemology of religious experience*. Ithaca: Cornell University Press.
- Blouw, Peter, Wesley Buckwalter, and John Turri. In press. Gettier cases: a taxonomy. In *Explaining knowledge: new essays on the Gettier problem*. Ed. Rodrigo Borges, Claudio de Almeida, and Peter Klein. Oxford University Press.
- David, Marian. 2005. Truth as the primary epistemic goal: a working hypothesis. In *Contemporary debates in epistemology*. Ed. Matthias Steup and Ernest Sosa. Malden, Mass.: Blackwell.
- DePaul, Michael. 2001. Value monism in epistemology. In *Knowledge, truth, and duty*. Ed. Matthias Steup. Oxford: Oxford University Press.
- DeRose, Keith. 2002. Assertion, knowledge, and context. *Philosophical review* 111.2: 167–203.
- Greco, John. 2010. *Achieving knowledge: a virtue-theoretic model of epistemic normativity*. Cambridge: Cambridge University Press.
- Huemer, Michael. 2011. The puzzle of metacoherence. *Philosophy and phenomenological research* 82.1 (January 2011): 1–21.
- James, William. 1897 [1948]. The will to believe. In *Essays in pragmatism*. Ed. Alburey Castell. New York: Hafner Press.
- MacIntyre, Alasdair. 1999. *Dependent rational animals: why human beings need the virtues*. Chicago: Open Court.
- Millikan, Ruth. 2005. *Language: a biological model*. Oxford: Oxford

University Press.

- Rawls, John. 1955. Two conceptions of rules. *Philosophical review* 64: 3–32.
- Richerson, Peter J. and Robert Boyd. 2005. *Not by genes alone: how culture transformed human evolution*. University of Chicago Press.
- Searle, John. 1969. *Speech acts*. Cambridge: Cambridge University Press.
- Simester, A. P., Ed. 2005. *Appraising strict liability*. Oxford: Oxford University Press
- Sosa, Ernest. 2007. *A virtue epistemology*. Oxford: Oxford University Press.
- Sosa, Ernest. 2011. *Knowing full well*. Princeton: Princeton University Press.
- Sutton, Jonathan. 2007. *Without justification*. Cambridge, Mass.: MIT Press.
- Tomasello, Michael. 2008. *Origins of human communication*. London: MIT Press.
- Tomasello, Michael. 2009. *Why we cooperate*. London: MIT Press.
- Turri, John. 2011a. Manifest failure: the Gettier problem solved. *Philosophers' imprint* 11.8 (April 2011): 1–11.
- Turri, John. 2011b. Believing for a reason. *Erkenntnis* 74.3: 383–397.
- Turri, John. 2011c. Mythology of the factive. *Logos & Episteme* 2.1: 143–152.
- Turri, John. 2013. The test of truth: an experimental investigation of the norm of assertion. *Cognition* 129.2: 279–291.
- Turri, John. 2015. Evidence of factive norms of belief and decision. *Synthese*, 1–22. doi:10.1007/s11229-015-0727-z
- Turri, John and Peter Blouw. 2014. Excuse validation: a study in rule-breaking. *Philosophical Studies*. doi:10.1007/s11098-014-0322-z
- Turri, John, Wesley Buckwalter, and Peter Blouw. 2015. Knowledge

and luck. *Psychonomic Bulletin & Review*, 22(2), 378–390.

Turri, John. In press a. From virtue epistemology to abilism: theoretical and empirical developments. In C. B. Miller, M. R. Furr, A. Knobel, & W. Fleeson (Eds.), *Character: new directions from philosophy, psychology, and theology*. Oxford University Press.

Turri, John. In press b. Knowledge judgments in “Gettier” cases. In *A companion to experimental philosophy*. Ed. Justin Sytsma and Wesley Buckwalter. Wiley-Blackwell.

Turri, John. In press c. The radicalism of truth-insensitive epistemology: truth’s profound effect on the evaluation of belief. *Philosophy and Phenomenological Research*.

Williams, Bernard. 1978. *Descartes: the project of pure enquiry*. New York: Routledge.

Williamson, Timothy. 2000. *Knowledge and its limits*. Oxford: Oxford University Press.