

The Concept of Legitimacy

N. P. Adams
npadams@virginia.edu

Draft 2022
Please do not circulate
Comments welcome

1. Introduction

What does it mean for something to be legitimate? Although the idea of legitimacy is commonplace and political philosophy takes state legitimacy in particular as one of its central concerns, the concept of legitimacy (LEGITIMACY) has received little attention.¹ In this article I offer a new understanding of LEGITIMACY by analyzing the role of legitimacy discourses in social practices.

To explain how legitimacy discourses function, I draw on sociologists' notion of boundary work, "one of our most fertile thinking tools" (Lamont and Molnár 2002, 169; Gieryn 1983). On this approach, social practices are the result of practitioners implementing normative boundaries, which result in significant patterns of behavior.² Social practices so construed can face a problem of bad boundaries. When practitioners follow bad boundaries, they act in ways that realize the extant version of the practice but fail to realize the practice's goals or teloses. Legitimacy discourses solve the problem of bad boundaries by providing practitioners minimal telic standards. Legitimacy discourses thus perform a gatekeeping function, ensuring that a minimally acceptable version of the practice is being implemented by correcting the extant version of the practice when it goes too far astray from its telos. Taking a functionalist approach to meaning (Chalmers 2011), we get the following proposal for the meaning of LEGITIMACY: E is a legitimate K if it meets minimal telic K-standards.

This general account of legitimacy has striking results when applied to the political domain. If the gatekeeping account is correct, then theorists have

¹ Following a convention, terms in small caps refer to the concept itself.

² Boundary work falls into approaches that focus on the production of meaning and identity rather than those that focus on rational choice (cf. Brennan et al 2013). As far as I can tell it has been ignored in philosophy.

misunderstood the nature of political legitimacy norms. Rawls and those following him present legitimacy as secondary, arising only as a solution to the problem of disagreement about justice. But legitimacy norms are necessary because of the nature of normative orders, not because of disagreement. Political legitimacy discourses are crucial for implementing political institutions because of their gatekeeping function.³ Even Rawls' well-ordered society that unanimously endorses justice as fairness would need political legitimacy norms.

Here's the plan. In section 2, I sketch a basic model of social practices as normative orders. In section 3, I make this model more realistic by showing how practitioners must constantly exercise discretion to implement social practices. In section 4, I introduce kind-relative properties so that I can argue, in section 5, that legitimacy is a kind-relative property that helps us solve the problem of bad boundaries by gatekeeping. In section 6, I apply this to law, first as the way societies regulate social practices generally and then as a social practice itself. In section 7, I reconsider the relationship between political legitimacy and justice.

2. Social practices

We are concerned with social practices in the broad sense of regularized patterns of behavior that are the result of people conforming to behavioral standards or norms (Bicchieri 2016). This expansive domain is appropriate given the apparently wide range of contexts where LEGITIMACY makes sense. Although details vary, starting from social practices in this approximate sense is common in social theory (e.g. North 1991; Tuomela 2002; Searle 2010).⁴ The basic model of social practices I sketch in this section shows how normative orders function and emphasizes the place of role-kinds.

Social practices are normative orders: they coordinate behavior by applying various norms to entities within the practice.⁵ Primary among these entities are practitioners, i.e. persons who follow the norms that apply to them. Practitioners put practices into practice. When practitioners share a significantly overlapping

³ This is about functional necessity. Nobody, of course, needs the English term 'legitimacy' and social practices gatekeep in a variety of ways.

⁴ Sometimes these are also termed social institutions but I prefer to reserve 'institution' for the type of formalized practice that is also called an organization.

⁵ I am using 'coordination' in a non-technical sense. Practices in my sense need not be collective or joint activity.

sense of what the norms require of them and accept that they should abide by the norms, they will generally act according to the norms in similar ways. (Here I leave aside practitioners' motivations for following the norms, which will vary widely, including internal norm acceptance and external incentives.) The resulting consonant behavior produces some purported good and characteristically aims at some goal or end. In this sense, social practices are teleological, although not in the sense of being intentionally created to pursue an end (Miller 2010).

When practitioners share norms, they have a common behavioral standard. Without shared norms, individual actions are unlikely to overlap in any interesting way; even if the actors all intentionally pursued a shared goal, conflicting individual judgments about how to proceed would lead towards haphazard and often contradictory behavior. The practice would not be able to succeed and would not exist in a significant sense without practitioners feeling bound by and sharing an understanding of the relevant norms. (Though the norms need not explicitly enjoin coordination.) Social practices are thus defined by behavior patterned by norms and only exist in virtue of their ability to establish and maintain such patterns.

Patterns are defined by their boundaries. A norm that people walk across a certain path in a field defines a boundary between acceptable and unacceptable behavior. If the norm is accepted within a social group, that normative boundary will result in a social boundary: when walking across the field, most people will follow the path most of the time. We can describe the pattern in terms of behavior that falls on one side or another of the boundary, either following the norm and walking along the path or violating the norm and walking elsewhere. This allows us to predict action, to show how this pattern will interact with others, and so on. We can explain why the grass is worn there but nowhere else or why people (lightly) sanction people who stray off the path. These patterns of behavior shape the social world and we can identify them by their boundaries.

Complex social practices mostly coordinate norms and behavior by defining role kinds. Assigning unrelated bundles of norms to practitioners would be very unwieldy and inefficient for a variety of reasons. Instead, practitioners take up roles that bundle norms into a comprehensible whole. Roles are the functional units of most practices. They enable the higher-order coordination of behavior that results in a successful practice, for example by reducing the demands of

interpersonal trust and cognitive load (Searle 2010; Ritchie 2020).⁶ Roles are partly defined by deontic status, a collection of incidents including rights, duties, permissions, and powers (Hohfeld 1919). Roles also include important elements such as ideals and more general expectations. The status, ideals, and expectations that constitute a role must be understood in light of the function of the role in the practice.

All kinds of entities are assigned roles (and status). Here we are primarily focused on practitioner roles, which have agential status. Entity kinds also include mere subjects (dogs), objects (stop signs), events (recessions), ideas (theories) and actions (voting), among others. I may ask a student whether their raised hand was intended as a vote; I am asking whether the action falls within the role and so should count in the way votes do. To see how practitioner roles must be understood functionally, consider the familiar roles of professor and student in the university. The status that defines being a professor must be directed at enabling the university to pursue some overall purpose or telos, for example education. Professors' teaching obligations are then matched to students' learning obligations. When kinds are defined and interrelated well, they enable us to pursue the practice's telos.

So we have a basic model of a social practice. Practitioners pattern their behavior according to the norms of the practice. In any given moment this takes the form of practitioners deciding how to act by applying the practice-defined kinds to assess their practical context. The routinized essence of social practices is practitioners determining whether an entity E is a kind K by applying K-standards and then, if E meets the standards, treating E as a K by following the relevant K-norms. For example, universities involve routinized classes. Practitioners gather with the expectation and intention that the standards for a class meeting will be met, take up coordinated roles, and engage in event specific behavior.⁷ If the practice is constructed and practiced well, recursive events of this sort will lead to a professor having taught a course and students learning.

⁶ Roles do more than name a status and my notion of a role is not identical to Searle's status function Y term. Searle's notions of roles and functions are, to my mind, unclear but at least clearly distinct from my usage, and my notion of a social practice is quite a bit broader than Searle's. [Acknowledgement removed for anonymous review.] The role can survive through changes in individual elements of status; this is not only a descriptive fact but a desirable feature of the role as a functional unit within a social practice.

⁷ The standards for class meetings will involve, in turn, nested judgments of Es as Ks: start times, classrooms, students, texts, and so on. Es and Ks range widely.

The basic model is a mechanical perfection. Internally, the normative structure is perfect. Roles are defined clearly, completely, and consistently such that any practical context a practitioner could face is subsumed into the practice. The norms are similarly clearly, completely, and consistently defined so that there is a univocal fact of the matter about what the practitioner should do. Externally, the normative structure is overriding. Practitioners are presented as practice automatons who do as directed without exception. They apply and follow norms as directed, they do not make mistakes, and they do not have other demands on their action. Considered as a whole, the practice appears autonomous with practitioners as cogs. This model illustrates how social practices function as role-based normative orders but it does not need LEGITIMACY. Legitimacy discourses are solutions to problems it does not have.

3. Boundary work

Let's introduce some reality. Practices do not appear or operate of their own accord. People enact the practice by taking up roles and patterning their behavior with norms. This is where boundary *work* enters the picture. Which norms people follow as practitioners and how they follow those norms are social questions that are only ever answered by particular people interpreting their practices in particular contexts. Social practices have to put into practice collectively, and so always contestedly.

Consider the classic sociological example of John Tyndall's 19th century boundary work on the social practice of science (Gieryn 1983). When a phrenologist was being considered for the Chair of Logic at the University of Edinburgh, a contest ensued over the boundaries of science. Tyndall advocated for a version of science that excluded phrenology while the phrenologists advocated for a more expansive version that would give them access to the social practice's status and resources. No matter where the boundary of science lies according to our preferred philosophical solution to the demarcation problem, science as a realized social practice is shaped by where practitioners believe the boundary of science lies and so what they will treat as a science (cf. Rawls 1999, 48). Eventually Tyndall and his allies convinced decision-makers at the university as well as members of the community to adopt their understanding of science and so to treat phrenology as a pseudoscience. One version of the social practice of science was implemented and

one was set aside.

Practitioners constantly face choices between versions of practices.⁸ Normative orders necessarily underdetermine practitioner choice because they function by intervening in the practical decision-making of agents. Norms are directions for exercising agential discretion, not programming. The basic model oversimplifies insofar as it treats social practices as static objects rather than processes, as automatons' mechanized behavior rather than as agents' patterns of choosing. Social practices only ever exist in virtue of people choosing, again, to follow practice norms by playing their role in the practical contexts they face over time. Discretion is intrinsic to social practices.

Practitioners have both internal and external discretion. Internally, practitioners have discretion in the sense that they can follow practice norms in many different ways and so must choose how to follow. Internal discretion is ineliminable because being a practitioner is not exhausted by simply conforming to fully programmed directives. Norms are necessarily underdetermined: they can never fully specify how they are to be practiced in particular contexts. The command "Open the door" underdetermines which part of the body is used, how much force is used, how many nanoseconds the door is opened for, how wide the opening is, and so on. It partly determines many of those elements but cannot fully specify them. Underdetermination is especially prominent in systems of codified rules like law (Schauer 1991). Further, practitioners rarely face decisions that are determined by a single element of their role; norms overlap in application, ideals and expectations exert pressure, and so on. Acting out one's role in a practice requires making the role's practical demands determinate in a specific context in ways that the role's norms cannot fully dictate. Call this contextualization.

Externally, practitioners have discretion in the sense that they can refuse to act as directed, instead acting against or outside the practice. External discretion is ineliminable because practitioners are not exhausted by their participation in any particular practice. Practitioners always act as whole persons whose actions

⁸ This notion of a version is quite vague. Clearly the version of science that includes phrenology is different from the version that excludes it. But not every variation counts as a new version; just as with software versions, when changes constitute a new version is a judgment call that partly relates to our understanding of the practice. Determining what counts as a contest between competing versions and what is just the normal temporal froth of recreating practices implicates the very contest it judges.

implicate their whole selves and there are always intersections with other aspects of their identity. Due to the spatiotemporally local nature of agency and action, when I stay late for a work meeting, I cannot spend that time with my family. Practitioners are thus always exercising their discretion about which practice's norms to follow in particular contexts and how to prioritize the various demands on their agency. Call this situating.

Contextualization and situating are ineliminable from practices. This is good—discretion is how we make social practices livable. For a social practice to actually be part of our lives, to be practiced by people like us in our circumstances, it must be flexible enough to be placed into our lives and into our contexts, against other practices and demands. Discretion enables practices to be adaptable to changing contexts, to update norms that have become sclerotic, to respond to the needs and desires of practitioners. However, discretion also goes hand in hand with hard cases.

Hard cases occur “when no settled rule dictates a decision either way” (Dworkin 1978, 83). Such cases are at the core of general jurisprudence, asking how judges should exercise discretion. These are hard cases partly because of courts' adjudicative function: not deciding is not an option and the contest is at least not trivial. Social practices are suffused with discretion but most discretionary choices are trivial. I get to choose at what decibel level I lecture; the university functions perfectly well as long as I choose within an understandable range. Neither the law nor the university take a stance on this discretionary choice, and neither needs to. But some choices matter for practices in the sense that leaving them up to individual discretion would harm the practice's functioning. Hard cases, then, are choices that the practice currently leaves open to practitioner discretion but where practitioner discretion is inadequate.⁹

The obvious way to solve a hard case is to introduce a new practice norm, since norms reduce discretion.¹⁰ Norms, and procedures for introducing norms, can be either formalized or non-formalized (cf. Brennan et al 2013, ch. 3).¹¹ Law

⁹ Lawsuits are a mechanism for demanding authoritative settlement, so making practitioner discretion formally inadequate.

¹⁰ Stating the norm reduces internal discretion for practitioners insofar as they are committed to the practice; it reduces external discretion insofar as it is attached to enforcement mechanisms of various sorts. Thus the distinct character of law as both fact and norm (Habermas 1996).

¹¹ The term ‘formalized’ emphasizes the necessary process, and so social practice, of formalization. The contrasting ‘non-formalized’ avoids the dismissive implications of informality; as Bourdieu

stands at the maximally formalized end of this spectrum, solving hard cases by introducing formalized norms in a formalized way. Courts' legal authority entails that their decisions bind, including by creating precedent. Law also has higher-order hard case settlement procedures; in the United States, for example, when federal appellate courts disagree, the Supreme Court takes over.

Reflection on this example demonstrates the limits of formalized norms. Notoriously, Supreme Court decisions are not immune from underdetermination. They render the law more determinate in some ways but also introduce new norms in new language. More importantly, Supreme Court decisions are not self-executing, so are open to pushback via contextualization and situating. The integration decision in *Brown v. Board of Education* resulted in decades of contested implementation. Legal decisions only pattern behavior in virtue of practitioners enacting them. Some legal practitioners have official roles and so play a special part in enacting law but, in the end, communities implement law by building legal norms into collective and individual decision-making. The relevant social practice is the legal order, a social order that is successfully organized via law. Supreme Court decisions change the law but how exactly that changes the legal order is an open question that can only be answered in practice by practitioners.

How do practitioners exercise coordinated discretion when formalized norms run out? A regress seems to threaten. Hard cases are solved by introducing new norms but all norms include discretion, so introducing a new norm also introduces new discretion. This can work by shifting the point of discretion, and perhaps reducing its scope. When the court decides, practitioners still have some discretion about how to implement the decision but at least there is an authoritative decision to implement. Norm introduction slows but does not halt this regression.

Social practices rely on non-formalized implementation norms to halt discretionary regression. To make practices collectively livable, practitioners coordinate discretion with what organizational theorists call the "logic of appropriateness" (March and Olsen 2011). Professors teach classes not by mechanically applying the formalized norms but by situating and contextualizing the norms following a shared sense of what is appropriate for professors. Non-formalized norms necessarily structure the implementation of formalized norms.

(1977, 164-71) emphasizes, non-formalized norms are often the most important and can be the most powerfully enforced precisely because they are non-formalized, as in the case of what he calls doxa.

These non-formalized norms arise from practitioners' overlapping, more holistic understanding of what professors should be like and the role they should play in the university (and so, what the university should be like).

The faculty handbook doesn't need to stipulate that I lecture at a certain decibel level because the non-formalized norms of appropriateness handle that. We understand that professors lecture to teach and so lectures need to constitute successful communication. The non-formalized norm that I teach at an understandable level is implicit in the very idea of the university and the role of professor. This norm is active; if I violate it, students or colleagues can bring it to bear. If I persistently violate this or many other implicit communicative norms, not only would I fail at being a professor, it is also doubtful that I even grasp what it is to be a professor. The broader social and linguistic context in which the practice exists provides resources for implementation that simultaneously channel the forms the practice can take. Competent practitioners necessarily possess a wide variety of appropriateness norms that guide choices when formalized norms run out. Instead of choosing randomly, practitioners can compare potential versions and choose to implement the better one. One way to identify which versions are better is to compare their kind-relative properties.

4. Kind-relative properties

Imagine I am at the store and offer what appears to be a twenty-dollar bill for my purchase. Sometimes people counterfeit money so cashiers have ways of testing putative dollars. These tests could be enabled with a new concept, SCHMOLLAR. The associated schmollar discourse would include standards for distinguishing dollars from schmollars and behavioral norms for how to treat each. By enabling comparisons between putative dollars, this conceptual discourse functions as a solution to the problem of fake money.

But the problem of fakes arises in a wide variety of practices. Thus, we have more general concepts to enable our solutions. Instead of asking about dollars and schmollars, the cashier can ask about genuine dollars and fake dollars, and we can ask about genuine and fake handbags, or genuine and fake friends, or a wide variety of other entities. But it also clear that what makes something authentic varies across kinds: fake dollars, fake handbags, and fake friends are very different entities in very different practices. Identifying each is a different project and we should also react

to each in quite different ways. Authenticity discourses apply the same concepts of GENUINE and FAKE across practices but use different standards of authenticity and license different responses to authentic or inauthentic entities. This enables us to use the same discourse to perform the same function across very different practices.

This social function is identifiable in language: “genuine” is an essentially attributive adjective (Sibley 2001, 169). Such adjectives ascribe properties that can only be assessed relative to membership in a kind. The claim that E is genuine must be understood as claiming that E is a genuine token of kind K. Famously, Geach (1956) argues that “good” is essentially attributive. Other essentially attributive adjectives include “real,” “tall,” and, “big.” As Sibley (2001, 175) puts it, “To ascertain that *x* is big, you must have a class of things in mind.” Bigness standards are relative to kinds. A big mouse is much smaller than a big elephant, which is much smaller than a big skyscraper. We do not always make the kind explicit because the context, including the entity itself, can make kind membership salient. But these comparisons necessarily assume that the entity is being evaluated as a member of a kind. Essentially attributive adjectives point us to kind-relative properties.

This ordinary usage calls for explanation, especially concerning the function of discourses that employ these concepts. Geach’s point about GOOD is not about language but about how we assess properties of this sort (Thomson 2009, 17). He argues that because goodness must be relative to kinds, it is a grave mistake to search for a reified property of goodness that all good things share. When we use adjectives attributively, we draw attention to membership in some kind and its associated standards because there is no way to assess the truth value of the attributive use without those standards.

Cross-practice discourses to assess kind-relative properties like authenticity and size are useful because in-kind comparisons of the same sort are useful. BIG and SMALL mean the same thing across practices; bigness and smallness pick out the same comparative properties. But the standards for assessing big and small are relative to each kind. If we only have one standard of size, it would be very difficult to, e.g., compare mice to each other because they would all be very small. Crucially, then, it is the standards of kind-relative properties that are relative to each kind, not the properties themselves or the function of assessing those properties.

BIG and SMALL differentiate members of the same kind. Do GENUINE and

FAKE? If we ask whether something is a genuine mouse and discover that it is a cleverly designed robot, we don't seem to be differentiating members of the same kind. Mice and robots are different kinds. This shouldn't lead us to conclude that GENUINE is not assessed according to kind-relative standards. Instead, it points us to the fact that different evaluations play different functions. Size evaluations are useful for some things, genuineness for others.

We use GENUINE when fakes or forgeries are salient possibilities. When asking whether something is a genuine mouse, the standards for GENUINE are affixed by the kind mouse but the class of objects we are differentiating among are not mice but putative mice. There's no sense in asking whether an apple is a genuine mouse because it's not even putatively a mouse. But it does make sense to ask whether the thing that appears to be a mouse but is also sparking a little around its tail is a genuine mouse or a fake mouse. A fake mouse is, of course, not a mouse at all; a fake mouse is a putative mouse that does not meet the kind-standards for membership in the biological category. So an authenticity discourse must have a class of evaluands that includes entities that are not part of the practice. By incorporating external entities into this liminal part of the practice, the authenticity discourse regulates inclusion in the practice, enabling it to function under the threat of fakes.

GENUINE is a conceptual tool that we use in authenticity discourses to solve the problem of fakes: the cross-practice property of genuineness is evaluated according to kind-relative standards. Fakes disrupt the routinized essence of practitioners evaluating whether Es are Ks by applying shared K-standards. Convincingly fake Ks will lead practitioners to treat non-Ks as Ks, thus impairing the practice's functioning, for example by diluting the resources and status of Ks. Fakes are a threat when the normal K-standards could be inadequate for identifying genuine Ks. So, to solve the problem of fake Ks, GENUINE discourses must employ a more stringent set of K-standards. While visual inspection of dollars may be a good enough dollar-standard for most transactions, enhanced standards might be necessary for higher denominations where fakes are more common.

We have detoured into authenticity because it gives us a clue: legitimate meant legally genuine, as in a legitimate child who thus has the legal status of heir (OED). This points to the idea that 'legitimate' is an essentially attributive adjective: "E is legitimate" means "E is a legitimate K." If legitimacy is a kind-relative property,

then we can only determine whether E is a legitimate K by reference to K-standards. The meaning of LEGITIMATE and ILLEGITIMATE is the same across practices; we can coherently talk about legitimate states, legitimate contracts, legitimate questions, or legitimate wins. But legitimacy standards must be kind-relative to play their distinctive function within each practice: gatekeeping.

5. Gatekeeping

Analogously to the problem of fakes, social practices face the problem of bad boundaries. Consider a version of football where the purpose of a season is to identify the best team. If the season consists only in a round-robin tournament, the team with the most wins is the best team. In this version of the social practice, the role of the kind “win” is to identify the best football team. But not all wins are created equal. If team A beat team B when B was missing their best players due to injury, we might say that this win shouldn’t count. Although A won, this win does not serve the purpose of counting wins in this context because this win is not a good indicator of A’s quality.¹² The extant practice boundaries define the kind “win” in a manner that doesn’t always help us identify the best team. Well drawn boundaries will define wins so that they indicate team quality. Bad boundaries will define wins so that they include some bad indicators and exclude some good indicators. So bad boundaries paradigmatically¹³ result in bad Ks: putative Ks that shouldn’t count as Ks because they don’t function well as Ks and putative non-Ks that should count as Ks because they would function well as Ks. For simplicity I’ll focus on the former kind of bad K (false positive or Type I errors).

The problem of bad boundaries is that badly drawn kinds result in bad Ks, which in turn result in a malfunctioning practice. Bad boundaries are bad from the internal perspective of the success of the practice. Other perspectives ground other problems of bad boundaries—how social practices handle such problems is an important topic, it is simply distinct from my concern here.

The problem of bad boundaries is closely related to practice

¹² This is a real problem, sometimes solved by weighting win-loss records by opponent strength.

¹³ More generally, bad boundaries lead to bad borderline cases. The paradigmatic bad borderline K is as the text describes: a K that’s on the wrong side of the boundary. But another bad borderline K is too close to the boundary. There should be no doubt that E is a K (or not a K). I can say that Martin Luther King, Jr. is a legitimate philosopher as a way of emphasizing that even raising the question of his status indicates a misplaced boundary.

underdetermination and so practitioner discretion. All social practices can be multiply realized and must be re-realized in changing contexts. Practitioners exercise discretion over which version to realize. This choice is often between (de re) better and worse versions. If practitioners implement bad boundaries, their routinized evaluation of Es as Ks will degrade the practice's functioning. To recognize this as a problem, practitioners need to be able to take up a stance internal to the practice but distinct from the extant version of the practice. That is, they need to appeal to a better version of the same practice. This enables immanent criticism by giving practitioners a perspective internal to a shared understanding of and commitment to the practice from which to criticize the extant version of the practice.¹⁴

Practitioners qua practitioners have a guide for evaluating better and worse versions ready to hand: teloses. Practitioners must have a practical understanding of the telos of the practice and of the kinds within the practice because they ground the logic of appropriateness. When practitioners situate and contextualize practices, they necessarily interpret the practice to implement a version that they offer as appropriate (Dworkin 2011). They present themselves as practitioners in good standing, so putting into practice a version that other practitioners recognize as fitting with the ongoing practice. All such interpretations can only be understood as an appropriate version of the practice qua pursuit of its telos. We get telic standards from teloses. The claim that A's win against B shouldn't count draws on telic standards for wins. The telos of the kind "win" is to measure the quality of the teams, serving the telos of the season to determine the best team. Wins against teams who are short-handed do not help us measure team quality, so they are degenerate instantiations of the kind: counting them does not serve the purpose of counting wins.

Legitimacy discourses are a solution to the problem of bad boundaries. Legitimacy discourses help us identify bad Ks and coordinate our response to them. Putative Ks that shouldn't be treated as Ks are illegitimate Ks. (The most salient way to be a putative K is to be an extant K but there are many other ways.) Illegitimate wins are wins that shouldn't be counted as wins. Illegitimate states are states that

¹⁴ So kind-relativity does not imply that the only available standards are practice-relative. Standards must be relativized to the nature of the practice, not relativized to standards practitioners currently endorse. Thanks to an anonymous reviewer for pushing me to clarify this.

shouldn't rule. Illegitimate children are children that shouldn't inherit.¹⁵ In each case, E is a putative K but should not be treated according to the status that a K normally has. Legitimacy discourses employ telic standards to identify those putative Ks that are inadequate *as Ks*.¹⁶ As noted above, bad boundaries also give rise to false negatives. Putative non-Ks that should be treated as Ks are legitimate Ks. Legitimate wins count towards our assessment of team quality and legitimate states have the right to rule.¹⁷

This explains what we add when we ask after an entity's legitimacy. When we ask after an entity's goodness, for example, we are admitting that they belong to the kind but are evaluating whether they realize the virtues of the kind. When we ask after an entity's legitimacy, by contrast, we have yet to admit them into the kind as a practical matter: we are considering whether we should treat them as a member of the kind at all. This presumes that there is some work to be done still. Simply asking whether they are a K is not enough; we imply that meeting the extant standards for K-ness doesn't settle whether they should be treated as a K, that another set of standards should orient our behavior.

Legitimacy discourses therefore serve as a corrective on the extant practice. When a practitioner judges that E is an illegitimate K, then they will not treat E as a K. If practitioners coordinate these judgments in a discourse, their pattern of behavior will shift its boundary to exclude E from K-ness. Treating this K as not a K is an inelegant solution. It handles a particular case but does not address the underlying problem of how this E became K. Presumably the processes that lead to this result will continue to generate illegitimate Ks. Trying to correct this on a case-by-case basis is a kludge that depends on a sufficiently shared non-formalized legitimacy discourse rather than the core machinery of a practice.¹⁸

This kludginess is mitigated by the internal logical of legitimacy judgments.

¹⁵ Our discomfort with this application confirms the theory. Since we think that the status automatically follows being a child, rather than being a child born within a legal marriage, we don't think there are any salient standards for evaluating merely putative children.

¹⁶ See Sangiovanni (2019) for another telic view of legitimacy. My emphasis on the function of telic standards in social practices is more general.

¹⁷ States may be an example where there is a presumption against legitimacy, so asserting that E is a legitimate state is always discursively sensible—all rulers are putative, their credentials always worth demanding.

¹⁸ This may not be correctible; there are limits of how ideally K-ness can be instantiated in an actual practice, due to K's relation to other roles, the practice pursuing multiple aims, the costs of actual institutionalization, and so on.

Telic standards appeal to practitioners qua practitioners: they say the practice is malfunctioning. Calling E an illegitimate K says not to treat E as a K because treating it as a K does not serve the purpose of the practice. This implies that the extant boundaries are insufficient for properly orienting ourselves to E. It also implies that extant procedures for correcting boundaries are insufficient, since we are having to correct the practice by violating it.¹⁹ When K-ness goes wrong in less worrying ways, social practices generally have other ways of handling bad Ks. Acting against the practice also risks destabilizing the practice. So the implication is that this bad K is such a bad K that we should fix the practice by shifting the boundaries ourselves but also that we should fix the underlying boundary-setting machinery of the practice. Due to the type of standards they employ and the discursive context in which employing these standards becomes useful, legitimacy judgments push reform both at the level of particular entities in the practice and at the practice itself.

For all this to make sense, legitimacy standards must be quite minimal. Just asserting that the practice has diverged from some standard is not enough. Practitioners are always instantiating the practice by interpreting it, situating and contextualizing it according to standards they endorse. There will be constant divergence, most of it trivial, and much non-trivial divergence will be handled by extant internal correction procedures.²⁰ If legitimacy standards were demanding, they would often be violated and practitioners would deviate so often from practice norms in an attempt at correction that the practice would be destabilized. Less demanding standards identify a line that cannot be crossed qua minimally acceptable version of the practice. Crossing this line warrants violating the extant practice norms from the perspective of the practice itself. Other telic standards will perform other functions. An EXCELLENT discourse gives practitioners a way of identifying Ks that realize the telos particularly well, giving us goals and other kinds of practical guides. The maximal telic standards associated with excellence enable us to improve and perfect our practices.

The minimal telic standards associated with legitimacy enable us to

¹⁹ Once this implication is well established, it can be the subject of strategic employment. I might assert that A is a legitimate K not to question the extant K-procedures but precisely to bolster them: it says that the extant practice fits its telos.

²⁰ So, *ceteris paribus*, legitimacy standards will be decreasingly demanding in proportion to increasingly effective internal correction procedures.

gatekeep our practices. Gatekeeping ensures minimally acceptable boundaries are being instantiated by correcting the practice when boundaries go too far astray. So here is the gatekeeping account of LEGITIMACY. That A is a legitimate K means that A meets some set of minimal telic K-standards and so practitioners should treat A according to the associated K-status. That A is an illegitimate K means that A does not meet those standards, so practitioners should not treat it that way. Legitimacy discourses coordinate these judgments to gatekeep social practices. This is a general account of the core meaning across any social practice; penumbral meanings of course follow. This generality explains the inclination to falsely ascribe a deep polysemy to LEGITIMACY.

6. Law

As I've described it, gatekeeping is a way that practitioners exercise control over the version of the practice they are instantiating by choosing among possible boundaries. There are many ways to perform this general function since choosing which version to put into practice is a standard part of boundary work. Tyndall's rhetorical boundary work appealed to a shared understanding of science and its telos without necessarily employing LEGITIMACY. We eventually developed PSEUDOSCIENCE as another tool for drawing some boundaries of science.

Gatekeeping's distinct usefulness becomes clear when considered against the background of all the other tools that we use to instantiate different versions of practices. Non-formalized telic norms are vague, not easily enforceable, and not guaranteed to be shared. They grant an insecure social grasp over the entities they apply to. In general, practices function by stably reproducing some default version, especially via what Bourdieu (1977) calls habitus: the internalized, embodied way that we situate ourselves into the social world, enacting roles by default. But precisely because of the facts of discretion and the need to shape practices differentially over time, practices also often have ways of enabling contests over versions. Highly formalized rule-making is the limit case.

Practitioners always situate social practices in and against other practices in ways that affect which versions get put into practice. In a particular course meeting, the professor and students are choosing which version of this course and courses in general to instantiate. But they cannot choose anything. They are constrained by the form that courses are given by their particular university and the broader

cultural understanding of courses. They are also constrained by what they are willing to choose given their understanding of their roles in the university, their identities more generally, the other roles they inhabit, and much else. And, of course, by law.

Law enables formalized, enforced, public control over social practices. Modern law plays an integrative function between various social domains, limiting which versions of social practices can be put into practice in order to ably situate practices against each other (Habermas 1996). For example, which corporate forms are legal constrains which versions are more readily put into practice and thus exercises some collective control over economic life. Practitioners are also always legal subjects in a modern regime and we put a legal order into practice partly by situating all our other roles as legal subjects. When Amazon workers successfully unionized in New York in April 2022, they used their powers as legal subjects to force Amazon into a version that bargains with workers as a unit. Amazon was aware that such a version was a possibility, so engaged in boundary work to persuade workers to vote against the union, presenting the union as undesirable. United States law mediates this contest over which version of Amazon practitioners would put into practice in virtue of its effective sovereignty claims.

Within any given domestic social practice, then, we must understand the function of legitimacy discourses in relation to legality. Law claims to set the outer limits of permissibility and possibility when practitioners are choosing which version of a social practice to instantiate. In a functioning legal order, that claim will be mostly effective due to both law-abidingness and enforcement. Thus, calling a potential version illegitimate when it is illegal is largely redundant; it may be true but it is mostly beside the point. Illegal versions will be difficult to operate in distinctive ways, including in light of practitioners' (variable) commitment to legality.²¹ So legitimacy discourses in this context are primarily useful to address telically bad but not illegal versions of the practice.²²

²¹ If society is largely constituted by sets of social practices, we can also see law as a way of deciding which version of society to instantiate. Law can declare entire practices illegal, i.e. there is no legal version of those practices, thereby instantiating a version of society where those practices do not exist or exist only in liminal forms.

²² This same embeddedness explains the difference between legitimacy discourses about practices as a whole and legitimacy discourses applied to particular kinds and entities within a practice. In the latter case, the practice as a whole sets some outer limits of practical possibility that needn't be

Of course, law is also a social practice. In a way, it is the social practice of social practices (cf. Rawls 1999, 462). Like all our social practices, practitioners always face an open question of which version to instantiate. But unlike all our other practices, law cannot stand outside this practice as arbiter of the outer limits of socially permissible versions. Law cannot stand outside itself. Law has highly formalized procedures for choosing which version to implement, up to and including a constitution and procedures of constitutional interpretation. There can be illegal versions of law but positive law is precisely changeable by the law-making regime, so illegality is not a very secure critical stance on that regime. Luckily, no social practice can close off practitioners' external discretion, our ability to step outside any practice and ask how we should act.²³ So the legitimacy discourse that we employ to solve the problem of bad boundaries of law has its own distinctive form and function.²⁴ We have discovered the legitimacy discourse that we call political legitimacy.

'Political legitimacy' is a term of art. It is not about all questions of legitimacy in the political domain, although such questions are coherent. Instead, it is about the regime-level claim to the right to rule. Political legitimacy discourses give citizens minimal telic standards for evaluating state rule. Lower-level questions of legitimacy within the regime take the same form as any other domestic social practice because they are also legally regulated. If someone questions Brian Kemp's legitimacy as the governor of Georgia, respondents can appeal to the fact that he won the election; if the question is pressed further, they can appeal to the legality of the election under state and eventually federal law; further again, to the constitutionality of such law. At this point, notoriously, positive law exhausts itself. The only place for this exchange to go is the nature of law and a legal order—perhaps officials' beliefs, or a grundnorm, or an act of popular sovereignty become the object of discussion. That is, where the formalized system of political contestation exhausts itself, telic standards on the ruling regime take over. Political legitimacy is where we turn when we need critical standards for gatekeeping states.

handled by the legitimacy discourse. Essentially, integration into the broader context is internal to the telos of a localized kind.

²³ Choosing how to act all-things-considered is always a matter of choosing how to situate practices and the demands made on us in virtue of our roles. [Reference removed for anonymous review.]

²⁴ Thus the gatekeeping account falls along similar lines to Buchanan's (2018) metacoordination view.

If we take a functional approach to legitimacy discourses, we must take account of practitioners' realistic alternatives. Precisely because they employ non-formalized telic standards, legitimacy discourses are an insecure solution to the problem of bad boundaries; in some ways they are last resorts. This is especially true of political legitimacy because it distinctively lacks legality as an external source of gatekeeping.²⁵ As classical social contract theory emphasized, practitioners also do not have the option of abolishing the political institution because organized community is morally and practically necessary (Christiano 2008, 237). So practitioners employ political legitimacy discourses to assess whether they should treat E as K (their state as having the right to rule) in a very particular context where they need a K and where they lack formalized external gatekeeping mechanisms.

The institutional form of the state further complicates gatekeeping. The sovereign state constitutively claims final and exclusive authority (Morris 1998). The state's monopoly on violence, in Weber's (1947) classic characterization, ensures these claims are effective. Effectively enforced final and exclusive norms guarantee "average norm compliance" (Habermas 1996, 31) and so result in a social order structured by law. Modern legal orders are characterized by successful integration of all the social practices in a territorial jurisdiction and the primacy of individuals as citizens above and beyond their participation in any particular practice.

For our purposes, the upshot is that state legitimacy is the limit case for legitimacy discourses. Attempting to fix the problem of bad boundaries by treating the sovereign state as if it lacks the right to rule is a fraught enterprise. By definition, there are no higher authorities to appeal to and no competitor institutions that are in place to carry out the same functions. As defenders of the status quo are quick to emphasize, correcting the state by disobeying it risks destabilizing the legal order. The alternatives are limited and by definition untested in that context—not to mention arrayed against whatever capacities for violence the state maintains and is willing to wield against challengers.

This does not mean that gatekeeping the state is impossible. It does mean

²⁵ This raises the possibility of international law as an external gatekeeper. Following Cohen (2012), I think this misunderstands the nature of contemporary sovereignty. Just as importantly, though, the gatekeeping account shows the limits of such an appeal. At some point, formalized standards run out; political legitimacy discourses would face the same challenges at the edge of international law as they do at the edge of the state.

that effective internal boundary repair procedures are more important for the state than any other social practice. In other words, the more a state enables citizens to legally contest and change the law, the less gatekeeping there is for a legitimacy discourse to perform (cf. Pettit 2019). It follows that democratic procedures, division of powers, and rule of law all make states more legitimate in the sense that telic failures can be handled without treating the state as if it lacked the right to rule.²⁶ Since these procedures give practitioners modes of effective boundary work within the practice, bad boundaries can be addressed without making the drastic judgment that the appropriate way forward is to treat this K as not a K.

On the gatekeeping account, then, a state is legitimate insofar as its laws should be treated as binding according to minimal telic standards, especially its higher-order laws concerning political contestation and change. We have alighted upon a familiar understanding of political legitimacy.²⁷ My goal in articulating the gatekeeping account is to refine our understanding of state legitimacy, not to overthrow it.²⁸ It is a virtue of the account that it can explain the language of legitimacy in a broad variety of contexts while remaining consistent with theories of legitimacy in the political context, where they have been most carefully considered.

7. Justice

A central question for political philosophy is how political legitimacy and justice relate (Larmore 2020). On a standard, Rawlsian understanding, legitimacy is useful because permanent reasonable disagreement about justice makes justice norms unsuitable as an object of stable agreement. On my view, by contrast, legitimacy norms are tools for implementing any kind of social practice, which take on particular importance in the context of the modern state. In short, even Rawls' well-

²⁶ This relies on some substantive views about the telos of political community. Of course, that's as it should be on my account.

²⁷ The view that states are legitimate if laws should be *treated as* binding is not quite the same as the view that states are legitimate if laws *are* binding. Without going into detail, I think this is why political legitimacy is not identical to the question of justifying the state, if what we mean is showing how its claims are morally justified. Rawls (2005, lv) hints at this but doesn't pursue it.

²⁸ Indeed, the implicit genealogy of my approach explains the broad meaning of legitimacy by appealing to the success of our political legitimacy discourses. My untested hypothesis is that legitimacy discourses migrated from the political domain as processes of modernity gave rise to increasingly functionally differentiated and legalized social domains, which have a greater need for telic gatekeeping.

ordered society would need political legitimacy discourses.

We can start with Rawls' (1999) appeal to the circumstances of justice. This presents norms of justice as solutions to problems that arise under certain conditions, following Hume. Different kinds of norms will be useful as solutions to different problems that arise under different conditions. The problem of justice as Rawls conceives it is contestation over moderately scarce resources. Norms of justice solve this by articulating acceptable public standards for distributing benefits and burdens. If people don't accept justice norms, then they won't accept the public authority's distributional decisions as binding, i.e. they won't reliably abide by law. This society will be deeply unstable because political decisions do not settle social contests. Rawls' worry that any liberal society is fatally unstable is answered by the possibility of the well-ordered society. Eventually, however, Rawls recognizes that norms of justice are insufficient grounds for a stable society because of their deep embeddedness in broader ideologies. Among others, Valentini (2012) narrates Rawls' famous political turn from justice to legitimacy as recognizing the condition of permanent reasonable disagreement about justice (cf. Buchanan and Keohane 2006). This circumstance of legitimacy presents a different problem and so demands a different solution: the liberal principle of legitimacy, public reason, and the overlapping consensus.

For Rawls and many following him, legitimacy norms are weaker justice norms (Rawls 2005, 428). They solve the same problem as justice norms but under different, especially non-ideal, conditions (Freeman 2007; Langvatn 2016). Their function is to stabilize a political community by providing standards for accepting political decisions. They must be less demanding than justice norms to be agreed upon even amidst disagreement about justice but they must also make sense as standards for accepting political decisions. Justice norms are preferable standards—after all, they tell us how people deserve to be treated—but disagreement rules them out as a practicable solution, so we settle for second-best.

The gatekeeping approach makes clear how this story mistakes the nature of legitimacy. As Waldron (1993) argues, justice must be administered, not merely meted out. Justice requires formalized rules of justice: institutionalized administration over a particular jurisdiction with publicly defined statuses and mechanisms for contestation (Rawls 1999, §1; Christiano 2008, ch. 2). This formalized social practice, as all formalized social practices, can only be put into

practice with the help of non-formalized implementation norms. Legitimacy discourses help practitioners implement social practices by solving the problem of bad boundaries.

The problem of bad boundaries is not a problem of disagreement. Consider Rawls' well-ordered society. Insofar as it is a human society, its law is underdetermined and practitioners have a great deal of discretion to implement their legal order in their changing contexts as they see fit. How they do so cannot be fully determined by justice norms since justice must be institutionalized. Formalizing implementation norms simply creates new underdetermined points of discretion. But implementing law via discretion can lead to versions of the practice that drift, including implementing the parts of law that serve as internal correction mechanisms. The possibility of sufficiently bad boundaries cannot be ruled out by stipulating compliance with principles of justice.

Consider civil disobedience as part of our broader theory of political legitimacy. Which injustices justify civil disobedience is underdetermined by the theory of justice because it is partly a strategic question of effective remedy. Citizens need non-formalized political norms to guide illegality as a way of riding herd on the state (Habermas 1985). It is precisely a question of exercising control over the form of our political institutions when formalized mechanisms have run out. As Rawls (1999, 342) puts it, "The final court of appeal is not the court, nor the executive, nor the legislature, but the electorate as a whole. The civilly disobedient appeal in a special way to this body." This appeal must be made to non-formalized telic standards. Appealing to formalized standards in this context means appealing to legalized standards, but making legal space for civil disobedience is conceptually incoherent since disobedience implies illegality in this context. The same dynamic of ineffective internalization haunts legitimacy standards. Formalizing them fundamentally changes them.²⁹

A human society that fully agreed on a theory of justice would still need political legitimacy norms.³⁰ The theory of justice cannot fully specify how to

²⁹ Thus insofar as legitimacy standards are institutionalized, as in measures of democracy, they lose some of their functionality as part of non-formalized gatekeeping tools. In essence, such formalization would shift the relevant gatekeeping to the legitimacy of the measurement and measuring institution.

³⁰ They of course don't need to call these legitimacy norms, nor would they need LEGITIMACY or any linguistic analogue. They need a normative discourse that plays this function.

implement a legal order. Legal orders are social practices that function via the enactment of shared, public norms, so are inevitably underspecified and require discretionary implementation. The function of political legitimacy norms is to ride herd on that practice by providing minimal telic standards for assessing versions and correcting sufficiently bad boundaries. Political legitimacy discourses provide the non-formalized norms that are necessary to implement the formalized administration of justice norms.

In summary, human social practices only exist in virtue of people exercising their discretion about which version of the practice to implement in their practical context. Practices guide this choice in many ways and so result in patterned behavior defined by boundaries. But explicit guidance always runs out and practitioners need ways to coordinate the non-formalized implementation of boundaries. One place such non-formalized guidance is especially important is in the correction of bad boundaries. Bad boundaries are a problem from the internal perspective of the practice but are also resistant to correction by extant procedures. Legitimacy discourses solve the problem of bad boundaries by identifying minimal telic standards that practitioners can employ in immanent, critical boundary work. Political legitimacy discourses enable citizens to ride herd on their states, coordinating corrective measures for when the state pulls away from our collective aims.

References

[Author references removed for anonymous review.]

- Bicchieri, Cristina. (2016) *Norms in the Wild*. Oxford: Oxford University Press.
- Bourdieu, Pierre. (1977) *Outline of a Theory of Practice*, trans. Richard Nice. Cambridge: Cambridge University Press.
- Brennan, Geoffrey, Lina Eriksson, Robert E. Goodin, and Nicholas Southwood. (2013) *Explaining Norms*. Oxford: Oxford University Press.
- Buchanan, Allen. (2018) 'Institutional Legitimacy'. In David Sobel, Peter Vallentyne, and Steven Wall (ed.), *Oxford Studies in Political Philosophy*, vol. 4 (Oxford: Oxford University Press), 53-80.

- Buchanan, Allen and Robert O. Keohane. (2006) 'The legitimacy of global governance institutions'. *Ethics & International Affairs* 20/4, 405-37
- Chalmers, David J. (2011) 'Verbal Disputes'. *Philosophical Review* 120/4, 515-66.
- Christiano, Thomas. (2008) *The Constitution of Equality*. Oxford: Oxford University Press.
- Cohen, Jean L. (2012) *Globalization and Sovereignty*. Cambridge: Cambridge University Press.
- Dworkin, Ronald. (1986) *Law's Empire*. Cambridge, MA: Harvard University Press.
- Dworkin, Ronald. (2011) *Justice for Hedgehogs*. Cambridge, MA: Harvard University Press.
- Freeman, Samuel. (2007) *Rawls*. New York: Routledge.
- Geach, P. T. (1956) 'Good and Evil'. *Analysis* 17/2, 33-42.
- Gieryn, Thomas F. (1983) 'Boundary-Work and the Demarcation of Science from Non-Science: Strains and Interests in Professional Ideologies of Scientists'. *American Sociological Review* 48/6, 781-795.
- Habermas, Jürgen. (1996) *Between Facts and Norms* 2nd ed., trans. William Rehg. Cambridge, MA: The MIT Press.
- Habermas, Jürgen. (1985) 'Civil disobedience: litmus test for the democratic constitutional state'. *Berkeley Journal of Sociology*, 30, 95-116.
- Hohfeld, Wesley Newcomb. (1919) *Fundamental Legal Conceptions as Applied to Judicial Reasoning*. New Haven, CT: Yale University Press.
- Lamont, Michèle and Virág Molnár. (2002) 'The Study of Boundaries in the Social Sciences'. *Annual Review of Sociology* 28, 167-95.
- Larmore, Charles. (2020) *What Is Political Philosophy?* Princeton, NJ: Princeton University Press.
- Langvatn, Silje. (2016) 'Legitimate, but unjust; just, but illegitimate: Rawls on political legitimacy'. *Philosophy and Social Criticism* 42/2, 132-153.
- "legitime, adj. and n." (March 2022) *OED Online*, Oxford University Press, www.oed.com/view/Entry/107120. Last accessed 11 April 2022.
- March, James G. and Johan P. Olsen. (2011) 'The Logic of Appropriateness'. In Robert E. Goodin (ed.) *The Oxford Handbook of Political Science* (Oxford: Oxford University Press), 478-497.
- Miller, Seamus. (2010) *The Moral Foundations of Social Institutions*. Cambridge:

- Cambridge University Press.
- Morris, Christopher W. (1998) *An Essay on the Modern State*. Cambridge: Cambridge University Press.
- North, Douglas C. (1991) 'Institutions'. *The Journal of Economic Perspectives* 5/1, 97-112.
- Pettit, Philip. (2019) 'The Control Theory of Legitimacy'. In Wojciech Sadurski, Michael Sevel, and Kevin Walton (ed.) *Legitimacy* (Oxford: Oxford University Press), 7-31.
- Rawls, John. (1999) *A Theory of Justice* rev. ed. Cambridge, MA: Belknap Press.
- Rawls, John. (2005) *Political Liberalism*, exp. ed. New York: Cambridge University Press.
- Ritchie, Katherine. (2020) 'Social Structures and the Ontology of Social Groups'. *Philosophy and Phenomenological Research* 100/2, 402-24.
- Sangiovanni, Andrea. (2019) 'Debating the EU's Raison d'Étre: On the Relation between Legitimacy and Justice'. *Journal of Common Market Studies* 57/1, 13-27.
- Schauer, Frederick. (1991) *Playing by the Rules*. Oxford: Clarendon Press.
- Searle, John. (2010) *Making the Social World*. Oxford: Oxford University Press.
- Sibley, Frank. (2001) *Approach to Aesthetics: Collected Papers on Philosophical Aesthetics*. John Benson, Betty Redfern, and Jeremy Roxbeer Cox (ed.). Oxford: Oxford University Press.
- Thomson, Judith Jarvis. (2008) *Normativity*. Chicago, IL: Open Court.
- Tuomela, Raimo. (2002) *The Philosophy of Social Practices*. Cambridge: Cambridge University Press.
- Valentini, Laura. (2012) 'Assessing the global order: justice, legitimacy, or political justice?' *Critical Review of International Social and Political Philosophy* 15/5, 593-612.
- Waldron, Jeremy. (1993) 'Special Ties and Natural Duties'. *Philosophy & Public Affairs* 22/1, 3-30.
- Weber, Max. (1947) *The Theory of Social and Economic Organization*, trans. A. M. Henderson and Talcott Parsons. Oxford: Oxford University Press.
- Williams, Bernard. (2009) *In the Beginning was the Deed*. Princeton, NJ: Princeton University Press.