

What Are Social Norms?

Franz Dietrich
Paris School of Economics & CNRS

Kai Spiekermann
London School of Economics

September 2024

Abstract: While the importance of social norms for shaping and transforming communities is uncontested, their nature and normativity are controversial. Most recent theorists take social norms to arise if members hold certain attitudes, such as expectations on others, perhaps along with certain behaviours. Yet attitudes do not create norms, let alone social norms or social normativity. Social norms are instead *made*: through a social process. Social norming processes are special communication processes, often non-verbal and informal. We present different versions of a process-based account of social norms and social normativity. The process-based view brings social norms closer to legal norms, as processes represent social ‘acts’, just as laws and contracts arise through acts rather than mere attitudes, for instances acts of voting or signing.

Keywords: social norms; normativity

1. Introduction

Social norms are key elements in our inventory of the social world. They are relevant as informal institutions that convey, entrench and sometimes transform a society’s central tenets. Philosophers have increasingly engaged with social norms: their nature, function and effects. In this paper, we propose a new and arguably more convincing account of social norms.

Understanding social norms is philosophically rewarding, but is also of significant practical importance. For instance, social norms help explain why injustice persists – they provide structural explanations that illuminate why the status quo can be sticky, and how it can be overcome. Since social norms underpin much of human behaviour, understanding and ultimately changing them will be key in the pursuit of justice.

A successful account of social norms must address two central questions:

1. What *are* social norms?
2. When, why and how are social norms *normative*?

Regarding the first question, recent philosophical accounts of social norms focus on the attitudes underpinning such norms. Some zoom in on beliefs, others on expectations, preferences, or commitments. Against the mainstream, we will reject this focus on attitudes. Instead, social norms are requirements generated by a social norming *process*. Attitudes, such as expectations and beliefs about others, are not the origin of social norms, but typical consequences of them. This inverts the direction of explanation between social norms and attitudes.

Regarding the second question, we will argue that social normativity is a distinct type of normativity, which is reducible neither to moral reasons nor to rational reasons created by sanctions against norm-violation. Social normativity is once again grounded in the social norming process, not in attitudes.

This is a programmatic paper aiming to intervene in a debate that, we think, could benefit from re-orientation. Our argument comes in different steps. After discussing some examples and existing philosophical accounts of social norms (Sections 2 and 3), we develop our process-based approach to social norms (Sections 4, 5, 6 and 7) and to social normativity (Sections 8 and 9), before concluding (Section 10).

2. Social Norms in Action

On a first pass, social norms are informal social requirements to behave or be in certain ways. For instance, in many societies it is socially required to participate in elections, and in many workplaces it is socially required to greet colleagues in the morning, and to avoid offensive language. Many approve of these requirements: citizens expect others to vote, and work colleagues expect others to give greetings and talk inoffensively.

Some social norms are created in explicit ways. Imagine a well-attended local neighbourhood meeting. At the meeting, one of the neighbours pronounces that ‘it has become good practice in this neighbourhood to keep the noise down after 10pm’. Everybody nods and sees everybody else nodding. In subsequent discussions, neighbours make regular reference to this moment. This is the birth or re-endorsement of a social norm in the neighbourhood, proscribing noisy activities after 10pm. What created (or re-endorsed) the norm is the process of agreement.

Many social norms are created through more implicit processes. People might express their agreement implicitly by following a social practice, while showing public disapproval when someone falls out of line. For instance, after having met in the ‘Black Eagle’ pub every Tuesday for the last three months to play darts, the regulars start criticising those who fail to show up without sending apologies. This process of consistently gathering and calling out digressions soon creates a social norm requiring attendance or apologies.

Social norms can emerge in radically implicit ways, as people express approval indirectly through sanctioning violations publicly. On London Underground escalators, one is supposed to ‘stand on

the right, walk on the left'. Transgressions are publicly sanctioned when locals make disapproving 'tutting' sounds, clear their way with a passive-aggressive 'excuse me!', or stand very close to the person impeding the way. This mixed sanctioning activity expresses the majority's view about how to use escalators, thereby creating or maintaining the norm.

3. Recent Theories of Social Norms

The examples above suggest that social norms are created or maintained by processes of various forms, later to be called 'social norming processes'. We now sketch some recent philosophical accounts of social norms. Strikingly, these accounts are not based on processes, but on attitudes.

The literature focuses largely on when social norms exist, not what they are. Occasionally one finds the surprising claim that social norms *are* (rather than exist under) certain combinations of attitudes. One might reinterpret this apparent ontological claim as a shorthand for expressing existence conditions within a largely non-ontological analysis. We shall later take an explicit position on what social norms are, not just on when they emerge.

Some accounts are preference- or desire-based. According to Cristina Bicchieri's (2006), social norms emerge under a particular constellation of empirical or normative expectations, conditional preferences for conformity, and special beliefs. Her precise 'Conditions for a Social Norm to Exist' (p. 11) single out social norms from the wider class of behavioural rules.

Another growing branch of the literature suggests that social 'norms are clusters of normative attitudes' (Brennan et al. 2013, 29, emphasis omitted). This is an example of an apparent ontological claim that might be reinterpreted non-literally, as a claim about when social norms exist, not what they are. This reading is supported when Brennan et al. turn explicitly to existence conditions:¹

'[a] normative principle P is a [social] norm within a group G if and only if:

A significant proportion of the members of G have P-corresponding normative attitudes; and

A significant proportion of the members of G know that a significant proportion of the members of G have P-corresponding normative attitudes.' (Brennan et al. 2013, 29)

In stating attitude-based existence conditions, they step into the Hartian tradition, which puts attitudes centre stage (Hart 1961).

Andrei Marmor (2023) and Margaret Gilbert (1999) also start from HLA Hart's (1961) classical discussion of 'social rules'. For Marmor (2023), a 'social rule' (roughly equivalent to a social

¹ Brennan et al. focus on a wider class of 'norms' in that definition, but our focus here is on social norms.

norm) exists under broadly the following conditions: (i) conformity with the content of the rule, and (ii) common knowledge that the group collectively intends the content, and that this fact provides normative reasons for compliance and enforcement (Marmor 2023, 53). This account of social rules (or norms) emphasises a practice and attitudes of two types, beliefs and collective intentions.

Margaret Gilbert's account also emphasizes joint or collective attitudes. For her, a social norm exists if, roughly, there is a joint commitment to accept the relevant requirement as a body (Gilbert 1999, 163). Her account superficially resembles Marmor's, but he ultimately reduces collective intentions to individual intentions, a move Gilbert rejects.

Raimo Tuomela (2007) presents another hybrid set of existence conditions: a social 'ought-to-do norm' exists if and only if there is collective acceptance of the normative demand, a practice, compliance for the right reasons, certain mutual beliefs, and some social pressure against deviation. This again requires certain attitudes and a practice.

Laura Valentini (2021), finally, suggests that social norms are 'requirements accepted as binding by a large enough number of people in a given context' (p. 386). In short, enough individuals robustly intend certain standards of behaviour to be binding. While once again explaining social norms by attitudes, these attitudes are interestingly different. Later, her 2023 book (Valentini 2023) avoids claims about what social norms *are* (p. 23), whilst stating 'social existence conditions' that require individual commitments to the norm and corresponding beliefs about the norm (p. 21).

Notwithstanding important differences between authors, a common theme emerges: the literature tends to tie social norms to attitudes, possibly jointly with practices.

4. The Process-Based Approach to Social Norms

The accounts of social norms sketched in Section 3 might be regarded as versions of a general approach:

Attitude-based Theory: A social norm is a requirement that is suitably generated by attitudes (or attitudes and practices).

To us, social norms are instead generated by social processes. Attitudes such as normative expectations do not create social norms, but are instead often created by social norms. No doubt, attitudes such as expectations can create a *social pressure* to behave or be in certain ways. Yet a pressure is not a norm: it is not a requirement, but a different influence or 'force' on people. An attitude-based social pressure can *co-exist* with a social norm, often as a sanction mechanism enforcing the norm. Some attitude-based accounts might be reinterpreted as accounts of *social pressure* – but this is not our focus. We propose:

Process-based Theory: A social norm is a requirement that is suitably generated by a social process. (The process is then called a social norming process or SNP.)

Note our basic ontological assumption that social norms are *requirements*, under both general approaches as defined here. This rendition of both approaches is thus ontologically explicit, whereas previous attitude-based accounts often avoid ontological commitments.

By a *requirement* we mean a requirement on people. It has a *content*: the thing being required. It typically concerns behaviour, and more rarely thoughts or attitudes. It can usually be expressed in the form ‘do (think, intend, ...) *X* in circumstances *Y*’, for instance ‘wear a tie at formal dinners’ or, in our earlier examples, ‘keep noise down after 10pm’, ‘show up at pub meetings or send apologies’, and ‘stand on the right on escalators’.

Some clarifications about the Process-based Theory are due.

First, a *social process* is an interactive causal process in a given group. Social processes exist in abundance. Many generate no requirements – gossip at a dinner party, for example, is a social process that normally generates no requirement.

Second, our notion of *generating a requirement* is broad. For instance, requirements can be generated in structured or unstructured ways, based on wide or narrow support, where support could be expressed implicitly, through behaviours or through doing nothing. A social process might even generate multiple conflicting requirements: imagine a tumultuous interaction in an overcrowded staircase where some shout out ‘Walk left!’, others ‘Walk right!’, yet others ‘Stand still!’.

Third, while ‘generating requirements’ is a broad notion, ‘generating requirements *suitably*’ is much narrower. Many requirements are generated non-suitably by a social process – then they are not social norms, and the process is not a SNP. The three requirements generated by the chaotic process in the overcrowded staircase are all generated non-suitably: they are not social norms. For another example, a social process in which citizens elect a dictator who then orders everyone to dress in blue has generated a requirement (‘one should dress in blue’), but the generation is non-suitable for a social norm. What exactly goes wrong in these examples? This is precisely what the following sections will aim to clarify.

Fourth, in what *metaphysical* sense does a process ‘generate’ a requirement? On a causal reading, the process *causes* the requirement. This reading, however, must be rejected. For one, causation normally happens between objects of same type, not between processes and requirements.² For

² For sure, the fact that some group members express (that they) want the requirement can be caused by earlier events, as part of the social process. But the requirement itself is arguably not caused by the process.

another, the idea that a (causal) process between people would, as a whole, cause a requirement seems like a misguided mix between two levels of causation. On our view, the process generates the social norm in the metaphysical *grounding* sense: the social norm exists *because* of the process (on grounding, see Bliss and Trogon 2024).

Fifth, do attitudes play any role at all? They can only play an indirect role, by driving the SNP, or instead resulting from the SNP or from the social norm. But the norm is grounded in the SNP, not in attitudes.

Finally, to set subtleties aside, let us henceforth think of a SNP as generating *one* social norm, and of a social norm as being generated by *one* SNP – so that a social norm has ‘its’ SNP and a SNP has ‘its’ social norm. One can always achieve this one-to-one correspondence by suitably individuating the SNP or the norm.³

We will now work out the process-based approach in more detail. It will provide for a parsimonious yet explanatorily powerful account of social norms, and a more compelling metaphysics of social norms and normativity. Let us anticipate five aspects. First, the approach blocks greedy reduction to attitudes when *explaining* norms. Second, it avoids identification with attitudes when *defining* norms – a category mistake already avoided by some existing attitude-based accounts and our rendition of the Attitude-based Theory above. Third, it can explain how social norms can exist without any corresponding attitudes. Fourth, our approach benefits from a helpful analogy with legal norms. Legislation is created by a suitable *legal* process, social norms by a suitable *social* process. We will draw on this parallel, by comparing legal and social norming processes. Finally, our approach will later lead us to a more compelling account of social *normativity*.

5. A Baseline Account

We now present a process-based account of social norms: a concrete version of the Process-based Theory. Variants of the account are discussed in Section 6. Like many attitude-based accounts, our account and its variants are given by a condition for when a social norm is generated. The condition will however require a certain social process, not certain attitudes. It will effectively define what a SNP looks like. Like for an attitude-based account, its precise statement is bound to be controversial – which is why Section 6 will offer alternative statements.

³ A social norm generated independently by different SNPs can be regarded as generated by the single super-process comprising those processes. Conversely, a SNP generating different social norms (e.g., norms about addressing insiders and about addressing outsiders) can be regarded as generating the single hybrid social norm given by the conjunction of those norms; alternatively, one can re-individuate the SNP more narrowly so that it generates only one non-hybrid norm.

Consider a grouping (such as a work community, family, nation, or even civilisation) and a suitably general proposition C about behaviour, thoughts or attitudes of persons (such as ‘we wear a tie at work’ or ‘one holds one’s parents in high regard’). On a first pass, C is the content of a social norm *if and only if and because* the following holds:

Informal Condition: Support for C is communicated sufficiently within the grouping.

This condition is deliberately vague, leaving open what ‘support’ and ‘communicated’ means. How should the condition be made precise? Here is our baseline rendition of the condition, with other renditions to come in Section 6:

Communication (COM): Enough members communicate to enough members that they want that C is obligatory.

Several clarifications are due. First, we start from a *grouping*. All groups are groupings, but groupings can have a less well-developed identity and demarcation. Nations, ethnicities and religions might qualify as groups, whereas the users of escalators in the London Underground form only a loose grouping. Social norms can already arise in loose groupings that change permanently.

Second, by ‘obligatory’ we mean ‘*normatively* required’. Analogously, ‘obligation’ stands for ‘*normative* requirement’. Members express that they want an obligation, not just a (possibly non-binding) requirement. Wanting this obligation differs from wanting the social norm, in two ways. In one sense, it is stronger, as obligations are normative by definition. In another sense, it is weaker, because a social norm, if normative, is a very special obligation, namely one grounded in a SNP. Yet members need not (and usually will not) express anything about the grounds of the obligation: they express that they want the obligation, period.

Third, we want COM to be read as: For all sets E and E' of enough⁴ members, each person in E communicates to each person in E' that she wants that C is obligatory.⁵ By ‘ A communicates X to B ’ we mean ‘ A expresses X to (at least) B and B perceives that A expresses X ’ (where X could be

⁴ We presuppose a notion of “enough”. Containing enough members could mean containing a majority of the members. (In principle, the notion of “enough” could even be a different one for expressing members and for perceiving members; we set this possibility aside.)

⁵ Rather than as: For all sets E of enough members, each member A in E communicates to each member of some set E_A of enough members that she wants that C is obligatory. This unintended reading is weaker, as E_A could depend on A . For instance, assume the grouping contains just three persons 1, 2 and 3, and “enough members” means “at least two members”. If 1 and 2 express, where 1’s expression is perceived by 1 and 2, and 2’s by 2 and 3, then COM holds only under the unintended reading, as too few members (i.e., only person 2) perceives both expressions. If 1’s and 2’s expressions are instead perceived by the same majority E' of members, then COM genuinely holds.

the will that *C* is obligatory). So, communication has a sending and a receiving aspect. According to COM, a will is widely expressed, which is then widely perceived. But what do we mean with ‘expressing’ and ‘perceiving’?

Expressing: We interpret ‘expressing’ broadly, as covering explicit and implicit expressions. One can express a will through speech, behaviours (single or repeated), facial expressions, and sometimes even silence. Typically, support for *C* gets expressed when the opportunity arises, often through acts of compliance with *C*, approval, or sanctioning of non-compliance. By contrast, it is hard to (credibly) express support while violating *C*.

Perceiving: We interpret ‘perceiving *X*’ as ‘coming to believe *X* as a result of *X*’. For instance, one perceives that it rains if the rain makes one realise that it rains.⁶ Perceiving *X* implies believing *X*. Why do we require a widespread *perception* rather than merely *belief* that the will is expressed? After all, attitude-based accounts of social norms often require a widespread *belief* that the relevant attitude (e.g., a normative expectation) is widely held.⁷ We would thus diverge less from some of the literature if we weakened COM such that the expressions of will are widely *believed* to happen rather than perceived. Yet the problem is that mere beliefs could exist by coincidence. Just imagine no one perceived any expressions of will, but everyone (unrelatedly) dreams that everyone expressed this will – and then forgets that it was ‘just a dream’, thereby coming to believe that the expressions really happened. In that case, the belief-based variant of COM holds (unlike COM itself). But arguably that process is flawed: it lacks interactions and is thus not a *social*. So, no social norm emerges.⁸

An interesting parallel arises with particular attitude-based accounts that merely require beliefs about certain attitudes (such as normative expectations), not these attitudes themselves. Such accounts typically make room for pluralistic ignorance, where widespread *false* beliefs about attitudes of others are supposed to ground a social norm.⁹ In our account, pluralistic ignorance can occur via a different route: Individuals need not actually have the will they express: they could

⁶ This mental notion of perception goes beyond mere sensory experience. Perception in our sense implies true belief: if someone perceives a proposition *X*, then *X* is true and she believes *X*.

⁷ Bicchieri (2006) demands such beliefs in her ‘normative expectation’ condition.

⁸ This objection resembles our objection against attitude-based accounts, which too are not properly ‘social’, since the communication element is missing.

⁹ In our opinion, such attitude-based view – based on attitude-beliefs rather than attitudes – confuse the question of whether social norms exist with the question of whether they are believed to exist. See also Valentini (2023), p. 32-3; and Brennan et al. (2013), p. 35. Of course, apparent and real social norms can have the same powerful consequences.

‘fake’ this will, consciously or subconsciously. If people perceive these expressions of will, and falsely infer that the wills exist, then COM holds and a norm emerges – under pluralistic ignorance.

The analogy to legal processes is informative in this regard: social norms and laws both come into existence once the right external ‘protocol’ has been followed, regardless of the ‘true’ wills. Indeed, laws arise once the legislators externally express their consent by voting, even if they secretly disagree. By allowing that people do not possess the attitudes they express, our process-based account again departs from standard attitude-based accounts.

Revisiting our examples, we can now see condition COM at work: The neighbours create a norm when enough of them communicate to enough others that they want certain noise-limiting behaviour to be obligatory. The dart players create a social norm if enough of them communicate to enough others that they want regular attendance to be obligatory. And the users of the escalators create or reinforce a social norm by repeated, public social disapproval of non-compliance, such that enough users communicate to enough others that they want standing on the right to be obligatory.

6. Variations of the account

Condition COM provides a particularly simple process-based account of social norms. Is the account satisfactory? Let us put up for debate some variations of the account, either strengthening COM or modifying COM by varying what exactly is being communicated, expressed or perceived. This will illustrate how the process-based theory could be refined. All variants represent ways to flesh out the earlier ‘Informal Condition’.

1. Truthful or credible communication. Suppose the local mafia boss orders that all front doors be painted in blue, his favourite colour. Fearing repercussions, nearly everyone complies and shows various signs of approval, thereby publically expressing the relevant will. This is widely perceived in the community. So COM holds. But is there really a social norm to paint doors blue? Two things are peculiar. First, the expressions of will are not truthful: members *pretend* to want this, out of fear. Second, those who perceive these expressions do presumably not believe what is being expressed: they realise that the others only pretend.

The account of social norms can be refined in two different ways to respond to the two concerns. We say that someone *communicates* X (e.g., a will) *truthfully* to someone else if she communicates X to him and X indeed holds.¹⁰ She *communicates* X *credibly* to him if she communicates X to him and he then believes X . The two refined conditions are:

¹⁰ Equivalently: she expresses X truthfully and he perceives the expression. Here, *expressing* X *truthfully* means expressing X where X indeed holds.

Truthful Communication (t-COM): Enough members communicate to enough members truthfully that they want that *C* is obligatory.

Credible Communication (c-COM): Enough members communicate to enough members credibly that they want that *C* is obligatory.

On the t-COM-based account, no ‘door colour norm’ exists in our mafia example, as the widely communicated will is fake (except for the boss). On the c-COM-based account, the norm again fails to exist, this time because people fail to believe that the widely communicated wills exist (except for the boss). While t-COM and c-COM are possible alternatives to COM, one might insist that truthfulness and credibility are not essential, though often present. Why? In a realistic version of the mafia story, members won’t manage to fake the will all the way: they will send mixed signals, including support (say, by complying with the order) and disapproval (say, by their facial expression or lack of enthusiasm for blue doors), overall *not* expressing the will. In consequence, COM fails, and *this* blocks a social norm. But if members do fake support consistently and these expressions of support are widely perceived, nobody questions their truthfulness, so that COM holds, and then the norm does arguably emerge. For a legal analogy, note that a contract comes to exist once parties express their will or agreement; speculations about ‘secret’ dispositions are irrelevant. Since, in our view, social norms are grounded in certain acts or processes (not secret attitudes or thoughts) – in the same way as contracts – it is a plausible implication of our account that social norms can emerge when the process succeeds – even if the attitudes involved are skillfully feigned.

2. *Communicating agreement.* Taking the analogy to contracts and laws further, some might argue that social norms arise through communication of agreement, not of will. They would replace COM with:

COM*: Enough members communicate to enough other members that they agree to *C* being obligatory.

Will and agreement are often linked, but agreement is more cognitive while will is more desire-based. One can want something without agreeing to it, and arguably also vice versa.

Expressions of agreement can again be highly implicit – and more than they would have to be for creating contracts or laws. Being silent, complying, or nodding (as in our neighbourhood meeting example) can all express agreement, not just will.

3. *Communicating basic attitudes.* In the accounts above, members do not communicate that they want (or agree to) *C* simpliciter, but that they want (or agree to) *an obligation* of *C*. So, they communicate an attitude about an obligation of *C* rather than *C* itself – an ‘indirect’ rather than ‘basic’ attitude about *C*, as we shall say. By contrast, many attitude-based accounts of social norms tie social norms to basic attitudes, such as normative expectations, demands, or preferences. For

many such accounts, one could consider a corresponding process-based account in which *this* basic attitude is communicated:

COM': Enough members communicate to enough members that they expect/demand/prefer/etc. *C*.¹¹

We are skeptical about a COM'-based account. Communicating mere approval of some behaviour without any obligation seems insufficient for a social norm. If everyone wants everyone to dress properly, and communicates this, but everyone also rejects an obligation to do so, then it seems implausible that a social norm was created. For comparison, to create a law (or contract), the legislators (or parties) must approve the law (or contract), not just what it requires. While COM* moves social norms closer to laws and contracts, COM' moves social norms further away from the legal sphere, because expectations (demands, preferences, etc.) do not ground laws or contracts.

4. Higher-order communication. Proponents of an attitude-based approach to social norms often do not stop with requiring that enough members (say) normatively expect some behaviour: the expectations must be widely believed to exist (first-order beliefs), widely believed to be widely believed to exist (second-order beliefs), etc.¹² COM only guarantees the existence of first-order beliefs: the expressions of will must be widely believed to exist. In fact, COM guarantees something stronger: the expressions are widely *perceived* (hence believed to exist). We have explained earlier why we require perceptions rather than mere beliefs. For analogous reasons, when going higher order, requiring higher-order perceptions is more adequate than requiring higher-order beliefs. We will talk of “higher-order communication”. We define communication of order 0, order 1, order 2, etc., as follows:

COM₀: Enough members express that they want that *C* is obligatory.

COM₁: Enough members perceive that enough members express this.

COM₂: Enough members perceive that enough members perceive that enough members express this.

...

Two facts stand out:

¹¹ COM' can be restated more precisely as a condition schema in which one can plug in any type of attitude *A*: *Enough members communicate to enough other members the attitude A towards C*.

¹² See Brennan et al. (2013), p. 31 for discussion.

Fact 1: For any order $k \geq 1$, COM_k implies $\text{COM}_0, \dots, \text{COM}_{k-1}$, under a non-triviality assumption on the notion of “enough” (namely that “no members” does not count as “enough members”).

Fact 2: COM_1 is equivalent to COM , under a plausible assumption on the logic of perception (namely that consequences and finite conjunctions of perceived propositions are perceived).¹³

Why does Fact 1 hold? COM_k implies COM_{k-1} because if enough members perceive something, say X , then *someone* perceives X (by the assumption about the notion of “enough”), which implies that X holds (following our notion of perception, defined earlier). Fact 2 holds by a more complex argument.¹⁴

By Fact 1, an account of social norms based on communication up to a certain order k (e.g., up to order 2) can be defined through a single unified condition, namely COM_k , which automatically subsumes $\text{COM}_0, \dots, \text{COM}_{k-1}$.

We leave open whether social norms require higher-order communication, and if so whether one needs communication up to a certain order k or of all orders k . If COM_k holds for all orders k , one might talk of ‘quasi-common perception’, in analogy with the notion of ‘common knowledge’ in logic. The qualification ‘quasi-’ reflects that the conditions COM_k are stated with ‘enough members’ rather than ‘all members’. Full-blown *common perception* holds if each condition holds in the stronger sense with ‘all members’.

¹³ This allows us to translate between perceiving enough expressions and perceiving that enough members express. For instance, perceiving that Ann expresses and that Indra expresses is equivalent to perceiving that Ann and Indra each express. Technically, our assumption requires that the set of perceived propositions is closed under taking consequences and finite conjunctions, i.e., forms a *filter*. This parallels the standard assumption that the set of *believed* propositions forms a filter.

¹⁴ First, if the notion of “enough” is degenerate in the sense that “no members” counts as “enough members”, then COM_1 and COM both hold trivially, hence are equivalent. Now let “enough members” exclude “no members”. COM says this: There exist sets E and E' of enough members such that, for all A in E and B in E' , A expresses and B perceives that A expresses. The clause after “such that” is equivalent to: all A in E express and for all B in E' it is the case that [for all A in E , B perceives that A expresses]. The clause in “[...]” is in turn equivalent to “ B perceives that all A in E express”, using the assumption on the logic of perception (cf. footnote 13). So, the clause after “such that” becomes equivalent to: all A in E express and all B in E' perceive that all A in E express. This is in turn equivalent to: all B in E' perceive that all A in E express. Here we could skip “all A in E express” because at least someone perceives that all A in E express (since E' is non-empty), which implies that all A in E express (since whatever is perceived is the case, as assumed earlier). In sum, COM is equivalent to: There are sets E and E' of enough members such that all B in E' perceive that all A in E express. In short, enough members perceive that enough members express. This is precisely what COM_1 says.

7. Social Norming Failures

Where needed, the remainder of this paper will assume the baseline account given by COM. Many social processes aim to generate a social norm, but fail, as COM is violated. COM can fail because of expression failures, or perception failures, or both. Examples will help.

Expression Failures. Expression failures happen, first, if too few members express approval, i.e., COM₀ fails. For instance, a lone supporter of orderly queuing is insufficient to instill a queuing social norm. Similarly, in the mafia example, members other than the boss will likely send mixed signals, overall failing to express the will.¹⁵ Second, expression fails if it is not public enough. No norm arises if everyone expresses a will *to herself* (even if this is covertly streamed on the internet, leading to wide perception). Similarly, if a social network is divided into many cliques where users only contact members of their own clique, then social norms cannot arise because wills are not expressed publicly across cliques – at most, social norms *within* cliques could emerge.

Perception Failures. Modifying the social-network example, imagine a different clique effect: members rarely see messages or posts from people outside their clique, be it because they chose certain settings or because the network provider promotes intra-clique communication. Then COM fails because expressions of will are not widely perceived. The same can happen if prejudices prevent members from perceiving wills of certain type or wills expressed by certain members.

Often, expression and perception failures combine. For example, 80 years ago, a social norm demanded short hair for men in Western Europe. Today, the SNPs that used to create and maintain this norm have largely disappeared, as the will is being insufficiently expressed (compliance declined) and insufficiently perceived (people increasingly ignore expressions). The norm has vanished.

8. The Normativity of Social Norms

A philosophically fascinating aspect of social norms is their role in creating normative reasons. The remarkable human ability to manufacture normativity on the fly is worthy of reflection: ‘we can create new normative truths merely by introducing, or getting some people to accept, some rules’, Parfit (2011) observes. Helpfully, we are already familiar with the creation of normativity in the legal realm. Legal philosophers continue to debate under which conditions legal normativity obtains – but at least the terms of that debate are now fairly well established.¹⁶ By contrast, the

¹⁵ Even if the will is initially expressed due to successful threats, so that COM holds, the norm might collapse later as members start sending opposite signals, so that COM fails.

¹⁶ See, for example: Hart (1961); Raz (1979); Enoch (2011); Kaplan (2017); Plunkett, Shapiro, and Toh (2019); Diamond (2024).

normativity of social norms is a recent field of investigation, with foundational questions still unanswered.

We are after a distinct type of normativity: *social* normativity, which differs from both moral and rational normativity. While moral norms are normative by default, the normativity of social (and legal) norms is not immediately evident. To distinguish social from moral and from rational or prudential normativity, we talk of ‘*social* obligations’, ‘*socially* permissible’, ‘*social* reasons’, etc., sometimes replacing ‘social(ly)’ with ‘social-normative(ly)’ for clarity.

We consider a social norm requiring *C*. There might exist a concurrent moral norm requiring the same *C* – but that moral norm gives *moral* reasons for *C*, which we set aside here.¹⁷

How about *rational* or *prudential* reasons for *C*? Social norms are often accompanied by a sanction or incentive mechanism (e.g., Ullmann-Margalit 1977; Schotter 1981). For instance, non-compliant agents could be excluded from the group (e.g., Fehr and Fischbacher 2004) or lose esteem (Brennan and Pettit 2000). Though intertwined with social norms and undoubtedly important for their enforcement, such mechanisms give only *rational* reasons for *C*, often driven by self-interest (Elster 1989). They create no *social* normativity.

Having avoided a confusion with moral or rational normativity, our question becomes more pressing: does social normativity actually exist, and under what conditions? Here again, the standard approach would tie social normativity to attitudes, possibly in combination with practices, whereas we will ground it in processes – the same contrast that arose between our and standard approaches to social norms.

Before developing our process-based approach to social normativity, let us briefly review attitude-based accounts.

For Brennan et al. (2013), if we interpret them correctly, social normativity comes ‘for free’: the normative attitudes underlying a social norm also make the norm normative.¹⁸ We will later argue that normative attitudes alone give no social reason to comply, especially to people not sharing those attitudes. For example, a majority’s attitudes in Italy about when to drink cappuccino

¹⁷ See Southwood (2011) for a related discussion of how to distinguish conventional from moral norms.

¹⁸ While Brennan et al. (2013) seem to have an initially simple view of the relation between normative attitudes and normativity (p. 28-9), their view has additional components that we set aside for simplicity: First, the normative attitudes constituting social norms in their account are, in turn, grounded on social practices, and second, there are expressive reasons to ‘*honour* the practice from the inside’ (p. 80).

arguably has no normative force: non-conformists can sip their cappuccino anytime and have no reason to do otherwise.¹⁹

Proponents of grounding normativity in attitudes might respond that we have reason to fulfil desires, satisfy preferences, or meet expectations. Such reasons can indeed exist – but they are reasons of the wrong kind, reflecting morality or duties of beneficence rather than *social* normativity. Our neighbour’s tastes or preferences might give us a (pro tanto) reason to paint our front door white. But this reason is not social-normative: it might reflect a duty of beneficence, or a consequentialist ethics combined with a preference-satisfaction account of goodness of consequences.

Margaret Gilbert (1999), Raimo Tuomela (2007), Olle Blomberg (2023), and Andrei Marmor (2023) all argue that *collective* attitudes tend to have normative force where individualist attitudes fail: they give a social-normative reason to comply. We, however, doubt that collective attitudes often exist, and that they bind individuals who do not contribute to them, such as non-conformist group members with opposed individual attitudes, or outsiders interacting with the grouping.

Laura Valentini proposes a different source of normativity (Valentini 2021, 2023). For her, social norms originate at least partly in individual *commitments*, and one ought to *respect* individual commitments by complying. We agree that there can be reason to respect individual commitments of others. Yet, to us, these are moral, not social-normative reasons.

9. The Process-based Approach to Social Normativity

To start with, is there such a thing as social normativity? We will argue that some but not all social norms are socially normative. Clearly, a social norm that lacks normativity fails to fulfill the expressed will of members, namely that *C* be obligatory, i.e., *normatively* required. Making room for an absence of normativity is important: individuals who aim for a *normative* requirement sometimes only partially succeed, by getting a social norm but no normativity.

What makes a social norm socially normative? Like the social norm, social normativity (if present) is generated by the SNP: the norm is normative *because* of the underlying SNP, in a grounding sense of ‘because’. In other words, the SNP *gives reason* for *C*, or simply is *reason-giving*. We said “the SNP”, but, more precisely, a SNP can be individuated in richer or narrower ways. Grounding social normativity requires a more richly individuated SNP than grounding the social norm: while a ‘thin’ process of will communication already grounds a social norm, additional process features must come on board to ground social normativity.

¹⁹ See Blomberg 2023 for discussion and Valentini 2023, p. 41 for a similar example.

Which additional process features ensure social normativity? This programmatic paper will briefly suggest three potential answers, in order to give a proof of concept for a process-based approach to social normativity. Fortunately, versions of these answers will already be familiar from debates about what justifies political and legal duties. We regard this parallel as an advantage of our process-based approach: we can rely on well-known reasons for cooperating and respecting law, rather than on a *sui generis* theory of social normativity.

The Autonomy Reason

On one view, a social norm becomes normative if an autonomy condition holds (besides COM):

Collective Autonomy (AUT): The grouping decides whether to require *C* through a process that constitutes an exercise of collective autonomy.

The conditions AUT and COM are independent. While both demand a decision process of whether to require *C*, AUT adds that this process meets the standards of an exercise of collective autonomy, while leaving open whether the requirement finds any support and thus whether the social norm emerges.

Why should AUT (jointly with COM) give reason for *C*? A familiar, broadly Kantian, thesis says that an exercise of *individual* autonomy commands respect.²⁰ Similarly, an exercise of *collective* autonomy in a SNP commands respect, by complying with the social norm. This is because there is value in groupings setting their rules autonomously rather than being governed by external forces. The value of such autonomous self-governance could be intrinsic or instrumental. Indeed, it may firstly be valuable *in itself* that members jointly author their shared environment.²¹ Secondly, this may have valuable effects, including benefits of social cooperation; in particular, members can pursue their goals and set terms of social interaction that would otherwise be impossible.²²

The parallel with legal normativity strikes again: laws owe their normativity partly to an underlying process of autonomous democratic will formation. In Waldron's words:

²⁰ Autonomous individuals are the author of their own life (Raz 1986). Others have a pro tanto duty not to interfere in this authorship.

²¹ Wellman (2003) helpfully conceptualizes group autonomy – though he also points to difficulties with defending such a value. These concerns go beyond the scope of the present discussion.

²² See Scanlon's (1998) 'Principle of Established Practices' and the discussion in Valentini (2023), p. 62-5.

‘A piece of legislation deserves respect because of the achievement it represents in the circumstances of politics: action-in-concert in the face of disagreement’ (Waldron 1999, 108).

If Waldron is right about legislation, an analogous argument for social norms looks promising: they deserve respect and become normative because a collective action problem was solved autonomously. This reasoning also explains why social norms can have force on people who interact with the grouping without being part of it: they should also respect the grouping’s autonomous self-governance.²³

Which SNPs represent an exercise of collective autonomy is debatable. Arguably, SNPs in which members fake or misrepresent their will, so that t-COM fails, do not succeed in generating social normativity. Members may then at best *believe* to have reason to comply, particularly if people fake their wills credibly, so that c-COM holds – the case of pluralistic ignorance discussed earlier. Other counterexamples are arguably SNPs in which members are brainwashed: the wills they express are genuine, but not “free”. Here normativity fails although t-COM holds.

In general, properties of a SNP promoting the exercise of autonomy include: transparency, inclusiveness, open deliberation, and positive responsiveness to individual inputs. Conversely, autonomy is undermined if the SNP is subject to malevolent or arbitrary influences from inside or outside.

In sum, the autonomy of the SNP and thus the social normativity stands and falls with the quality and integrity individual will-formation, will-expression, and will-aggregation.

The Public Reason Reason

Alternatively, a social norm becomes normative if a deliberative-democratic condition holds (besides COM):

Public Reason (REA): The grouping decides whether to require *C* through a process that meets the standards of public reason.

REA requires an exchange in which members justify their wills publicly, by offering publicly acceptable reasons. A public reason exchange would then ground social normativity – just as

²³ For an ambitious argument that justifies democracy on the basis of collective autonomy, see Lovett and Zuehl (2022). However, Lovett and Zuehl commit to a notion of collective democratic authorship based on joint intentions – a step we avoid because it is unclear whether SNPs require joint intentions to succeed.

public reason processes can justify democratic decisions, following a vast literature on public reason and deliberation that we cannot review here.²⁴

On most accounts, public reason processes require high-level cognitive engagement. Few SNPs will clear this hurdle. Examples of SNPs violating REA are non-verbal SNPs, verbal SNPs where members express wills but not reasons, and SNPs where members express publicly unacceptable reasons, such as racist or dishonest reasons.

The Consent Reason

On an entirely different view, a social norm becomes normative if members give their unanimous consent:

Consent (CON): All members give their consent to all members that *C* is obligatory.

The idea is that a requirement binds someone if she has given her consent – just as contracts bind each party because they have given consent, say by signing. Giving consent is an *act*, in fact often the same sort of act by which members also express their agreement in COM* or their will in COM, such as speech acts or acts of nodding. Arguably, giving consent *always* expresses agreement, truthfully or non-truthfully; and so CON implies COM*.

A CON-based account of social normativity faces two problems. First, CON rarely holds, and so few social norms would be normative. Second, explaining why the norm should bind outsiders interacting with the group becomes hard, unlike for an AUT-based or REA-based account. Two unsatisfactory possibilities present themselves. Either the norm does *not* bind outsiders – but this is ‘selective normativity’, which leaves outsiders and newcomers beyond the force of social normativity, against common views and intuitions about the normativity of social norms.²⁵ Or the norm *does* bind outsiders – but this view is question-begging, if not incoherent. The problem is familiar from consent-based views on political obligation (e.g., Simmons 1981).

Additional Background Conditions

Normativity requires not only one of these substantive reasons, but also some background conditions. We propose two such conditions:

Feasibility (FEA): *C* is feasible.

Intra-collectivity (ICO): *C* is only about people in, or interacting with, the grouping.

²⁴ But see Quong (2022) for an overview.

²⁵ See Valentini (2023), pp. 81-2 for discussion.

FEA excludes norms requiring infeasible actions, attitudes or thoughts. Its exact meaning depends on the notion of (in)feasibility, ranging from logical or physical notions to psychological, economic or political notions (e.g., Southwood 2018). One could defend FEA using the principle ‘ought implies can’, applied in a social rather than moral setting.

By ICO, the norm only speaks to people in, or interacting with, the grouping. *C* could say ‘we eat with knife and fork’ or ‘we and our guests eat with knife and fork’, but not ‘all humans eat with knife and fork’. Norms violating ICO will be called ‘imperial’, as they interfere with other groupings. Imagine, say, a social norm created among UK residents prescribing cutlery norms for people living in Vietnam. This norm is *exclusively imperial*: it is *only* about outsiders. Exclusively imperial social norms are rare. *Partially imperial* ones, about insiders and outsiders, emerge in abundance, however. Why? Sometimes real-life SNPs are simply too unsophisticated to set scope-restricted requirements.²⁶ But often the imperial nature of norms is intended: many groupings strive for universal rules, valid for everyone, not just themselves.

Imperial social norms arguably lack normativity, at least for outsiders, because such norms are a form of unjust interference.²⁷ In fact, ICO might *follow* from AUT, REA, and CON, since interfering with others’ self-governance cannot be classified as an exercise of collective autonomy or defended by public reason. Moreover, the interference would not have been unanimously consented to.

Though lacking social normativity, imperial social norms can – unfortunately – still exist and be enforced through sanction mechanisms, creating rational rather than social reasons to comply.

Social normativity summarized

We have proposed three potential sources of social normativity, and some background conditions. A narrow but unified theory of social normativity would accept just one source. For instance, by accepting only the autonomy reason, a normative social norm would arise if and only if this combined condition holds:

$$\underbrace{\text{COM (or a variant)}}_{\rightarrow \text{social norm}} \quad \& \quad \underbrace{\text{AUT}}_{\rightarrow \text{normativity}} \quad \& \quad \underbrace{\text{FEA \& ICO}}_{\rightarrow \text{background conditions}}$$

²⁶ It is easier to agree (through a SNP satisfying COM) that anyone shakes hands with guests than that anyone *in or interacting with our grouping* does so.

²⁷ For analogy, if a national parliament passes a law that requires something from citizens of foreign states outside its jurisdiction and without a relevant connection to the legislating state, then, even though such a law might exist, it would normally lack legal normativity for those ‘outsiders’.

Perhaps more plausibly, one could accept different sources of social normativity. For instance, by accepting all three sources discussed, a normative social norm would arise if and only if this combined condition holds:

$$\underbrace{\text{COM (or a variant)}}_{\rightarrow \text{social norm}} \quad \& \quad \underbrace{[\text{AUT or REA or CON}]}_{\rightarrow \text{normativity}} \quad \& \quad \underbrace{\text{FEA \& ICO}}_{\rightarrow \text{background conditions}}$$

10. Conclusion

To conclude, let us highlight some advantages of our process-based approach.

First, we get the basic ontology right. Social norms *are* neither social processes, nor (as is sometimes suggested) clusters of attitudes. They are *requirements*: requirements grounded in social processes. Precisely which social processes generate social norms is debatable. We have proposed a baseline account (condition COM), and several variants of this account. All these accounts focus on a communication process between members, but what ‘communication’ means varies. Typically, members express certain attitudes to others who perceive these expressions. Whether members truly possess these attitudes is typically irrelevant, as long as they are expressed to others.

Second, the relation between social norms and attitudes is different than usually assumed: rather than grounding social norms, attitudes are typical effects or causes of social processes grounding social norms. For instance, normative expectations are typical effects of social norms, not their grounds.

Third, we can clarify the source of social *normativity*: the normativity of social norms depends on the ‘quality’ of the social norming process. Normativity might arise if the process constitutes an exercise of collective autonomy (condition AUT), meets the standards of public reason (condition REA), or is based on mutual consent (condition CON). Attitude-based accounts struggle with explaining why social norms are normative for those individuals who disagree. To us, attitudes such as disagreement (or agreement) are irrelevant in the first place. What matters is the social norming process.

Fourth, we clearly distinguish social from moral (and other kinds of) normativity. Something can easily be required *social-normatively* but forbidden *morally* – think of social norms with immoral content. While an attitude-based approach can (and should) also maintain this ‘separation of normativities’, our process-based approach emphasises the divide, because social and moral normativity have clearly distinct grounds. Social normativity is *made*, whereas moral normativity is normally taken to exist *independently* of social processes.

Fifth, the connection to legal norms (and normativity) is clarified. Social and legal norms are both grounded in processes. Although social norming processes differ from legal norming processes, some interesting parallels have been highlighted.

Last not least, what is ‘social’ about a social norm? The process-based approach has a clear answer: social norms are generated by an interactive process. The attitude-based approach also has an answer: the attitudes in question are attitudes about others. Yet nothing prevents people from forming other-regarding attitudes in total isolation from one another. Genuinely *social* norms arise from interacting with, not merely thinking or feeling about one another.

References

- Bicchieri, Cristina. 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge, MA: Cambridge University Press.
- Bliss, Ricki, and Kelly Trogdon. 2024. “Metaphysical Grounding.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta and Uri Nodelman, Summer 2024. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/sum2024/entries/grounding/>.
- Blomberg, Olle. 2023. “Tuomela on Social Norms and Group-Social Normativity.” In *Tuomela on Sociality*, edited by Miguel Garcia-Godinez and Rachael Mellin, 219–41. *Philosophers in Depth*. Cham: Springer International Publishing.
- Brennan, Geoffrey, Lina Eriksson, Robert E. Goodin, and Nicholas Southwood. 2013. *Explaining Norms*. Oxford: Oxford University Press.
- Brennan, Geoffrey, and Philip Pettit. 2000. “The Hidden Economy of Esteem.” *Economics and Philosophy* 16: 77–98.
- Diamond, Alma. 2024. “Shadows or Forgeries? Explaining Legal Normativity.” *Canadian Journal of Law & Jurisprudence* 37 (1): 47–78. <https://doi.org/10.1017/cjlj.2024.1>.
- Elster, Jon. 1989. “Social Norms and Economic Theory.” *Journal of Economic Perspectives* 3 (4): 99–117. <https://doi.org/10.1257/jep.3.4.99>.
- Enoch, David. 2011. “Reason-Giving and the Law.” *Oxford Studies in the Philosophy of Law* 1: 1-38.
- Fehr, Ernst, and Urs Fischbacher. 2004. “Social Norms and Human Cooperation.” *Trends in Cognitive Sciences* 8 (4): 185–90.
- Gilbert, Margaret. 1999. “Social Rules: Some Problems for Hart’s Account, and an Alternative Proposal.” *Law and Philosophy* 18 (2): 141–71. <https://doi.org/10.1023/A:1006141723526>.
- Hart, H. L. A. 1961. *The concept of law*. Clarendon law series. Oxford: Clarendon Press.
- Kaplan, Jeffrey. 2017. “Attitude and the Normativity of Law.” *Law and Philosophy* 36 (5): 469–93. <https://doi.org/10.1007/s10982-017-9304-6>.
- Lovett, Adam, and Jake Zuehl. 2022. “The Possibility of Democratic Autonomy.” *Philosophy & Public Affairs* 50 (4): 467–98. <https://doi.org/10.1111/papa.12223>.
- Marmor, Andrei. 2023. *Foundations of institutional reality*. First Edition. Oxford scholarship online. New York: Oxford University Press.

- Marmor, Andrei, and Alexander Sarch. 2019. "The Nature of Law." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2019. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2019/entries/lawphil-nature/>.
- Nieswandt, Katharina. 2024. "Must I Honor Your Convictions? On Laura Valentini's Agency-Respect View." *Analyse & Kritik* 46 (1): 51–65. <https://doi.org/10.1515/auk-2024-2008>.
- Parfit, Derek. 2011. *On What Matters*. I vols. Oxford: Oxford University Press.
- Plunkett, David, Scott J. Shapiro, and Kevin Toh. 2019. *Dimensions of Normativity: New Essays on Metaethics and Jurisprudence*. Oxford University Press.
- Quong, Jonathan. 2022. "Public Reason." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2022. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/sum2022/entries/public-reason/>.
- Rawls, John. 1964. "Legal Obligation and the Duty of Fair Play." In *Law and Philosophy*, edited by Sidney Hook. New York University Press.
- Raz, Joseph. 1979. "Kelsen's Theory of The Basic Norm." In *The Authority of Law: Essays on Law and Morality*. Oxford: Oxford University Press.
- . 1986. *The Morality of Freedom*. Oxford: Clarendon.
- Scanlon, Thomas. 1998. *What we owe to each other*. Cambridge, Mass., London: Belknap Press of Harvard University Press.
- Schotter, A. 1981. *The economic theory of social institutions*. Cambridge [England]: Cambridge University Press.
- Simmons, A. John. 1981. *Moral principles and political obligations*. Princeton, N.J.: Princeton University Press.
- Southwood, Nicholas. 2011. "The Moral / Conventional Distinction." *Mind* 120 (November): 761–802. <https://doi.org/10.1093/mind/fzr048>.
- . 2018. "The Feasibility Issue." *Philosophy Compass* 13 (8): e12509. <https://doi.org/10.1111/phc3.12509>.
- Tuomela, Raimo. 2007. *The philosophy of sociality: the shared point of view*. Oxford: Oxford University Press.
- Ullmann-Margalit, Edna. 1977. *The Emergence of Norms*. Oxford: Clarendon.
- Valentini, Laura. 2021. "Respect for Persons and the Moral Force of Socially Constructed Norms." *Nous* 55 (2): 385–408. <https://doi.org/10.1111/nous.12319>.
- . 2023. *Morality and Socially Constructed Norms*. Oxford University Press.
- Waldron, Jeremy. 1999. *Law and Disagreement*. Oxford: Oxford University Press.
- Wellman, Christopher Heath. 2003. "THE PARADOX OF GROUP AUTONOMY." *Social Philosophy and Policy* 20 (2): 265–85. <https://doi.org/10.1017/S0265052503202119>.