University of Wisconsin Milwaukee
UWM Digital Commons

Theses and Dissertations

May 2015

Constitutive Inescapability and the Search for Normative Authority

Henry Argetsinger
University of Wisconsin-Milwaukee

CONSTITUTIVE INESCAPABILITY AND THE SEARCH FOR

NORMATIVE AUTHORITY

by

Henry Argetsinger

A Thesis Submitted in

Partial Fulfillment of the

Requirements for the Degree of

Master of Arts

in Philosophy

at

The University of Wisconsin-Milwaukee

May 2015

ABSTRACT
CONSTITUTIVE INESCAPABILITY AND THE SEARCH FOR NORMATIVE
AUTHORITY

by

Henry Argetsinger

The University of Wisconsin-Milwaukee, 2015
Under the Supervision of Professor Luca Ferrero

A recent family of views known as constitutivism represents a novel attempt to
ground metaethics in the nature of action. A key move constitutivists make is to ground
"normative authority" in the nature of our practical commitments – or, in other words, in
the inescapability of a practical point of view. In the following paper, I argue that
normative authority *can* emerge from inescapability, and articulate the strongest form of
this constitutivist strategy – one that sees the aim of action as self-understanding. I then
explore a recent set of objections that claims talk of "inescapability" cannot get us
normative authority. The upshot of this view is that any constitutivist strategy not
supplemented by a traditional metaethical account will give us contingent and non-
normatively authoritative aims. I argue that such contingency worries are largely
toothless. Noting, however, that the worries *do* capture the problematically thin nature of
the norms we can derive from a constitutivist project, I sketch a solution. I show that we
can transform the constitutivist view into a constructivist metaethical account and that
this provides the metaethical substance constitutivism might otherwise lack. Finally, I
complete this metaethical project by showing what moral reasoning will look like, and
how it is justified for the metaethical constructivist.

TABLE OF CONTENTS

<u>Introduction</u>

Consider a common moral claim such as, "It is wrong to cause unnecessary suffering." Such a claim appears normative, objective, categorical and authoritative. That is, it gives us a prescriptive "ought" (normativity), which appears to be independent from any particular agent's standpoint (objectivity), seems to apply to all hypothetical cases regardless of circumstance (categoricity), and has a sense of power and proper origin that commands obedience (authority). I will use the phrase "**normative authority**" to refer to this cluster of concepts.[1] Showing that our moral claims *do* have normative authority, and how this might be so, is a central task of metaethics, and the subject of this paper.

A recent family of views known as constitutivism represents a novel attempt to ground this metaethical task in the nature of action. Matthew Silverstein calls the constitutivist method the "Kantian Strategy," as it seeks normative authority in the universal and inescapable nature of our practical commitments (2012, 2). He identifies the "Kantian" position with the claim that, "[e]thical judgments correctly invoke normative authority not by corresponding to or picking out anything in the world, but rather by expressing demands or commitments that are, from the practical point of view, *inescapable*" (2012, 2). In the following paper, I argue that normative authority *can* emerge from inescapability, and articulate the strongest form of this constitutivist strategy.

Chapter One is devoted to constitutivism. In it I briefly lay out the constitutivist landscape before arguing, for the version of constitutivism I find most plausible. I then explore a recent set of objections to the "Kantian Strategy" raised by Silverstein and

---

[1] These bear some obvious similarity to something like R.M. Hare's prescriptivity, universalizability and overridingness. See Hare (1981), esp. Part 1 Chapter 3.

David Enoch. They argue that talk of "inescapability" simply cannot get us normative authority. The upshot for them is that any constitutivist strategy not supplemented by a traditional metaethical account will give us contingent and non-normatively authoritative aims. I argue that such contingency worries are largely toothless. Noting, however, that the worries *do* capture the problematically thin nature of the norms we can derive from a constitutivist project, I sketch a solution in **Chapter Two.** I show that we can transform the constitutivist view into a constructivist metaethical account, and that this provides the metaethical substance constitutivism might otherwise lack. In **Chapter Three** I complete the metaethical constructivist project by showing what moral reasoning will like from this point of view.

My argument proceeds as follows: A successful constitutivist project gives us a substantive aim that all agents share. Such an aim has normative authority in virtue of its inescapabilty. However, given the thin nature of the aim, we cannot expect it to generate robust moral norms. Metaethical constructivism gives us a shared point of view for moral reasoning. We might all share, for instance, the practical point of view of being a "valuer."[2] However, on such a view, the values we happen to have are contingent matters of psychology, taste and evolution. They do not have sufficient normative authority. Noting the strengths and weaknesses of each view, I suggest a combination. The constitutive aim of action can serve as the foundation of our practical point of view. Such a viewpoint allows us to build robust moral norms, and we can ground their normative authority in the type of inescapability that constitutivism provides. I argue that our moral

---

[2] This is the view of Street (2010), who I take as the main exemplar of the type of constructivism that seems workable here, and whose view I discuss later in the paper.

reasoning proceeds from the aim of action (constrained by the consistency and coherence of our beliefs, reasons and desires) towards a recognition of our diachronic nature. This recognition leads us, finally, to the connection between our practical reasoning and the interpersonal social and moral spheres.

I. Constitutivism

_____

1. The Appeal of Constitutivism

As the term is used in a wide variety of ways, I'd like to first be clear about what type of "**constitutivist**" view this paper defends. Constitutivists hope to connect the fundamental features of intentional action to substantive norms. They seek an **aim,**[3] shared by all actions, in order to: 1. Differentiate intentional actions from mere bodily activity, and 2. Provide a standard of assessment by which to judge those actions.[4] Insofar as each and every action shares the aim, each and every action may be assessed in light of it. Constitutivism thus promises a basis for universal norms of assessment grounded in nothing more than the nature of agency itself.[5]

Constitutivists often utilize the analogy of games to explain what such standards look like.[6] Consider a competitive soccer match. There are many rules and regulations, but to engage fully in the activity of playing competitive soccer, one must attempt to score more goals than one's opponent.[7] This aim seems to constitute the practice –

---

[3] I am aware that there is some debate around the use of aims rather than something like "principles." I am not engaging in this debate here, and will continue to use the concept of aims.

[4] See Katsafanas (2013, 1) for a good explanation of this.

[5] It may also partially solve a puzzle about distinguishing true agential action from mere activity going back at least to Harry Frankfurt (1978, 157-162).

[6] Another helpful parallel is the case of belief. It seems that belief constitutively aims at truth. See Velleman (1996, especially 705-706).

[7] We could give a more complex and dense definition that might be more accurate: "the aim is to score more goals than your opponent in the correct way – while following a certain set of rules and regulations, while using a soccer ball, scoring the goals during the correct period of time, etc." But the idea should be intuitive – if you aren't trying to get the ball in the back of the net, you aren't truly playing the game of soccer.

without it, one is simply not in the business of playing soccer. Such analogies illuminate the way in which a constitutive aim both guides the actions of agents and provides a standard of success and assessment.

Consider, for example, a team who approach a match with a strategy of total defense – that is, they do not attempt to score goals at all, but merely try to prevent the opposition from scoring. Commentators will often moan that they are, "not truly playing soccer at all." These pundits are making an evaluation based on the constitutive aim of the game of soccer. If you are not trying to score goals, then there is a real sense in which you are not playing the game, or at least doing a bad job of it.

An objection might seem obvious here: we can do things that look an awful lot like playing soccer, but don't involve trying to score goals. For example, many professional teams set out for a draw against a better club. Are such activities soccer poorly played? Are they soccer at all? If an aim is constitutive, does it demarcate some level below which one is no longer engaged in that activity? Lavin (2004, 424-457) and Korsgaard (2009) have good accounts of this line of questioning. I will address these questions, for the most part, indirectly. It is important to note at least that agency may be importantly different from other enterprises (I expand upon this below), and the example of soccer is merely illustrative of how a constitutive aim might operate. Here I will simply assume that such questions can be answered in a satisfying way – my goal is to consider more closely the normative implications of constitutive theories, so I will be largely leaving puzzles in the philosophy of action to the side.

Return to our example of a competitive soccer match. Following David Velleman, we might say that any activity or enterprise has both a **formal** and a **substantive** aim

(1996, 700). In seeking the constitutive aim of soccer, we seek a substantive aim, rather than a merely formal one. The formal object of the game is to win. But this is uninformative – it is true of any competitive game. The formal object doesn't answer the question, "*how* does one win a soccer match?" I posited that the substantive aim of the game of soccer is to score more goals than one's opponent. This aim tells us *how* to win; it tells us what to do if we attempt to play soccer successfully.

Identifying a constitutive aim provides a more precise understanding of an enterprise, but it doesn't yet get us normative authority. Although useful up to a point, the game analogy invites a rejoinder made famous by David Enoch: even if a game has a constitutive aim, this gives us no motivating reason whatsoever to play the game, or any binding reasons to follow the rules (2006, 169-198; 2011). It might be true that to play soccer we must try to score goals – but who said we wanted to play soccer? The fact that soccer *has* a constitutive aim holds no normative force unless we have decided to play – unless we *share and accept* that aim.

Enoch thus argues that constitutivist views do not adequately answer skeptical challenges to agency. If someone doesn't care about being an agent, how could the constitutive aim of agency have normative authority for them? The skeptic pushes us to admit that the constitutive aim of action itself must have further normative grounding outside of the practice. Constitutivists attempt to counter this push by appealing to a feature unique to agency: its **inescapability**.

The inescapability of an enterprise suggests that we cannot exit it and, at its strongest, that no other enterprises exist outside of it. We can get the gist of inescapability if we consider what is special about being an agent. Luca Ferrero has argued, for

instance, that agency has two special features which grant it inescapability (2009, 308-310).

Special Feature One: All other enterprises fall under agency's domain. (Ferrero 2009, 308) We can leave off playing soccer, but we would still be agents. And so too with any more limited enterprise that can have a constitutive aim. Engaging in any projects or enterprises puts us within the scope of agency. Attempts to renounce or re-evaluate our projects, or *do* or *think* anything *at all*, must occur from within the enterprise of agency.[8]

Special Feature Two: We cannot imagine a challenge to agency from *outside* the perspective of agency. Even if we are considering whether or not to be agents, we can only do so by engaging in agential activity. Thus, even our skeptical questions about the nature of agency occur from the perspective of being an agent (Ferrero 2009, 309). If we ask questions about the normative grounds of our agential commitments, or attempt to consider whether and how we could get beyond the grip of the constitutive aim of action, we do so *as agents*.

Taken together, these special features show that agency is non-optional. Putting the same point in slightly different terms, Velleman claims that agency is "inescapable, because it is fixed by an aim that is naturally inescapable for us as human beings and constitutively inescapable for us as agents" (2009, 140). Agency is naturally inescapable because as human beings we have the kind of mind that must have such an aim, and asking whether we have the aim is a moot question – there is no getting around it (2009, 137). Agency is constitutively inescapable because any attempt to justify its aim involves

---

[8] This doesn't imply that agency is inescapable in some ontological sense. We can, for instance, "leave" agency when we go to sleep, or if we were to die or commit suicide.

our practical reasoning – involves utilizing the very aim in question (2009, 137). To ask for a reason why we ought to be agents is to show that we are susceptible to reasons. If we care about reasons, then we already care about the inescapable practical-reasoning that the constitutive aim of action is connected to.

We now have a sense of the two important dialectical moves employed by constitutivists: **1. Identifying an aim that both explains what an activity is constituted by, and offers a standard of assessment for that activity**, and **2. An appeal to the inescapability of that activity/aim.** By utilizing these moves, constitutivists hope to show that the inescapable aim of action can ground our further normative commitments.

2. The Best Form of Constitutivism

Given the dialectical moves available to constitutivists, I argue for the following as desiderata of a workable constitutivist account: (**D1)** The view must be one that **links practical reasoning and action** in some plausible way that explains why and how we have reasons for action. (**D2)** The aim must be one that **has a formal and substantive object** in order for it to explain *how* we are to act. Finally, (**D3**) It must advance an **aim that is normatively thin**; for two reasons. First, if it is inescapable, it ought to be malleable in light of the many different ways, open to agents, of acting and living rationally. Second, because in searching for the non-normative *facts* about the aim of action we ought to resist building in substantive normativity from the outset.

As an example of how these desiderata can be fulfilled, and to provide the reader with a concrete account of one specific constitutivist project, I will focus briefly on the

work of David Velleman.[9] Velleman, begins by considering the nature of action. He

argues that whenever we act, we do so under the conscious control of reasons (1996).

Action, as opposed to mere activity, is necessarily linked to practical reasoning.[10] We

ought, therefore, to determine the aim of practical reasoning, as this aim will tell us

something about action itself. I agree with Velleman that reasons and action share a

necessary link, which his account explains, thereby satisfying D1.

Recall that any enterprise has both a formal and a substantive aim. The formal

goal of practical reasoning is, Velleman tells us, "figuring out what to do" (1996, 701).

This sounds plausible enough, but the important point is that "what to do" doesn't yet tell

us anything substantive. We don't have a way to *figure out* what to do simply by noting

that practical reason *aims at figuring out* what to do. We cannot simply be looking for the

rational thing to do full-stop, without any reference to a substantive standard (1996, 702)

– this would be like looking for a way to win a soccer match without any reference to the

rules of the game.

Having given us a formal aim, he advances his favored substantive counterpart –

intelligibility or self-understanding (Velleman 2009). There is a pressure towards

making sense to ourselves which seems embedded in the nature of intentional action and

in figuring out what to do. Intending what to do next and knowing what to do next both

seem to require that we understand what we are doing. Action, for Velleman, "is thus

behavior aimed at intelligibility" (2009, 133). Velleman claims that intelligibility is a way

---

[9] I do so mainly as a matter of preference, as I find Velleman's project the most coherent and sensible extension of the constitutive strategy currently in the offing. I think, however, that my preference is justified, and my reasons for rejecting other accounts will become clear as I work through the desiderata of a good constitutive account in this section.

[10] This does not mean that we always plan out our actions beforehand, or that our actions are intrinsically reasonable, or even that an intentional action must have the "weight of reasons" on its side. It simply notes that we must explain actions in terms of reasons.

of holding ourselves together [11] – of telling ourselves a story that makes us a single unified agent; a character that makes sense. We now have the formal and substantive objects of practical-reasoning: *to figure out what to do, by figuring out what it would be intelligible to do*. Thus, Velleman's account satisfies D2. The constitutive aim of action is intelligibility or self-understanding.

To get a sense of this, consider an episode of practical deliberation. Although many of our actions occur non-deliberatively (Arpaly and Schroeder 2013 and Velleman 2014), deliberation might be thought to lead to "full-blooded" actions or actions *par excellence.*[12] When we deliberate, for instance, about going to a movie this afternoon, our internal monologue proceeds in such a way that we both narrate it and witness its progression. We don't know exactly what will come next; yet not just *anything* can follow. We suggest options to ourselves, and our reactions color those options as live or dead. "Shall I go see Movie A?" we ask. "No, I wouldn't like it" we reply upon reflection – and we agree – for if we did not there would be a strange feeling of alienation.

In fact, it would be hard to reach a conclusion that we felt alienated from. We can only decide what movie to go see given what we know about ourselves and the way our deliberation is proceeding. Thus, if we find ourselves walking to go see a film we have just decided against seeing, we would be alarmed, disoriented and conclude that something had gone wrong.[13] We know what it makes sense for us to do, and if we do

---

[11] Sarah Buss, as I understand her, argues for a similar conclusion. Her idea is that we cannot wittingly violate the principle of instrumental rationality; we cannot think of ourselves from more than one point of view. In arguing for this claim she utilizes the idea of "superficial unity." Whatever psychic disharmony we possess, we always strive to tell ourselves a story about ourselves that unifies us (however superficially) into a whole. A key difference in her argument is that we don't have a *desire* for unity or intelligibility – instead, it is just not conceptually possible to be disunified (Buss, 2014).

[12] Velleman's terms, used by Arpaly and Shroeder as well, introduced in Velleman (1996, 714).

[13] I am tempted to say that such an action simply wouldn't make sense, but perhaps this is too strong. See Velleman 2014, especially 13 for an in-depth explanation of this supervisory role of practical reasoning.

otherwise we will be confused and may find we cannot follow through with the action. At the very least, activities that issue *from* us and yet are non-intelligible *to* us will usually not count as full-blooded agential actions.

We have a potentially compelling version of constitutivism on the table. At its most basic, Velleman's claim is that when we engage in practical reasoning about action we are pressured towards accepting and acting for reasons that make sense to us. This satisfies D3, as there may be many different ways of making sense to ourselves – as many ways as there are unique agents. The concept does not seem to build in any substantive norms from the outset. It gives us a coherent account of an inescapable feature of agency: an aim that explains action in terms of the reasons we have to make sense to ourselves.

3. Objections to Inescapability

How can the inescapable aim of self-understanding ground the normative authority of our commitments? The skeptic will not be satisfied by appeals to inescapability. In showing that agency is inescapable, they argue, we have not yet given a normatively authoritative reason for why we *ought* to be agents. Silverstein (2012), for example, criticizes appeals to inescapability as a grounds for normativity in several ways. First, he claims that knowledge of the inescapability of some aim cannot in any way assuage worries about the validity or truth of such an aim.[14] For him, "[o]ur hope to vindicate ethics consists of more than the desire to know that we are all on the same page; we also want to be assured that we are on the *right* page" (Silverstein 2012, 5). I take this

---

[14] Roughly speaking, these are "Mackie-style" worries that our ethical practice is, in principle, unjustifiable and may well be in error. Put forward in Mackie 1977, and similarly in Joyce 2001.

to be the key expression of his critique of constitutivism. For Silverstein, facts about

ethical commitments need to be objective in some deeper sense than being universally

shared. If the universal nature of our commitments is merely contingent, it provides no

normatively authoritative reason to be so committed. Even if some universally shared

moral code existed, this would not be enough to get us the proper objectivity of morality

– an objectivity that calls for the kind of foundational metaphysical facts that only

traditional metaethics can provide.

Silverstein argues that the constitutivist is stuck on the horns of a dilemma:

(**H1**) We might admit that there is an inescapable but *non-*

normative aim of action, and try to infer from this aim some

normative conclusions, or:

(**H2**) We might build normativity into the constitutive aim itself.

If we accept the first horn, there is no compelling reason why we *ought* to hold

that aim. As Silverstein notes, such inferences from "non-normative premises to

normative conclusions [would be] ungrounded" (2012, 8). If we accept the second, it

seems we attempt to sneak substantive normative assumptions in from the start – without

properly justifying them. In order for a constitutivist theory to gain normative authority,

Silverstein argues, it must seek out the resources of traditional metaethics. As he puts it in

his conclusion, constitutivists have no "route around" traditional metaethics (2012, 24).

They must wade into the muck along with the rest of us.

There is good news for the constitutivist: the argument is not as decisive as it

might first appear. It seems clear that we must reject the second horn of the dilemma. As

I noted in the desiderata of a constitutivist account, constitutivists (if they are being

intellectually honest) ought not to begin with robust normativity built into the aim of action. So, we must begin by grasping the first horn – but we can reject Silverstein's characterization of it. In doing work in the philosophy of action, constitutivists are searching for the non-normative facts about the aim of action – facts that do not in and of themselves provide normative authority. Silverstein overlooks a key step in the "inference" from a non-normative premise to a normative conclusion. The inescapable nature of the aim quickly brings to bear two types of normativity.

To begin with, we can make a normative assessment about any action given the enterprise's constitutive aim. Constitutive aims bring to bear standards of success, giving us a mode of assessment, and assessment is a normative tool.[15] Yet, such assessment need not bring to bear normative authority on any particular agent. Some actions may successfully meet their aim and others fail, but what's it to you or me? I might well give a negative assessment of my actions vis-à-vis satisfying some constitutive aim, and yet be unmoved to wish I'd acted differently. Take, as an example, writing a paper. Such an enterprise has a constitutive aim. Let's just say that it's something like "making a compelling argument by doing x and y in the proper way z." We can easily imagine the following: I realize that my term paper is poorly written, and yet, being perfectly happy to get merely a passing grade in the class, I'm unmoved to improve the writing. It seems that I am making no real errors in practical reasoning here.[16] Certainly the realization of

---

[15] Another way of putting it is that agency, on the constitutive account is a goodness fixing kind, see Judith Jarvis Thompson: *Normativity,* 2008.

[16] Assuming some non-perfectionist account of practical reasoning.

the constitutive aim of paper writing doesn't compel me to do better absent some antecedent desire on my part.[17]

The example above assumes that an agent can distance herself from the constitutive aim of her activity. If I recognize that I am a bad paper writer, but don't care about doing a better job, the constitutive aim of paper writing seems to have no normative authority over me. True, I can use the norm as one of assessment, but only in a slightly removed, third-personal way. Yet, even here we might argue that the aim has some force. There is a pressure to conform my actions to the aim at a minimal level that makes my actions recognizable as *paper writing* rather than something else. If I don't care about writing a *good* paper, but still want to pass, I'd better care at least a little about the aim – enough to write a paper at all. Yet, notice that I *needn't care at all*. The enterprise is entirely escapable – I can simply choose not to write a paper.

There is, I have argued, one enterprise this is not the case of: agency. Things get interesting if the constitutive aim is what we might call *internally* inescapable. If the inescapability of the constitutive aim of action is necessarily tied to our own goals and commitments, we feel the authority of such a commitment by our very natures. There is no distance between recognizing the aim and finding it normatively authoritative. It is one thing to say "I recognize the aim of paper writing, but I don't want to write a paper." It is another entirely to say, "I *have* the aim of writing a paper, but I don't want to write a paper." Such a statement shows that something has gone wrong – in much the same way as walking to the movie that we'd decided not to see would show that something had

---

[17] This is assuming some very mild form of Humeanism or an internalism requirement: if the aim doesn't resonate with me, I may have no reason to accept it.

gone wrong. [18] Statements that deny our commitment to the constitutive aim of agency

are of this form. Being an agent is not like building a house, baking a cake or writing a

paper. Agency is inescapable, and its aims originate in our natures – one cannot try to not

be an agent and succeed. Trying to do any intentional action requires agency. Therefore,

no normative assumptions are loaded into the *aim of action* from the outset – rather the

practice and experience of *being an agent* (whatever its constitutive aim might be) is

itself normatively forceful.

This discussion points us towards the second error Silverstein commits in his

characterization of the dilemma. It is not always clear what *kind* of normative authority

Silverstein is searching for. He writes, at times, as if he is unsure how we can get any

normative grip or force at all from inescapability. Yet, this minimal requirement is easily

met. We all feel the normative push and pull of our own commitments. If we are

committed to something inescapably, it is not hard to see why we would be normatively

gripped by it, at least to some extent. If Silverstein cannot see how *this* type of

normativity comes from inescapability then his argument seems off the mark.

At other times, however, he seems worried that such normativity will not have the

right kind of authority. This is a much more compelling worry, but is also the juncture at

which Silverstein and the constitutivist must part ways. He insists that the objectivity

sought by traditional metaethics, the kind of objectivity that posits ethical facts as part of

the fabric of the world, is the only acceptable kind of normative authority. Yet, the

constitutivist argues that this grounding of authority is unnecessary. Ethical claims can

---

[18] What has gone wrong appears to be very similar to what Frankfurt (1971, 10) describes when he claims that "It could not be true both that *A* wants the desire to *X* to move him into action and that he does not want to *X*."

have normative authority by being inescapable. Constitutivists can happily admit that we will not get the type of normative authority Silverstein is after. Such authority is unnecessary, and probably illusory.

Silverstein is, nevertheless, on to something. Building too much normativity into the constitutive aim of action is unconvincing and problematic, as his dilemma showed. So the aim must be rather minimal. Yet, this gives us no clear sense of how it might generate robust normative or ethical principles. Silverstein's critiques suggest the point at which the constitutivist project in the foundations of action might need a supplementary metaethical theory: to move from a minimal aim to more robust norms.[19]

---

[19] Silverstein even alludes to the fact that constructivism is the most natural ally for such a task. Yet, he goes on to insist that constructivism is a reductionist project and "no metaethics at all" (2012, 16).

II. Metaethical Constructivism

_____

1. Constitutivism as Constructivism

We have in hand a particular brand of constitutivism, and some reasons to think that Silverstein's objections against inescapability as grounds for normative authority fall short. The criticism that has yet to be fully discharged is whether such an aim can get us more robust, normatively authoritative ethical commitments. My claim is that supplementing constitutivism with the resources of constructivism can answer this remaining challenge.

What sort of constructivism? I have in mind something like Sharon Street's (2010) version. She argues for a form of **metaethical constructivism** defined by its reliance on the "practical point of view" and its non-restricted or thoroughgoing nature. I will define each of these terms in turn. On Street's view, what is at the "heart" of constructivist theories is not their reliance on the idea of a correctly characterized procedure[20] but, "the notion of the practical point of view and what does or doesn't follow from within it" (2010, 366). What is this **practical point of view**? It is the standpoint of being a *valuer*, and of:

---

[20] We might imagine John Rawls' veil of ignorance as a good example of such a procedure. Whatever conclusions emerge from behind the correctly executed veil on the Rawlsian view constitute the facts about how a just society ought to be set up. Such a proceduralist view is to be rejected in the metaethical realm. If we understand constructivism this way we will not be able to answer the charge that we are merely disagreeing on substantive first-order matters. We can always say, "I see that from your procedure come these facts, but why that procedure and not some other?"

> [taking] at least some things in the world to be good or bad, better or worse, required or optional, worthy or worthless, and so on – the standpoint of a being who judges, whether at a reflective or unreflective level, that some things call for, demand, or provide positive reasons for others. (366).

To operate from this point of view, we need only *engage* in the activity of valuing. Insofar as we engage in reasoning, thinking, caring, arguing, and other typically human activities, we are unable to operate outside the practical point of view of valuing – such a point of view is inescapable. The point of view is shared by all valuers – however much their values might differ.

What might the practical point of view show us? Street argues that we can see what normative conclusions follow from any specific combination of values and facts about the world (367). However, it's important to keep in mind that metaethical constructivism is after a *single* point of view. Building in substantive normative content will make the viewpoint one possible point of view among many; or in Street's words **restricted** (368).[21] In contrast, **thoroughgoing** constructivist views are content-neutral; they concern normativity as entailed by the practical point of view as such (368). The thoroughgoing constructivist view offers a structural characterization of the practical point of view, showing us something about the nature of valuing itself.

An initial skeptical objection looms: such a point of view might tell us how we *do* value, but it does not grant any normative authority to the contingent values we have. Our own contingent points of view as valuers often come into conflict with the contingent valuings of others. We pick out certain objects as valuable; other agents value different

---

[21] Rawlsian political constructivism, for instance, is obviously restricted, taking as it does the basic structure of a liberal democratic society as its area of interest and building in many normative judgments that are widely held in the public political culture - see Rawls (1999), particularly sections 2 and 29.

sets of objects entirely. Without more foundational metaethical facts, how can we say who is right and who is wrong? This skeptical objection is a fair one, yet notice that the nature of valuing compels us to have an opinion on the matter, and the practical point of view compels us to recognize that all human beings share the experience of having such opinions and values. This inescapable valuing *is* the practical point of view. The key question, then, is whether there is any way to move past the realm of arbitrary, if inescapable, opinion to the realm of normatively authoritative, shared commitments.

In the remainder of the section I attempt to answer this question. Here are the crucial points to keep in mind about metaethical constructivism (hereafter simply **MC**): **1. It is best understood in terms of a correctly defined "point of view." 2. Such a point of view is practical;** it can show us what is entailed from a given combination of values and non-normative facts. **3. MC is thoroughgoing** – it gives us a structural characterization of the practical point of view that contains no substantive normative content built in from the start.

The constitutivist and MC views thus share much in common.[22] Both rely on the "Kantian strategy" of locating an inescapable point of view from which to derive norms, and insist that the nature of our practical reasoning can tell us something about normative authority. Yet, it seems that constitutivists who accept the first horn of Silverstein's dilemma must admit that morality will not be entailed by the constitutive aim of action.[23] I suggest that such a constitutivism can be rearticulated as a form of MC – thus diffusing Silverstein's remaining objections.

---

[22] Street herself notes that Velleman's view might count as a version of metaethical constructivism (Street 2010, 382, endnote 28).

[23] Velleman's aim, for example, remains thin and points us in a "pro-moral" direction at best (2009, 2).

In order to generate reasons for action an activity must have both a formal and a substantive constitutive aim. Recall that MC views are thoroughgoing insofar as they give a formal (but not substantive) characterization of the "practical point of view." It may, therefore, seem as if the two views are in tension. MC gives us a structural characterization of valuing. It shows us what it is to be a valuer. In order for the view to show us something of *substance* about that valuing, we must plug in substantive content: various normative commitments. But, we noted above, we lack substantive commitments that might lead to a first order moral theory *beyond merely contingent ones*. If all we have are contingent and differing evaluative starting points, then the morality we are able to extract using the constructive power of the practical point of view will be relativistic in nature, and Silverstein's criticisms about contingency are apt indeed. So MC has a problem: even if it is thoroughgoing, it appears it will be relativistic.

Once again, there is hope. Constitutivism articulates the aim of action, and in showing that it is inescapable, ensures that all agents cannot help but share it. The practical point of view shows us how we inescapably value our commitments – and thus that all agents value the object of the aim of action. Recall, however, the problem with such an aim: we worried that it would be too thin to generate further substantive normative commitments. Notice how, if we combine the two views, both problems are solved: If we have supplied a suitably thin aim of action, then the formal and thus thoroughgoing practical point of view *already* includes the constitutive aim of action as a *non-normative fact* about agential psychology. This leads to a non-contingently shared substantive value in any agent's point of view. We all value "self-understanding" or whatever action correctly aims at. And we can use this shared value to construct further

shared reasons. So the practical point of view, with the constitutive aim of action embedded within it, strikes down worries concerning the contingency of constructivist views and the lack of constructive power of constitutivist views at the same time.

Street worried that adding in substantive norms would make constructivism restricted rather than thoroughgoing. But I've just shown that this isn't so for the constitutive aim of action. This aim is simply part of the full picture of what it means to inhabit the practical point of view. To engage in practical reasoning at all, we must share the constitutive aim of the practice. Thus, if we all inhabit that viewpoint inescapably, the inclusion of such an aim is neither arbitrary nor normatively substantive in such a way as to make the view restricted. This is why I've suggested that what is occurring is not, strictly speaking, a combination of two disparate views, but a more accurate and full account of the structures of constitutivism and MC.

2. Objections Reconsidered

Let us return now to Silverstein's objections and determine whether the rearticulated account moves us forward. Silverstein claimed that the normative authority of our ethical commitments could not be vindicated by constitutivism. Yet constitutivism as the basis of a shared constructive viewpoint seems to lead to sufficient normativity, objectivity, categoricity and authority. Normative authority thus emerges precisely *from the point of view* of agents as valuers.

We must remember, however, that the skeptic objects to the normative authority of the practical point of view itself. Silverstein insists that MC, "involves a kind of *reductionism* about normativity. Normative facts turn out to be facts about what is

entailed from within the practical point of view" (2012, 17). He worries that this leaves

the account dangerously open to the same skeptical worries he pushed against

constitutivism. If facts about normativity reduce to facts about what's entailed from the

practical point of view, we might naturally ask, "why assume that *that* point of view has a

special privilege to the facts? Why not some other?" Or even more incisively, "why care

about these facts at all?" It might be the case that many human beings share the same

contingent set of values, but we can always ask what makes those values good, and

recognize that they are no better than any other psychologically contingent values we

might have started with.[24]

In fact, these challenges are unsuccessful. The first question, regarding other

points of view, fails to ask something meaningful. MC's facts about normativity are

entailed by the structure of the practical point of view.[25] But this type of "reduction" is

not problematically susceptible to the types of skeptical worries pursued by Silverstein.

In saying that normativity is entailed by the practical point of view, we are suggesting

that that's *just all there is* and all there possibly could be about normativity. Normativity

cannot be accessed from any other point of view. So we cannot coherently say, "I see that

the practical point of view entails this set of normative facts, but why not use another

point of view to get another set of facts?" There are no other sets of facts to imagine –

because there are no other points of view to imagine. If agency is inescapable (or to put it

---

[24] Here Silverstein seems to be following the lead of Darwall, Gibbard and Railton (1992) in their survey of metaethics, in which they insist that constructivism is simply a substantive first order normative view. Curiously, Silverstein himself suggests that conceiving of the "Kantian strategy" as an "exercise merely in normative ethics," is a mistake (Silverstein 2012, 9).

[25] As Street explains it, using an analogy about "baseball facts" and what makes it the case that someone is safe on any given play, we have two options. We can give an account where an ideal observer (umpire) of some kind determines the truth about whether a player is safe, or we can give an account of how the rules of baseball and the given set up non-normative facts *entail* that the player is safe (Street 2010, 373).

another way, agency is the *only* point of view we can value from) then a certain type of skeptical challenge does no work here. On the metaethical constructivist's view, an account of normativity divorced from the practical standpoint of a valuer engaged in valuing would be false.

The trickier question is why we ought to *care* about those facts – even if no others exist. This type of open-question challenge might remain: "I see that the practical point of view entails this set of normative facts, but why should I care about them?" If we are inescapably committed to the practical view point, and thus to caring about something or other, then these are the facts we are committed to caring about as human beings. Keeping in mind that there are no facts from some other point of view that we *could* care about (if we are engaged in the enterprise of agency), do we require *further* reasons to care about them? To my mind: no. To Silverstein's: yes. Here, perhaps, we are at an impasse. I take it that I have provided sufficient reasons to think that to continue to demand more is unreasonable.

Let me end this section by making the rearticulated account explicit once more: Constitutivism tells us that all actions share an aim. Such an aim is inescapable. If we are agents, we have this aim, and if my arguments about inescapability go through, we value what it aims at. Thus, the practical point of view can be imagined as a complex web of non-normative facts, and contingent values, where one of those non-normative facts is the aim of action. It follows then that one of the values is non-contingent: the object of the constitutive aim of action. The formal characterization of the practical point of view provides the structure we need to understand how a constitutivist view can be seen as a

compelling, metaethical constructivist view – and such a view, in providing a shared

value, can be used to construct shared norms and reasons.

III. Moral Reasoning

_____

1. The Scope of this Section

At this point, something more needs to be said to give color and content to the above proposal. I do not think we will achieve, without much work, a rich first-order ethics from the constitutive aim of action in the practical point of view alone. Morality won't just fall out of the view for all to admire. If the theory seems thin, remember that it would be trivial and uninteresting to get normativity from something that was normatively thick to begin with – the trick would be too good to believe. In fact, both Street and Velleman argue that we will not get "enlightenment morality" from the practical point of view at all (Velleman 2009, 148-151). I am optimistic that we can get more than they settle for. We ought to remember that we can't move too quickly from a minimal aim to substantive moral content. With that said, the rest of this paper will show that there *is* a rational process getting us from point A – a minimal aim, to point Z – substantive moral reasoning.

We have, thus far, taken one large step in this process: we have shown the source of normative authority derived from the practical point of view. We have shown that an evaluative norm (the aim of action) can, in a suitable fashion, lead to reasons for action that have normative authority. Now we must turn to the substance of such moral reasoning; what will it look like? MC gets us enough to make universal normative claims. The process can now be described as an opening up from internal to external. That is, we want to show that the purely internal constraints of the aim of action – constraints that are

best construed as being about an agent's practical rationality as she herself observes it – can be expanded to the external, social constraints of public reasoning about morality. Such reasoning will not be moral unless it concerns the relationships we have to our fellow agents in social life.

I'd like to begin with a few caveats about what this section will avoid taking a stance on. One way of construing the remainder of the process is as an argument about moral motivation. To avoid the many debates that surround moral motivation, I will generally not cast the process in these terms. With that said, it is fair ask the following question: If we can connect motivating reasons to the aim of action, will these correspond to our intuitions about the scope of morality? In other words we might ask: will an intuitive moral claim like the one that began this paper (it is wrong to cause unnecessary suffering) be one that the aim of action properly motivates? I'll having something to say about this motivational question in **§III.4** below.

I'd also like to avoid any strong stance on the nature of norms. In some obvious sense, they are those rule-like things that guide everything from our use of language, to our practical reasoning, to our actions, to societal sanctions, to our reactive attitudes. But we are concerned here with *moral norms*. I am using a familiar sense of this concept where I mean that a norm is a moral norm just insofar as it carries the normative authority I described at the beginning of the paper (normativity, objectivity, categoricity and authority). It will be important to show that explicitly moral rules can be derived from MC, not purely social norms or norms of self-interest. All that said, I will be considering

the constitutive elements of norms (desires, reasons, intentions, etc.) rather than norms themselves.

This is for two reasons. First, the construction of actual norms is the work of first-order ethics, and I cannot provide a substantive first-order view here (although, as I've said, I endeavor to point the way there from the metaethical work I'm doing). Second, the best way to get a sense of the shape of these moral norms will be to do some constructive work concerning moral reasoning. Therefore, although I will be showing what moral reasoning looks like, I will not go very far into the content of what reasons will actually emerge from the first-order process.

With all this in mind, we can sketch the features of MC that any substantive derivation of moral reasons must put to work. These features serve, primarily, as constraints on the type of reasons that we are able generate. Here is a first approximation: Intelligibility as the aim of action in the context of the practical point of view gives us the following three constraints: **1.** We are unable to accept self-violating maxims. **2.** We are under pressure to have coherent and consistent psychic economies and to avoid certain strains of cognitive dissonance. And **3.** Pragmatic and social norms gain an important weight and relation to our reasons insofar as they are tied to the aim of action. Much more needs to be said about each of these features, and I will take them in turn. The upshot will be that we can connect the particular reasons we have for action to more universal reasons, and that such reasons, being shared, universal things, will move us beyond purely self-interested practical reasoning and towards a moral sphere.

2. Self-Violating Maxims

If we are to succeed in our practical reasoning and in satisfying the aim of action, we cannot accept norms that straightforwardly violate this aim. Take as an example a maxim of this kind: "act in self-deceptive ways." We cannot both (intentionally) aim at understanding ourselves *and* aim at deceiving ourselves.[26] Such a maxim would be self-defeating or self-violating – it aims at an action that frustrates the aim of action. This category of norms seems, on the one hand, relatively uninteresting: it tells us a class of fairly obvious and straightforward rules. On the other hand, it is powerful, in that it gives specification to the source of authority of some norms that we would like to claim border on being self-evident.

There are authoritative reasons against actions that undermine our self-understanding. If our actions aim at intelligibility, then insofar as actions frustrate this aim they will be ruled out. This may take the form of a certain kind of impossibility, but more often than not, they are ruled out as irrational, even if they are practically possible. If we value intelligibility, then we have reason to negatively assess actions, reasons, desires and intentions that interfere with our self-understanding. We do not need to rely on mere intuitions here, strong as they may be. Instead, we can give a formal reason as to why these actions violate a deeper constraint. Aiming at some end while working to undermine that end ought to be ruled out by a theory of practical reasoning – and on this theory we can easily see why this is so.

---

[26] Remember what I've said about self-deception above. Sometimes self-deception appears to aid our understanding – we try to unify ourselves by painting a deceptive picture. But this is different than *knowing* we are lying to ourselves and trying to present that knowledge as truth.

I say, however, that such reasons and norms may not be of great interest. This is so because, first of all, rival theories can generate a similar class of reasons without trouble from merely pragmatic concerns. Something like efficiency or means-end instrumental rationality can quite easily accommodate them. In any case, it is not yet clear if we are saying anything moral. All of these reasons concern self-preservation or self-interest. It would be a much harder sell to claim that an other-regarding reason could be derived from straightforwardly self-violating maxims. How, for example, would we derive a rule like, "one ought to be respectful of the autonomy of others" on purely formal, self-interested grounds? It is not clear that we are violating the aim of our own self-understanding in some obvious way by being disrespectful of the autonomy of others.

There are familiar ways of making such a formal argument. Resources might be mustered from the theories of Kant or Hegel, for instance. Hegel, famously argues that it is a form of false and self-defeating self-consciousness to try to rule over another, and Kant explicitly formulates the categorical imperative to make respecting others as ends in themselves a requirement. But such formalism can seem problematic (among many other reasons) in that its high-minded, conceptual norms have a tenuous relationship to the lived experience that underlies our ordinary assumptions about the permissibility or impermissibility of actions. We would certainly hope that our theory could be accepted by those who aren't particularly sympathetic to Kantian or Hegelian metaphysics and morality.

In fact, although it might be tied to this formal structure in some sense, the real motor of derivation for moral reasons on a constructivist view ought to be something that

does take into account our moral intuitions – perhaps a model somewhat like reflective equilibrium. The formal structure of self-violation sketched here puts itself further beyond suspicion if it is backed up by down-to-earth psychological constraints. When Velleman and others claim that constitutivism will be merely "pro" or proto-moral, it is this formal structure of self-violation they have in mind. We should turn then, to the second constraint, as it is the driving force in the move from merely self-regarding maxims to more general moral norms.

3. Consistency and coherence constraints

The constitutive aim of action, as a model of practical reasoning, carries with it consistency and coherence constraints. Psychological consistency constraints are familiar enough in epistemology, but in this realm we are not merely discussing beliefs. Our view of the mental shouldn't matter overmuch here – the idea is simply that our psychic economies cannot function well with too much disorder or too many conflicts. [27] I call this state of dysfunction one of **cognitive dissonance**, using the phrase to note the very real phenomenological state that occurs within an agent with a messy or clashing psychic economy. In a formal sense, there are constraints that suggest we shouldn't have baldly self-defeating or self-violating aims. In a more substantive sense, part of being an agent involves at least some psychic house-cleaning due to the pressures of avoiding cognitive dissonance.

Even so, consistency constraints can seem a rather poor tool for generating moral content. Are we, for example, supposed to imagine that the reason I shouldn't kill an

---

[27] By psychic economy I mean something like Donald Davidson's notion of the intentional cognitive capacities: the propositional attitudes we hold.

innocent person has anything to do with the confusion it might bring about in my psychic economy? At this stage – perhaps. We are trying to show how to move from purely self-interested practical rationality towards a moral sphere. It is important to be clear, therefore, about what role the avoidance of cognitive dissonance plays in a moral theory that traces its roots back to a constitutive aim of self-understanding.

Our intelligibility suffers insofar as our beliefs, desires, intentions, and reasons for action conflict with one another.[28] The most obvious case is that of conflicting beliefs, which seem to incur the most significant costs. It is widely argued, for example, that we cannot knowingly hold both the belief p and the belief that not p. This is a paradigm case of the type of psychic constraint I'm imagining. There is some sort of mental impossibility implicit in belief formation that restricts such conflicting beliefs.

Far more interesting for our purposes are reasons, intentions, and desires. Such mental states do not have the same flavor of impossibility as in the case of belief. MC can be compatible with a good number of theories about what reasons and desires are, as long as we have some idea of the way in which we rationalize our desires, intentions and actions. That is, as I've argued in Chapter One, a key component of self-understanding is the ability to give reasons for our actions. If it becomes clear to me that I hold conflicting reasons concerning an action (say, some reasons for it and some against) in mind, I am under pressure to resolve this tension. This might take the form of giving up one of the reasons, but, since we are not dealing in impossibilities, it needn't be at all like the case of belief.

---

[28] A well-ordered psychic economy is, therefore, not worth working towards merely due to some fetish for consistency, but a necessary (if always incompletely realized) aim of a creature that is seeking self-understanding.

What seems necessary is *some* resolution *if* the reasons are going to issue in an action. This resolving might take a straightforward route of deciding that some reasons are good and the conflicting ones or bad, or silencing some reasons, or unifying our reasons into a coherent set – but this won't necessarily be so. Such a resolution might take many forms, and in the case of reasons, it might simply involve inaction.[29] But note something here that Christine Korsgaard argues about "our plight" as agents. We must choose what to do – and even choosing to do nothing is a choice (Korsgaard 2009, 1-2.). Thus, when I say that we must resolve these tensions insofar as we are going to act, the door might seem open for a massive amount of psychic conflict in an agent who simply doesn't do a lot of choosing (or acting). But this is a false picture. Even that agent is choosing and acting. By taking a skeptical or stoical stance towards her reasons and desires and trying to "take things as they come," she is exhibiting one way in which we can reach a resolution between conflicting reasons.

This style of resolution is, to some degree, true of all non belief based psychic conflicts. Take the case of a desire, for example. Again, as with reasons, there is nothing explicitly wrong with holding conflicting desires. Yet, the aim of intelligibility puts us under pressure to privilege one desire or another *insofar as we are going to act.* We ought to be careful with this claim – and we mustn't read *too* much into it. It is perfectly acceptable and possible to maintain conflicting desires – we needn't ever choose between them – and there may even be good reasons to maintain some conflicting desires for a

---

[29] One might note here the kind of problems Holly Smith has recently raised about acting on less than perfect or even sufficient information (Smith 2014). If MC requires us to act so as to increase our understanding, we might wonder about its efficacy. In much the same way that act utilitarianism asks us to calculate the overall good produced by an action (a seemingly impossible epistemic task), doesn't MC ask us to calculate what will most increase our self-understanding? I think not. MC requires that we act in a way that we *can* understand – and only in a very strict teleological reading that we do what promotes *most* understanding, although surely there is more to be said in this area.

while.[30] All I'm arguing is that we are under *some* pressure to do psychic house-cleaning, eventually, insofar as we are going to act, and insofar as we are going to be a well-functioning agent. Too many conflicting desires makes life as an agent who *must act* less intelligible.

We are, therefore, focusing on reasons and desires as they relate to intentions and actions. This shouldn't be surprising. A theory that derives moral norms from norms of action must focus in on desires and beliefs that affect our dispositions to act. Let's take an example. Say I desire to eat this piece of cake in front of me, and I also desire, for health reasons, not too. Now imagine that I am in the process of forming an intention – either to eat the cake or not. Part of aiming at self-understanding is rationalizing our intentions. Thus, if I examine my desire to eat the cake, I ought to be able to give a reason for why I might do so that squares up with the rest of my reasons. Say that I cannot find a good reason beyond something like, "because it would feel good in the moment," and I am the type of person that, in general, doesn't take this to be a very good reason to do something – then I will feel some psychic conflict. This explains my guilt, perhaps, at doing something I know I have no good reason to do (or even good reason not to do). In other words what drives us towards unified and well-ordered psychic economies is the interplay of reasons and the aim of intelligibility.

Perhaps the importance of consistency and coherence constraints is now clear. There will be certain reasons with normative authority – the authority granted by the aim of self-understanding.[31] I've noted that such reasons are overriding of our other mental

---

[30] If we don't have a good deal of evidence about what we'd actually enjoy, say, or if we value being able to make spontaneous decisions at the last moment.

[31] What about norms? Here having a fully specified theory of what norms are might change our view, yet what I've said about psychological consistency constraints is intended to show that there will be costs no

attitudes, such that when they conflict with other aspects of our psychic economies, we have reason to accept these authoritative reasons and prune our other mental attitudes. This all adds up to the claim that psychological consistency constraints make a much stronger difference than we might initially have imagined. I shouldn't unify my beliefs, reasons and desires because of some fetishization of consistency. We need to unify, to a degree, for reasons of practicality – insofar as we are able to act. If a conflict of some mental attitudes leads me towards violating my self-understanding, it is rather like the self-violating maxims discussed above. The key move is in tying our mental attitudes to rationalizing reasons by way of the pressures against psychological incoherence. In accepting and following reasons we are utilizing our practical reasoning, and are thus under the sway of the aim of action. We are directed at self-understanding and away from cognitive dissonance.

We've moved from the aim of action, through self-violating reasons, to norms of consistency and coherence in practical reasoning; another step towards an other-regarding normative stance. When we talk about normativity we talk about what reasons we have. If we are talking about reasons, we are talking about things that are, I will argue, universal. This need not mean that being rational implies one exact course of action in any given situation, or one truly best way to live as a human being. It certainly need not be read as a reduction of "us" to some Cartesian agent with no identifying features or particular relationships. Recall that one of my original desiderata for a constitutive

---

matter our theory of norms. If a norm gives us reasons, desires or beliefs that conflict with our mental attitudes we will either be unable to accept the norm, or we must change our mental attitudes. If we view norms as standing beliefs, clearly we cannot comfortably hold a belief that conflicts with a norm we hold (so if we hold the norm that it is wrong to kill, we cannot believe easily in capital punishment). If they are closer to intentions or rationalizations of desires cognitive dissonance still has a great deal of power (so if we derive the norm that we always have reason to help those in immediate danger, our desire to pass by a motorist with a flat tire on a cold and isolated road might be overruled).

account was that it should be flexible in accommodating various forms of leading a life or aiming at the good. While this flexibility is crucial, it also need not imply some kind of particularlism about reasons.

4. Universal Reasons, and the Diachronic Turn

Here, a close parallel can be drawn to some of the concerns that run through the work of Michael Smith. I'd like to briefly discuss the implications of his view of rationality and desire. Doing so will help make it clear why any pressure towards rationality, and thus the shareable nature of reasoning, commits us to a shared social point of view. I note that earlier versions of Smith's view contained several large problems that constitutivism could help solve. In fact, he has recently begun arguing that his view is a constitutive view of agency and reasons. And seeing why he makes this transition will also make explicit the transition to the moral realm I'm after.

Universal Reasons

For Smith, we have a reason to act insofar as we have multiple options for action, and the outcomes our various actions might produce are more or less desirable. The reasons we have are, to put it simply, just reasons about how desirable those outcomes are (Smith 1995, 2013, 1). The more desirable an outcome is, the more reason we have to do it.

On Smith's view, we can show that reasons are universal by asking ourselves, "if I was her, in her situation, would that be a reason for me?" (1995, 121-122) Thus, a reason is universal in the sense that it comes packaged in a recognizable, sharable form.

Although we may not share the actual desires of another agent, we can imagine what life would be like if we did, and see if we would derive the same reasons. This can seem like an ingeniously simple solution. It presents a view where any given combination of desires and beliefs will produce a certain set of sharable reasons. Notice how very similar it is to the practical point of view, as defined by Street, where any given combination of non-normative facts and values will produce a set of normative conclusions – sharable reasons. Smith takes this further than we might initially expect, by imagining that, theoretically, the reasons we *really do have* reduce to the reasons an idealized version of ourselves in the closest possible world [where they are "fully informed and our desires are maximally coherent and unified" (2013, 1)], would advise us to have. And, if Smith is correct in thinking that our idealized selves would all converge on similar desiderative sets in such a world, we won't just be able to see what reasons others would have in principle – we will *share* the very same reasons.

One might be worried by the idea of such convergence. On his 1995 model of universalizability, Smith imagines a universe of shared reasons at the level of propositions that "we would express by using words like 'to help my children', 'to promote my welfare', and the like" (1995, 122). Here, I think, is one location where Smith's earlier view runs into trouble.[32] The underlying problem is that 1995 Smith hasn't given us a compelling picture of why my desires will indeed converge with those of others. It seems right to say that asking for or offering a reason is asking for or offering something universally understandable. But reasons like "it promotes my welfare," that is,

---

[32] His claim would certainly draw the ire of certain strains of feminist epistemology and ethics that claim who we are is integral to what we can know, and our particular relationships are crucially important in telling us what we have reason to do, but let's bracket this concern for now.

reasons at the highest level of generality, are extremely non-specific, and this generality causes several problems. First, such reasons do not really pick out the content that Smith is after. There is a big difference between the meaning of, "I ought to help my children, whoever they may be" and the meaning of "I ought to help that specific actual child Martha who is mine right there." Perhaps, when I imagine the desires and reasons I would have if I was someone else, Smith is after the kind of full imaginative empathy Lewis describes as "gained by difficult exercises of imagination, carried out perhaps in a philosopher's or a novelist's armchair" (Lewis, Smith and Johnston 1989). But this seems both practically and conceptually difficult. Take me as an example: not being a mother, not having children at all; in fact, not caring overly much for children in general – it would be quite a feat for me to truly understand what it feels like to want to save *my own* drowning child. It seems that there is *at least* "one thought too many" here, and that the reasons such exercises would generate don't have the right kind of feel.

The question of how to make my own self-interested practical reasoning amount to something more moral rears its head again. If reasons need to be connected to particular desires in order to properly motivate me to action, it is unlikely that they will be universal in the way Smith wants them to. Yet, if reasons, in presenting something universal and categorical, are divorced from desires, then they do not motivate us to action in the proper way. The generality of the content Smith proposes faces a difficulty: it doesn't seem to pick out the correct content after all.

More importantly, convergence on such general principles is not as rationally guaranteed as Smith suggests. What we have reason to do, on Smith's view, is motivated by what we most desire to do. But what we desire will only be as universal as what we

value. I have no desire to do those things that I find dis-valuable, either instrumentally or inherently. The reasons we share will only be as universal as the values (and thus the desires) that underpin those reasons. This is why I argued in **II.2** that we do all share the value of self-understanding in the practical point of view.[33] And it is the reason that Smith feels the push to "go constitutivist" with his theory.

To reiterate, the question is how to reconcile the universal nature of our rationality with the facts of our contingent valuations and desires. Here is where constitutivism can reinforce a position such as Smith's. Showing that we share the practical point of view and the aim of action is meant to show that we do share an underwriting value (as I've claimed).[34] This shared value moves us beyond the contentless point of view of full rationality, full stop. This is why, for Smith, agents are "desire realizers" – a view of agency with a teleological function. Agents aim at realizing their desires, and thus there is constitutive pressure towards doing so as fully and "optimally" as possible (2013, 3). It is this teleological desire realization which guarantees, Smith thinks, rational convergence.

As I have argued that action has a different aim, I suggest a different route – one that I believe be stronger.[35] The route makes use of a concept that's been in the background of several of my arguments so far: reasonability. That is, I argue that we can

---

[33] Another worry at this point is that one shared value isn't enough – or more precisely, that it doesn't *get* us enough. Elizabeth Anderson, for example, argues that: "Our evaluative experiences, and the judgments based on them, are deeply pluralistic" (1995). If something like value pluralism is true, then we shouldn't expect that all human agents will converge on the same set of values – and this is no surprise given our description of the practical point of view. It seems that if we are pluralists Smith will have to argue for the more complicated and difficult position that actual human beings will converge on the same set of values – not just their idealized selves. Thus, Smith cannot get away from contingency as quickly as he likes – he will have to do something to show that the values we share truly are universal, which may be more or less difficult depending on the theories of value and agency that are in play.

[34] See pages 22-23.

[35] That being said, I'd like to point out that this paper, read as a general argument for the possibility of seeing a constitutivist view as a metaethical constructivist view, might see Smith's view working out fine.

find the content we are seeking in a point of view we might call that of Full Reasonability. Reasonability here is contrasted with rationality, in a way that is indebted to a Rawlsian view. On such a view reasonableness is concerned with being responsive to the reasons of others (Rawls 1999, 358-360; Rawls 2001, section 59).

This viewpoint of Full Reasonability calls for more articulation. To be reasonable, according to something like Rawls' theory, is to be open to reciprocity. It is to have a "disposition to answer in kind." That is, reasonableness is concerned with the equality of other people's reasons with your own. It is concerned with the recognition of your reasons by others, and their reasons by you. When someone is being reasonable, they are responsive in the right kind of ways – and you too will be responsive to good reasons if you are being reasonable. This articulation may seem maddeningly vague, but there it is. If you are reasonable, you are disposed to reciprocate; to treat reasons given by others seriously, and to be fair in asking for and giving reasons.

Let's do some work to tie the various threads of my argument together. The normative authority of the aim of action connects, by way of psychological consistency constraints, to what reasons we can accept. Such reasons are sharable with other agents who share the aim of action (in other words – with all rational agents). But there is a large class of these reasons that are purely self-interested, or at the very least particular to my own point of view; thus we don't yet have a way to show why we ought to regard the reasons of others as (equally) important. My claim is that if we can show that we ought to be reasonable, we can show that the aim of action puts us under pressure to treat seriously the reasons of other agents; not just our own.

The Diachronic Turn

The crucial move, I think, is to prove that the aim of action shows that we should (and must) be reasonable in just this way. Here's how such a proof works: to accept the aim of action is already to accept a diachronic notion of agency. Without taking any strong stance on personal identity over large swaths of time, or on the reality or fictive nature of the self, I argue the following: we relate to ourselves qua self over time. And such a diachronic nature is inexorably bound up with reasonableness. Thus, we are under pressure to be reasonable to ourselves. And if, as Smith suggests, reasons are universal, then we have no less reason to be reasonable towards the reasons of others than our own.

Let's fill out the argument a bit. When I contemplate a reason for a future action, I extend myself diachronically. When I consider the reasons I had for past actions I do the same. The diachronic nature of agency shows that we must have concern for our past and future selves. One way to get at this claim is to make an argument similar to that of Derek Parfit in the Rationality and Time section of *Reasons and Persons*; another would be to follow Nagel in *The Possibility of Altruism*. We can thereby show that we are reasonable to *ourselves.* Intelligibility to ourselves requires intelligibility across time. We must be concerned with the reasons we had in the past and will have in the future. We are called upon to offer up reasons to ourselves – to explain and rationalize our actions to ourselves.[36] It may have seemed that if what I was after was just *self* understanding it would be impossible to generate other regarding norms. But it turns out that such self-understanding is itself reciprocal. And if my self-understanding depends on other selves –

---

[36] If this reminds us of Velleman's improvisational actor, it is no coincidence.

my past and future selves, it may well depend on other selves that are not so intimately my own. Our own self-understanding read as a coherence constraint demands that we respect our future self(s). But why then, should we not extend this consideration to others, whose reasons are just as good and, as generators of good reasons, who might be just as important to my self-understanding?

Smith arrives at this claim from the other direction – for him idealized agents, in converging, are in some sense inseparable from one another, so that "The lesson we learn from the fact that agents can exist as one among many is that the coherence of one agent's psychology must not come at the cost of the incoherence of another agent's psychology. …ideal agents must therefore treat other agents on a par with the way in which they treat future stages of themselves…" (2013, 5). Whichever direction we come from, if our own inter-subjectivity demands that we treat reasons as consistent across time for ourselves, there is no reason to think we shouldn't treat reasons as consistent across space for other people. If we were unreasonable it would involve recognizing the reasons of others as good reasons for them, but claiming that they were not good reasons for me if I were in their situation. And this type of cognitive dissonance impedes self-understanding. When someone else has a reason, I need to treat it seriously.

Let me sum it up in a slightly different way: In moving towards a concern with diachronic agency and universal rationality, we are concerned with the reasons that support certain attitudes and courses of action over time, not just the brute conflicts between them at a given moment. Any kind of reasoning cannot trip over itself; it will involve certain attitudes holding over time – and this shows something about the nature of reasons. I know that I do not want my own self-understanding and integrity impeded –

I have a reason to accept consistency constraints that mediate in favor of self-understanding. I do not want to be harassed or impeded or lied to or killed – and one reason is that such violations of my autonomy interfere greatly with my intelligibility to myself (here construed very broadly). If, in recognizing that other agents also share these concerns, I attempt not to give their similar reasons weight, I am opening myself up to irrationality and, thereby, interfering with my intelligibility. I am saying, "yes, this is a reason for me, and I see that you are in the same circumstances, but I do not count it as a reason for you." To hold these conflicting concepts in mind is to seriously undermine our own self-understanding. Treating reasons as particular rather than universal, I argue, therefore violates the aim of practical reasoning.

5. Social Costs and Rewards

We have thus far shown that the aim of action prevents us from accepting norms that straightforwardly violate self-understanding, and that the practical point of view of valuing self-understanding compels us towards consistency, particularly with regards to reasons for action. In showing this, I've claimed, we have now made ourselves vulnerable to social norms and the reasons of others.

I'd like to pause to answer the following question. Have I been working too hard to prove something rather trivially true - or at least already widely acknowledged? That is, most moral and political theorists are no longer (if they ever truly were) so cavalier about insisting on individual autonomy as somehow divorced from relationships, social custom, culture and community. There is a venerable, and incredibly diverse tradition that sees practical reasoning, moral reasoning, and even autonomy itself as bound up in

the social. Thus, Charles Taylor writes "A great deal of human action only happens insofar as the agent understands and constitutes himself as integrally part of a 'we'" (1995, 168). So too, Wilfrid Sellars tells us:

> [T]he conceptual framework of persons is the framework in which we think of one another as sharing the community intentions which provide the ambience of principles and standards (above all, those which make meaningful discourse and rationality itself possible) within which we live our own individual lives (1963, 39-40).

And of course, there is the Millian idea of social unity ["the social feelings of mankind; the desire to be in unity with our fellow creatures" (Mill 2001, 32)], propounded in *Utilitarianism* and *On Liberty,* to which modern liberal discourse owes so much. Mill tells us that social feeling is a natural sentiment, and the "firm foundation" of utilitarianism, claiming:

> the social state is at once so natural, so necessary, and so habitual to man… [that it] becomes more and more an inseparable part of every person's conception of the state of things which he is born into… The deeply-rooted conception which every individual even now has of himself as a social being, tends to make him feel it one of his natural wants… To those who have it, it possess all the characters of a natural feeling (2001, 32)

Finally we might note that the tradition of seeing the social nature of man as a natural part of his good goes at least as far back as Aristotle, who wrote that friendship "is not only necessary but noble" (1998, 192). Note the order of the pair above. Aristotle didn't feel the need to argue that although friendship was clearly noble it was also necessary, but rather the other way around, even claiming that, "with friends men are more able both to think and act" (1998, 191).

My incomplete survey of the vastly different figures over the course of philosophical history that have deemed the individual's relation to society to be

importantly reciprocal could go on at length. Yet there are reasons why I've tried to do some substantive and formal work to arrive at these conclusions, rather than just noting that they are commonly held. Any ethical theory can show that we live in a society and a linguistic community that prescribes norms and assigns costs and sanctions to those who violate them. Theorists can note that the individual autonomy of an agent bears a strong (if somewhat confusing) relationship to the strength of her society. What I hope the above arguments show is that, although other theories can appeal to such pragmatic costs and naturalistic explanations as reasons to obey moral norms, MC can appeal to an underlying justificatory currency that better explains and motivates them.

There is a rich strain of liberal (and, for that matter communitarian) thought that presupposes man and society as *naturally* mutually supportive, but it would be nice to have theoretical underpinnings that allow us to tell more than an evolutionary just so story. We needn't take Mill or Aristotle at their word that social feeling is a natural sentiment if we can show that it is also rationally required. Or, perhaps less antagonistically, it would be good to show that both the naturalistic and conceptual projects, what Sellars in the article quoted above calls the scientific and manifest images of man at least in this case, converge. Finally, for the non-Hobbesians among us, it would be good to show that we do not have *merely* pragmatic reason not to violate these norms. We'd like to show that violating societal norms in a community of shared reasoners and shared reasons is, in the sense described above, a violation of something moral in nature - something that concerns other people and what we owe them, as well as being a violation of our own self-understanding.

We now have a full sense of the type of Metaethical Constructivist theory I am proposing, and reason to think that a properly construed constitutive account can move us towards a shared moral sphere. Given all that's gone before, one objection in particular, which will appear familiar, calls for some treatment. As has been the case at each step of the process, a skeptical or "open question" type worry might pop up. That is, we might have wondered, as I noted in Chapter Two, whether there was a Moorean open feel to MC's account of motivation. We could ask, "I see that this action promotes self-understanding, but do I have a reason to do it?" I've already given some responses to this line of questioning above, but we now we have another potential reply.

To get a better sense of it, consider for a moment the contractualist moral system of T.M. Scanlon. Scanlon's basic idea is that what morality consists in is that set of principles that no other similarly motivated agent could reasonably reject. MC is set up so that it leads to a very similar first-order view. Now, Scanlon's view also has something to do with his theory of goodness, which he has called "buck passing." In such a theory, the rightness or goodness of an action doesn't provide some further reason to do it – it is those properties themselves that make the action right or good that provide reasons for it.[37]

I mention this because MC has a very similar account of where the buck stops for the normative authority of reasons. In other words, when we note that a reason or norm has properties such that it promotes self-understanding or that it increases intelligibility between ourselves and others, we are not given some new, supervening reason to endorse the reason, perform an action or accept the norm. The reasons we have gain their force

---

[37] See Stratton-Lake and Hooker, "Scanlon versus Moore on Goodness," esp. 150-153 in *Metaethics After Moore* for more discussion of how to formulate and understand Scanlon's Buck Passing account.

from our recognition that the properties of the reason that make it conducive to our self-understanding just *are* what it is for a reason to be have normative authority.[38]

Recognizing that something promotes self-understanding isn't recognizing some *new* reason to do it – there isn't some further reason supervening on natural facts. Recognition of an action's link to self-understanding just *is* recognizing it as reason giving. Put more intuitively: recognition that something promotes the aim of action is recognition that we, as a community, have reason to value it. If we accept the broadened account above, such that the scope of the reasons I care about are social as well as personal, or to put it in Smith's terms, I recognize that I have as much reason not to interfere with the self-understanding of others as I do my own, we can see how social and moral norms get their purchase. If my intelligibility and self-understanding are bound up with the lives of those around me, I do not need to go hunting for extra reasons to acknowledge what I owe to them – the authority of those reasons is bound up in the nature of rationality and reasonableness.

So Scanlon notes that, "when I reflect on the reason that the wrongness of an action seems to supply not to do it, the best description of this reason I can come up with has to do with the relation to others that such acts would put me in: the sense that others could reasonably object to what I do" (1998, 155). Scanlon and I disagree about the ultimate source of such moral requirements. He writes that, "The special force of moral requirements seems quite different from that of, say, principles of logic, even if both are, in some sense, 'inescapable.' And the fault involved in failing to be moved by moral requirements does not seem to be a form of incoherence" (Scanlon 1998, 151). I would

---

[38] This helps answer the motivation problem I noted in **III.1**.

certainly agree with the first part – moral requirements have a special force that is very different from logical requirements – it is the force of normative authority I've been describing in this paper. This is so because, in caring about the judgments of others, we make ourselves socially and interpersonally vulnerable. Such vulnerability carries with it very different sorts of costs than the costs associated with logical mistakes. Yet I disagree with his second claim. In the end, the guilt and shame we feel when we violate a moral requirement are intimately related to a form of incoherence. And it is the inescapability of the aim of action that shows how failing to respond to moral requirements is incoherent.

6. Conclusion

We ought to note that in making first order moral claims, the content we are seeking is no longer metaethical content. Metaethically speaking, we know that any norm we construct must be underwritten by its plausible connection to the constitutive aim of action. But this isn't all we can say about moral prescriptions. The details of moral claims are the purview of first-order ethics and, therefore, I don't have much more to say about them. Constraints of time and space do not allow me to work out a fully formed first order ethics here. I hope to have shown, at least, that what makes something a good piece of moral reasoning is not only whether it promotes the aim of practical reasoning, but also whether it conforms with what rational agents would accept in the practical point of view – with what we owe to one another from the standpoint of Full Reasonability.

We may not have anything so sweeping as Kantian ethics here. But neither do we have to be relativists. The job of our first order ethical theory is going to lie in figuring out what system of norms best expresses the inescapable commitments of the practical

point of view. This will involve figuring out those commitments that all rational agents must accept. In considering Kant's moral theory, Velleman appeals once more to inescapability to do some explanatory work. "When an agent makes an objective practical judgment," he claims, "he recognizes that there is no getting around it, because there is no practical perspective from which it can be reversed. Having no prospect of successful appeal, he has no rational alternative but to accept the judgment, and seeing that there is no alternative is what motivates his obedience" (2009, 167-168).

Although Velleman doesn't have something like first-order constructivism in mind in the above quote, I think we ought to read the spirit of the MC project in this way. From the moral point of view we are considering questions concerning what we owe to each other and what reasons we have that cannot reasonably be rejected by any other agent. Such a point of view is possible because we all value intelligibility. Our desire to make sense to ourselves and one another means that we are moved towards shared ends and shared understandings. When we moralize, we make sense to each other by attending to the shared aims of that project and by what it is reasonable to say and do from such a perspective. Now, such judgments sound to my ears to be as normatively authoritative as we can hope for. Finding out whether we are on the "same page" with our fellow agents, and figuring out how we can *get* on the same page, are at the heart of our ethical endeavors.

REFERENCES

Aristotle. 1998. *Nicomachean Ethics,* translated by David Ross. Revised by J.L. Ackrill and J.O. Urmson. Oxford: Oxford University Press

Arpaly, Nomy and Timothy Schroeder. 2014. *In Praise of Desire*. Oxford: Oxford University Press.

Buss, Sarah. 2014. "Norms of Rationality and the Superficial Unity of the Mind." Presentation, University of Wisconsin Milwaukee Philosophy Department Colloquium. Milwaukee, WI, April 25th.

Darwall, Stephen, Allan Gibbard and Peter Railton. 1992. "Towards Fin de Siècle Ethics: Some Trends." *The Philosophical Review* 101, no. 1: 115-189.

Enoch, David. 2006. "Agency Schmagency: Why Normativity Won't Come from What Is Constitutive of Action." *The Philosophical Review* 115, no. 2: 169-198.

____. 2011. "Schmagency Revisited." In *New Waves in Metaethics,* edited by Michael Brady. Palgrave Macmillan.

Ferrero, Luca. 2009. "Constitutivism and the Inescapability of Agency." *Oxford Studies in Metaethics* 4: 303-333.

Frankfurt, Harry. 1971. "Freedom of the Will and the Concept of a Person." *The Journal of Philosophy* 68, no. 1: 5-20.

____. 1978. "The Problem of Action." *American Philosophical Quarterly* 15, no. 2: 157-162.

____. 1982. "The Importance of What We Care About." *Synthese* 53, no. 2, Matters of the Mind: 257-272.

Joyce, Richard. 2001. *The Myth of Morality*. Cambridge: Cambridge University Press.

Katsafanas, Paul. 2013. "Constitutivism About Practical Reasons." In *The Oxford Handbook on Reasons and Normativity*, edited by Daniel Star. Oxford: Oxford University Press, draft of May 6th.

Korsgaard, Christine. 2008. "The Activity of Reason." *Proceedings and Addresses of the American Philosophical Association* 83, no. 2: 23-43.

____. 2009. *Self Constitution: Agency, Identity, and Integrity*. Oxford: Oxford University Press.

Lavin, Douglass. 2004. "Practical Reason and the Possibility of Error." *Ethics* 114, no. 3: 424-457.

Lenman, James. 2014. "Moral Naturalism." *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), edited by Edward N. Zalta. URL = <http://plato.stanford.edu/archives/spr2014/entries/naturalism-moral/>.

Mackie, John. 1977. *Ethics: Inventing Right and Wrong*. New York: Penguin.

Mill, John Stuart. 2001. *Utilitarianism, Second Edition,* edited by George Sher. Hackett Publishing Company, Inc.: Indianapolis.

Rawls, John. 1999. *A Theory of Justice.* Cambridge: Harvard University Press.

Scanlon, T.M. 1998. *What We Owe to Each Other.* Cambridge: Harvard University Press.

Sellars, Wilfrid. 1963. "Philosophy and the Scientific Image of Man." In *Empiricism and the Philosophy of Mind*, London: Routledge & Kegan Paul Ltd, 1-40.

Silverstein, Matthew. 2012. "Inescapability and Normativity." *Journal of Ethics and Social Philosophy* 6, no. 3: 1-26.

Smith, Michael. 1994. *The Moral Problem*. Oxford: Wiley-Blackwell.

____. 1995. "Internal Reasons." *Philosophy and Phenomenological Research* 55, no. 1: 109-131.

____. 2013. "The Magic of Constitutivism." *American Philosophical Quarterly* 52, no. 2.

Street, Sharon. 2010. "What is Constructivism in Ethics and Metaethics?" *Philosophy Compass* 5: 363-384.

Taylor, Charles. 1995. "To Follow a Rule" *Philosophical Arguments*. Cambridge: Harvard University Press).

Velleman, David. 1996. "The Possibility of Practical Reason." *Ethics* 106, 694-726.

____. 2009. *How We Get Along*. Cambridge: Cambridge University Press.

____. 2014. "Time for Action." Written for the Agency and Time conference at George Washington University (Draft of July 18th).

Williams, Bernard. 1981. "Internal and External Reasons." In *Moral Luck: Philosophical Papers 1973–1980*, Cambridge: Cambridge University Press.

____. 1995. "Ethics and the Fabric of the World." In *Making Sense of Humanity*, Cambridge: Cambridge University Press.