

Imperativism and Pain Intensity

Colin Klein and Manolo Martínez

1 Introduction

Pains vary in intensity. A good philosophical theory of pain should say something about those variations. For one, variations in the intensity of pain are, from a practical standpoint, almost as important as the presence or absence of pain itself. An ibuprofen may dull but not eliminate the pain of your sore muscles. That will have consequences for how you feel, what you do, and what your pain is like.

For another, if you *don't* say something about degrees of pain, your philosophical opponents will tend to assume that you *can't*. We have both recently defended versions of an *imperative* theory of pain.¹ On such a view, pains have imperative contents that express commands to do or avoid certain actions. For example, in Martínez's version, the content of pain experiences is analogous to:

See to it that this bodily state does not exist!

Similarly, in Klein's more recent formulation (forthcoming), a pain is an imperative along the lines of:

Protect this body part (in this manner)!

¹See Klein 2007 and Forthcoming, and Martínez 2010. Richard Hall has also defended an imperative view in his 2008.

On both accounts, a pain is an imperative which would be satisfied if the sufferer of pain ceased to be in a particular bodily state. That bodily state is often one that is threatening to life or health. On the whole, then, pains command us to do things that will keep our bodies intact and well-functioning. Not all pains need do so—the system might misfire in a variety of ways—but the *function* of the pain system is to keep us healthy.

Pains work this way because of the nature of imperatives. Imperative content prescribes some action to be taken, or some goal state to be attained. Pains are not the only sensation with imperative content. The difference between pain and many other sensations can be explained in terms of the differences in the imperatives that constitute them (Hall 2008). Itches command *scratching* regions rather than protecting them. Hunger commands *eating* rather than an action directed at a particular body part. Finally, imperative modalities differ from representational modalities like vision and touch: the latter *represent features* of the world and have truth-apt contents, while the former *command an action* and so do not have truth-apt content.²

In defending imperativism, neither of us has said much about pain intensity. Some have suggested that this is because there is nothing to be said: imperativism, for principled reasons, cannot give a story about intensity. That would appear to be a point in favor of representationalism about pain. After all, many facts about the world vary in magnitude. If pains represent something, then it is easy to see how that something might vary in magnitude, and thereby how pains themselves might have different intensities. The imperativist—who does not think that pains are in the business of tracking properties in the world—has no corresponding facts to appeal to. At least, that’s how the objection goes.

Our opponents are too pessimistic. We can give a perfectly satisfying account of imperative intensity. We’ll present the case in two steps. First, we must establish the general fact that imperatives can vary in intensity. We’ll do that in section 2. Along the way, we will respond to several critics of the imperative theory. Second, we must give a model for variations in intensity, on which the variations are part of the *content* of the imperative. Imperativism is most exciting as an intentionalist theory, one on which

²Recently, it has been suggested that the content of certain perceptual states might have imperatival aspects—see, e.g., Siegel (forthcoming) on *experienced mandates*, or Bengson (in progress) on practical perception. Regardless of whether such suggestions are correct, no one denies that much of the content of perceptual experiences is not imperative.

phenomenal properties supervene on intentional contents. Failure to reduce variations in intensity to variations in content would weaken the appeal. We give such a model in section 3, and conclude with a few brief reflections on the naturalization of imperative intensity.

2 Imperatives come in Degrees

2.1 Imperative intensity

Commands come in degrees. All things being equal, shouting:

Pass the salt now!

conveys a more urgent command than does the calm

Please pass the salt at your earliest convenience!

We will call these variations in the *intensity* of the imperative. Variations in imperative intensity have practical consequences. All things being equal, a more intense imperative should be more likely to make you change your plans, to perform the commanded action sooner, to weight the commanded action higher when deliberating among mutually exclusive courses of action, and so on.

More formally, we will assume that the content of an imperative can be partially identified with a set of *satisfaction conditions*, here modelled as a set of possible worlds $W_{s.c.}$.³ “Pass me the salt!” is satisfied just in case you pass me the salt. So the content of the imperative is, at a first pass, the set of worlds in which you pass me the salt. A command from a source

³We assume this without argument; see Hamblin (1987) for some nice ones. We will refine this quick version, but we will not discuss all of the refinements that might be necessary. Vranas in particular has given compelling arguments that imperative content requires both satisfaction and violation conditions, an alternative that we won’t pursue here (2008).

that you take to be legitimate gives you a reason to act so as to bring about the satisfaction conditions. Such commands will thus have practical consequences for your actions.

We will treat the intensity of a command as a further weighting on the set of satisfaction worlds. At its most general, then, the proposal is this: the content of an imperative includes both a set of satisfaction-worlds and a degree of intensity. Intensity simply functions as a weighting across possible worlds, dividing them into better and worse, in ways to be elaborated shortly.

Intensities might be relatively vague, as in the case of the salt commands above. Or the degree may be more precisely specified, as it is in certain more formal contexts: triage classifications in emergency rooms, the rubber stamps used by enthusiastic managers to prioritize their edicts, or the priority codes on the Autovon telephone network.

The Autovon priority system is an especially nice example of what we have in mind. The Autovon network was developed during the Cold War to provide communication during a nuclear attack. An Autovon keypad had an additional column of four keys which allowed the user to specify the precedence level of the call. The dialed number and the precedence level thus ordered the network to route the call to a certain location, with a certain precedence. (Note that the priorities relate to the *routing of the calls*, not to the contents of the calls themselves. The White House could use Autovon to order a pizza during a nuclear war. While the call itself would trump all others, delivery wouldn't be high on anyone's list of priorities.) Higher precedence calls would, if necessary, kick lower precedence calls off of the trunk to ensure that a call went through. The highest level, which would guarantee that a call would go through if any could, was restricted to the White House.⁴ We will argue that imperative intensity can be generally modeled by something like the Autovon's ranking system. Before we get there, however, we want to say a few words about what imperative intensity is *not*.

⁴A similar, more complex system is still in place. For historical details on the precedence system for Autovon, see page vi of the 1978 "Global Autovon Telephone Directory", digitized online at <http://hdl.handle.net/2027/mdp.39015078412346>.

2.2 What intensity isn't

That commands *can* vary in intensity is, we take it, obvious enough. Further, the extension to pains seems straightforward enough: pains are imperatives, like other imperatives they can vary in intensity, and that variation has exactly the same sorts of practical consequences. A more intense pain in your foot will, all things being equal, cause you to forgo more activities in order to tend to your foot, to tend to your foot more urgently, and so forth.

There is still work left to do to ensure that intensity can be usefully modeled. Before we do it, though, it's worth distinguishing intensities from a few things which look a lot like intensity but aren't. Failure to do so has, we think, led to some of the objections to the imperative account.

First, the intensity of a command ought to be distinguished from the *illocutionary force* of an *utterance* of an imperative. Imperatives can be polite or nasty; they can be phrased as requests or pleas or straightforward commands. None of these are variations in intensity *per se*.

Failure to distinguish the two may lead to the impression that variations in intensity are not part of the content of imperatives. In a recent paper, Cutter and Tye have objected that an appeal to intensity amounts to abandoning intentionalism. Regarding the proposal that imperatives vary in intensity, they write that:

On such a proposal, the difference between [two pains] is analogous to the difference between the following two imperative sentences:

- (Please) stop that bodily disturbance.
- Stop that bodily disturbance!!! (Cutter & Tye 2011)

They then object that such an account is inconsistent with intentionalism because

[N]ow the phenomenal character of an experience does not supervene on its content alone; rather it supervenes on its content together with its degree of urgency. (Cutter & Tye 2011, p. 104)

We agree with Cutter and Tye that the two sentences above may not differ in content. The variations in the two sentences above are variations in *politeness*, however, not the intensity of the imperative.⁵ More generally, two imperatives may differ in illocutionary force without differing in content. The very same command may be ordered, demanded, requested, or politely suggested. Why? Because commands are individuated by their content — on our account, the satisfaction-conditions of the command plus a weighting over worlds — while differences in illocutionary force depend on the social circumstances under which imperatives are uttered. So the very same command, expressed in varying external circumstances, might do different things without changing in content.

There is the potential for confusion, we suggest, because sometimes facts about intensity can be expressed by uttering imperatives with the right illocutionary force — rude commands tend to express high intensity imperatives, for example, and polite ones low intensity. But note that that’s entirely compatible with intensity being part of the *content* of the command. Compare: I might express the relative temperature of a 110-degree day by using a variety of choice expletives. The vigorousness of my swearing might accurately convey the degree to which some external magnitude varies. Yet my claims about the temperature still depends, in an important sense, on what I have said. So too with imperatives.

Given the potential ambiguity, though, we will separate *commands* (a certain type of content) from *orders* (a type of speech act which can convey a command). Note more generally that commands might need to be distinguished from related contents like *permissives*, which permit but do not obligate action (Hamblin 1987).

Second, intensity ought to be distinguished from the *temporal priority* or *urgency* of the command. Some imperatives ought to be satisfied sooner than others. This temporal ordering may depend in part upon the intrinsic content of an imperative. However, priority itself seems like an *extrinsic* property of imperatives. That urgent bit of dusting becomes less urgent when the house is on fire; the urgent ache in your toe suddenly becomes less so when the bear appears from behind the tree. Of course, priority depends on intensity: all things being equal, a more intense imperative should also be

⁵Note that Cutter and Tye use ‘urgency’ rather than intensity. For reasons to become clear shortly, we want to distinguish the two.

satisfied sooner. Our account will preserve this feature. But it also depends on what else is going on, and to which other imperatives you're subject. So while imperative priority is tightly related to intensity, it is not the same thing.⁶ Hence while urgency is extrinsic, intensity need not be.

Third and finally, imperative intensities are features internal to imperatives. They need not track any magnitude in the world. That should not be surprising: the imperativist account, remember, says that pains are primarily spurs to action rather than states which track the world.

Adam Pautz uses this fact to object to the imperative theory. He writes:

While negative imperatives admit of degree, it is hard to see how their degrees might match up with degrees of painfulness. What in the imperative contents of [an agent's] two consecutive pains... might determine that the second pain was roughly twice greater than the first? (Pautz (2010) p. 364, fn. 36)⁷

We think that this is to misunderstand the imperativist position. For all we've said, there's a simple answer to this question: what determines whether one pain is twice greater than another is simply that it is twice as intense. That answer will need some refinement, but in general — the fact that imperatives come in degrees opens the possibility of comparing those degrees.

Of course, pain intensities differ from the intensities of sensations like light and sound. The latter track changes in the magnitude of something in the world. The former don't. Insofar as judgments about pain intensity

⁶Contrary to what one of the present authors once defended, in (Klein 2012).

⁷Pautz actually raises several objections to imperativism that aren't relevant to the main point of the paper but are worth mentioning. He notes that that (i) imperativists have not yet presented a theory of pain locations (true, but a topic for another time; see chapter seven of Klein, forthcoming), (ii) that imperativism implies that increasing thermal sensations actually change the type of content they have (true, and here our intuitions simply differ on how strange this is—after all, the sensation presumably changes for a good reason, namely that the cause has become threatening), and (iii) that imperative theories cannot solve certain problems about magnitude of taste, sound, and color experiences (true, but imperativism doesn't claim that *all* sensations are imperative; these are among those which aren't, and so require a different treatment.).

are about anything, they are about *pains*—and their effects on our motivational state—rather than about the world.⁸ But this seems like an attractive feature of our account, not a problem. The intensity of pain roughly tracks intensity of (say) injury — but only very roughly. Pain from the very same injury may ebb or fall in intensity through the day as we rest, take painkillers, or simply by its own obscure logic. We aren't inclined to think that the *injury* has changed, only that the *pain* has. More generally, it's always an open question whether changes in the intensity of pain correspond to some change in our injuries. It is not always clear to us whether our pains correspond to any injury at all. Indeed, there is an extensive empirical literature, in the form of the gate control theory of pain and its successors, devoted to elucidating the mechanisms of this dissociation.

So there is no obvious problem here. There is, however, a more subtle problem lurking in the vicinity. Subjects will consistently *judge* that two pain intensities are ratio multiples of each other in laboratory settings (Price et al 1983). This is often taken as showing that pain intensity itself is a proper psychophysical magnitude: that is, that it has a theoretically motivated zero point (Nunnally 1967) and that it makes sense to be realists about subjects's judgments because they pick out actual ratios of a real quantity. If so, then pain intensity is comparable to (say) felt intensity of loudness or brightness.

However, the question of whether pain intensity is really a true magnitude, and so *properly* measured by a ratio scale, is a contentious question. The alternative would be that pain intensities, might be merely *orderable*, rather than true magnitudes. That is, it would make sense to talk about pains being stronger or less strong than one another, but not for one pain to be twice as bad as another.⁹ The mere fact that subjects can respond to ratio

⁸In that sense, it's even roughly congruent with Pautz's larger project, which is an attack on 'tracking externalism' (Pautz, 2014). That includes most representationalist theories of pain but arguably does not include imperativism. For a discussion of Pautz's more general argument, as well as a general response on behalf of the representationalist which is also partially applicable here, see (Hilbert and Klein, 2014).

⁹Here is one, admittedly very speculative, argument. A true magnitude requires a zero point. While it might seem obvious that there is a zero point for pains – that is, their absence – the idea of a zero point of pain *intensity* is actually quite odd. Intensity is a quality that all pains ought to possess. If the intensity scale has a zero point, then it should be possible to have something which is still a pain, but which has zero intensity. That's a really odd possibility, and would need some theoretical motivation. The ordinal version accounts for our intuitions far more naturally: there are no cases of pains with zero intensity, but rather of pains that are simply *absent*. When pains *are* present, they

questions and that their response fits a power curve does not itself show that the scale is a ratio scale. For as Hall notes, this is a task “which *pre-supposes* that subjects can judge ratios.” (1981, p. 103). The fact that the resulting judgments fit a power function (especially across the limited range available to ethically acceptable laboratory experiments of pain) may thus show “nothing more than the flexibility of the power function as is evident in its capacity to fit any of the monotonically increasing curves common-sense would expect to describe the relationship between numerical judgments and stimulus intensity.” (p. 104). Ordinal pain intensity would not necessarily prevent people from making judgments about ratios of pains *when instructed to do so*. That might even come out as consistent, especially across limited ranges of stimuli (such as are ethically inducible in a laboratory). Nevertheless, the ultimate question is not which scale *works* but which scale is appropriate for measuring pain.

We think it wise to sidestep this debate. Our goal in the next section, when we model pain intensities, will be only to give a theory that is consistent with subject’s *judgments* of pain intensities, rather than one which assumes something more controversial about the structure of what underlies those intensities themselves. That should make the resulting theory compatible with both magnitude or ordinal theories of pain intensity. It’s worth noting at the outset, however, that magnitude theories get our proposed structure on the cheap. Magnitudes can be represented by simple cardinal numbers on a scale with a true zero. Those should always permit the sort of inter-comparisons that we will propose. If we were sure that pain intensity were a magnitude, it would be tempting just to represent it by a simple number and move on. Since this isn’t obvious, it’s worth building an account that works even on ordinal theories. Further, even if *pain* intensity happens to have a simple structure, *imperative* intensity more generally probably needs an ordinal treatment, so the account should be of broader interest.

always have some intensity that’s comparable (i.e. greater than, equal to, or less than) any other pain. Indeed, the imperativist probably ought to say that intensity might be vanishingly mild, but it’s never zero—the whole point of imperative sensations is, after all, to motivate you, and a zero-intensity pain wouldn’t do that.

Of course, a similar argument could be run for many representational sensations, and there is it more common to assume that they fall on a ratio scale (though this is occasionally debated). Whether this is a reasonable argument, would then depend on whether the disanalogy between pain and straightforward representational sensations is strong enough to block the analogy. That is a question for another paper.

3 Intensity and Content

3.1 Preliminaries

Imperatives vary in intensity; so do pains. A natural conclusion: variations in intensity of pains are due to variations in the intensity or degree of the imperatives that constitute them.

That is not quite enough, however. As we've emphasized, imperativism is a species of intentionalism. Intentionalism claims that phenomenal content supervenes on the content of experiences. We must therefore show that intensity is part of the content of imperatives.

That is not so obvious. One might hold, for example, that imperative intensity is an extrinsic property (like urgency). We have said this is distinct, but that's far from obvious. If it wasn't, that wouldn't necessarily be fatal to imperativism: one could still claim that pains are imperatives but that their intensity depends on their merely functional (rather than content-bearing) relations to other mental states. That would be a sort of quasi-intentionalism, which might in turn still be amenable to a naturalistic treatment.¹⁰ Nevertheless, we take it that one of the attractions of imperativism is its ability to link content and consciousness in what have been traditionally hard cases for the intentionalist, and so it's worth trying to preserve that feature.

The imperativist must also avoid the appearance of an *ad hoc* solution. So, for example, one might treat the content of a pain as simply the ordered pair:

$$\langle W_{sc}, d \rangle$$

where W_{sc} is the set of worlds in which the imperative is satisfied and d is the intensity. That would solve the surface problem neatly. But without saying more, it would be under-motivated. Ideally, the imperativist ought to get the intensity of pain to fall out of a more general theory of imperatives and their structure.

In the remainder of this section, we sketch a model of imperatives on which

¹⁰Thanks to Todd Ganson (personal communication) for this suggestion.

intensity falls out as a natural, intrinsic part of content. We will then return to pain, using our framework to explain some puzzling features of pains, and reflecting a bit on the naturalization of intensity.

3.2 Imperatives and Ranking

Remember that we initially identified the content of an imperative with the set of possible worlds in which the imperative would be satisfied. This can't be sufficient, however. The command

(O) Raise money for Oxfam!

admits of two readings: one on which you are commanded to raise *any* amount of money, another on which the more money you raise the better. Both imperatives would be satisfied in the same worlds: namely, ones where I raise any amount of money for Oxfam. But the two clearly express different commands. This is because the second *ranks* satisfaction-worlds: the more money I raise, the more preferable the world. Similarly, "Clean your room!" expresses a different command from "Clean your room, and the sooner the better!" Both are satisfied in the same worlds. The latter contains an additional exhortation though, which ought to get you moving faster.

This suggests that the content of an imperative is larger than just the satisfaction-worlds: it is instead a set of satisfaction worlds plus a *ranking function* \succsim defined over the set W of all possible worlds. The ranking function fills the gap identified above by introducing a mechanism whereby imperatives satisfied in the same worlds might be compared. For any two possible worlds w_i and w_j , $w_i \succsim w_j$ means that w_i is ranked at least as high as w_j . Two worlds are equally ranked if they are ranked at least as high as each other, and a world w_i is ranked strictly better than w_j just in case w_i is ranked at least as high as w_j and they are not equally ranked.

We can thus model the content of an imperative i as the ordered pair

$$\langle W_{sc}, \succsim \rangle$$

For most simple imperatives, satisfaction is better than non-satisfaction. So each world in W_{sc} is ranked strictly better than each world not in W_{sc} , each pair of worlds in W_{sc} is equally ranked, and *mutis mutandis* for worlds not in W_{sc} . So \succsim partitions W into exactly two equivalence classes: the satisfaction-worlds and the rest, with worlds in the former preferable to the latter. On the first reading of (O), \succsim will simply define a partition with two equivalence classes: one which includes all and only those worlds in which the addressee of the imperative raises any amount of money for Oxfam, another which includes every other possible world, and where any world in the former equivalence class is strictly preferable to any world in the latter.

More complex imperatives partition W into more equivalence classes. On the second reading of (O), \succsim will provide a more complex ranking—the worlds in W_{sc} are better than the worlds outside, but the worlds within will also be further partitioned into smaller equivalence classes, each composed by all and only those worlds in which the addressee of the imperative raises a certain amount of money.

This captures the subtle distinction between the two *meanings* of (O). That distinction is grounded in the differences in the \succsim of the two readings, not in the satisfaction-conditions. That suggests that \succsim is part of the content of imperatives, and necessary to make certain distinctions between them.

Given this expanded content, two observations are the key to a theory of imperative intensity.

First, note that the \succsim for an imperative i ranks over all possible worlds. This means that it will also include worlds in which *other* imperatives are satisfied.

Second, note that \succsim is not restricted to ranking worlds in W_{sc} as higher than worlds not in W_{sc} . That is, it is possible to have an imperative which ranks the satisfaction of some *other* imperative as strictly better than its own satisfaction. That may seem odd. But note that it is precisely the case with the formalized structures of imperatives noted in section 2.1. An Autovon call with status IMMEDIATE ought to be routed to its destination. Further, it ought to be routed preferably to calls with status ROUTINE. So worlds where it is routed and a ROUTINE call is not are ranked as strictly

better by \succsim than worlds in which the reverse is the case. And finally, the call ought to be *dropped* in case of conflict with a call of status FLASH. So there are non-satisfaction worlds (in which the call fails in favor of a FLASH call) which are ranked by \succsim as strictly better than satisfaction worlds (i.e. where the call goes through at the expense of the FLASH call).

We propose that this sort of structure is actually constitutive of imperative intensity. Say that an imperative i is satisfied *at the expense of* some imperative j just in case either i is satisfied and j is not or i is satisfied temporally prior to j being satisfied. Say that an imperative i is *semi-preferable* to an imperative j relative to some (possibly hypothetical) distinct imperative k just in case either:

1. Both i and j rank all worlds in which i is satisfied at the expense of j as better than worlds in which j is satisfied at the expense of i
or
2. i and j cannot be satisfied at the expense of one another, but k is such that
 - (a) i ranks all worlds in which it is satisfied at the expense of k as better than worlds in which k is satisfied at the expense of i and (2) j does not rank all worlds in which it is satisfied at the expense of k as better than worlds in which k is satisfied at the expense of j .
or
 - (b) i ranks some worlds in which it is satisfied at the expense of k as not worse than worlds in which k is satisfied at the expense of i and (2) j ranks all worlds in which it is satisfied at the expense of k as worse than worlds in which k is satisfied at the expense of j .

Given this, we can define imperative intensity as follows:

Intensity An imperative i is more intense than an imperative j just in case there is some imperative relative to which i is semi-preferable to j and no imperative relative to which j is semi-preferable to i .

Let's unpack that a bit. Clause one covers the simple—and by far the most common—cases where i ranks its own satisfaction as better than j 's,

and j agrees. How do we know which of two Autovon phone calls is more important? Well, a FLASH call i ranks its own connection as better than some ROUTINE call j , and j also ranks its completion as less important than the completion of i . So i is semi-preferable to j . As that structure is consistent across the different Autovon levels, j is never semi-preferable to i . Hence i is more intense than j .

The third imperative k , and the indirect structure it makes possible, are necessary for a tricky subset of cases (ones, however, that pains arguably exemplify). Consider the imperative “Repent of your sins, and the sooner the better!” It seems possible that the pastor and the prophet might issue these commands in a way that varies in intensity—the prophet utters it as a matter of gravest importance, while the pastor is more understanding of the complexities of modern life. Yet the two imperatives have the same satisfaction conditions W_{sc} . Further, they always rank earlier satisfaction as *ceteris paribus* better than later satisfaction. So it’s not obvious how to drive a wedge between the two. Note that this problem appears whenever we have two imperatives with the same W_{sc} : it is not possible to construct worlds in one but not the other is satisfied, or where one is satisfied earlier than the other, and so it’s hard to see how they might be ranked.

That’s the point of the second clause. For we can still construct semi-preferability by reference to *other* things one might do. Consider some third imperative k – say, to move your car out of the fire lane. The prophet cares not for fire lanes; $i_{prophet}$ ranks repentance higher than satisfying k . The pastor understands the importance of rendering unto Caesar and all that. So i_{pastor} might rank satisfying k as more important than immediate repentance. (Note that the ranking stipulated by k is irrelevant: it is only the ranking of various forms of i relative to the *satisfaction* of k that is at issue.) This is consistent with the stipulation that the pastor wants you to repent sooner rather than later, so long as that is read with an appropriate *ceteris paribus* clause. For it is true that, everything else being kept fixed, earlier repentance will be better than later. If so, then $i_{prophet}$ is semi-preferable to i_{pastor} . Assuming that this structure holds generally, then the prophet’s imperative is more intense than the pastor’s, even though both are satisfied in exactly the same conditions. Hence for very similar or identically-satisfied imperatives, their intensity can still be compared by triangulating against *other* possible things you might do.

3.3 Advantages of our approach

This approach to imperative intensity has several advantages. Most happily, imperative intensity ends up depending on content. Furthermore, it does so by depending on content in a principled way: the inclusion of \succsim is necessary to capture the differences between the readings of imperatives like (O). Second, we have done so using a purely ordinal measure, \succsim , which leaves open the possibility that the extension to pains might treat pain intensities as merely ordered.

The theory on offer also explains why imperative intensity makes the most sense in limited domains—or, conversely, why many imperatives have *incommensurate* urgencies. If I utter (O) and your department chair orders you to repaint your office, there may not be a well-defined sense in which one of those commands is more intense than the other. That falls out nicely from our theory. The two imperatives are mutually selfish: each ranks its own satisfaction-worlds higher than worlds in which it fails to be satisfied and the others are. So neither is semi-preferable to the other.

Many simple imperatives might be like this, which partially explains why imperative intensity has received so little treatment in the literature. It's also possible for two imperatives to be incommensurate because each is semi-preferable to the other with respect to different third imperatives. Such cases may play a more interesting role in action deliberation when the third imperative is not merely hypothetical but also one that the agent must try to satisfy.

Of course, one must decide which imperative ought to take *priority* in your actions. Priority, as we established previously, may depend in part on facts extrinsic to the imperative—your other plans and desires, the source of the commands, and so on. So priority can be sorted out even in cases where urgencies are incommensurate. However, when two imperatives are commensurate, then the more intense one ought to be (all things considered) satisfied before, and in preference to, the lesser.

Finally, the present theory can account for cases where intensities are judged on a ratio scale even if they are in fact merely ordered. Suppose \succsim has relatively fine-grained rankings over worlds, such that (for example) it distinguishes worlds where you satisfy i and forgo one unit of some good from

worlds where you satisfy i and forgo two units of some good. Suppose further we can identify points of indifference where satisfying i is as good as some quantity of x , better than any lesser quantity, and worse than any greater quantity. Given this, suppose we have two imperatives i and j each of which have this structure relative to the x s, and further that i ranks its satisfaction as on a par with n units of x while j ranks itself as on a par with $2n$ such units. Under such circumstances, we can say that j is not just more intense but *twice as intense as i* .¹¹ One would, of course, need such a structure to be in place consistently—but if it were, it would be natural to speak of ratios of urgencies, not simply ordinal comparisons between them. Note here that the relevant ratio measure has no relation whatsoever to properties, but only to the rankings of different satisfaction and non-satisfaction worlds by the imperative. Hence we can have complex differences in the intensity of an imperative without having to find a sense in which that imperative somehow tracks a magnitude in the world.

3.4 The Intensity of Pain

The extension to pain is straightforward. Pains are imperatives. As part of their content, they have an especially rich and complex \succsim , one rich enough to make the sorts of comparisons noted above. So, for example, one pain is more intense than another just in case their world-rankings agree that one should tend to the body part involved in the first pain at the expense of the one in the second. Differences in intensity of pains in the same body part can be cashed out (for example) in terms of standard gambles against other valued goods.

This structure explains differences in intensities of pain. It also allows plenty of flexibility to account for oddities of cross-modal judgments of intensity. I can judge the intensity of my pain and my hunger; further, I can judge that my hunger is more intense than my pain. It is not clear to us, however,

¹¹Or more precisely, under conditions where both rankings satisfy the axioms of the von Neumann-Morgenstern theorem. If so, then the theorem would ensure that we could find a suitable cardinal function that described an agent's behavior, and that would in turn be enough to account for ratio judgments of intensity as well as standard bets over merely likely satisfactions of i . If you don't like the arbitrary good x , you might also construct a cardinal function out of those standard bets. Note that this is consistent with, but does not imply, the presence of an actual real-valued magnitude which gives rise to \succsim (Ellsberg 1954).

whether we're ever in a position to judge that some hunger is *twice as intense* as a pain. If that is a puzzling thing to say, the phenomenon can be captured by our account: it will occur if \succsim has the requisite structure to make ratio judgments only within modalities, while between-modality imperatives have enough structure to \succsim to allow ordinal but not ratio judgments. This would be the case if, for example, a set of imperatives had a structure like the following (where \sim indicates 'is ranked equally with').

$$\begin{aligned} \succsim_1 & \text{ I don't have pain}_1 \sim [\text{I have pain}_1 \text{ and } \$2] > \text{ I don't have pain}_2 \sim [\text{I} \\ & \text{ have pain}_2 \text{ and } \$1] > \text{ I don't have hunger}_1 > \text{ I don't have } \textit{hunger}_2 \\ \succsim_2 & \text{ I don't have pain}_1 > \text{ I don't have pain}_2 > \text{ I don't have hunger}_1 \sim [\text{I have} \\ & \text{ hunger}_1 \text{ and } \$2] > \text{ I don't have hunger}_2 \sim [\text{I have } \textit{hunger}_2 \text{ and } \$1] \end{aligned}$$

That is, both rankings place pain only as more intense than hunger, but each of them has richer structure within just one modality. Hence cross-modality comparisons are merely ordinal, but intra-modality ones can be cardinal.

Further elaborations of the model could account for other pain phenomena. Here is one intriguing possibility. Merely ordinal rankings—even ones that can be used to derive cardinal ranking functions—are not easily aggregated. Even if, say, the pain in my feet is more intense than both my hunger and my thirst, it's not obvious whether I should stop hiking or continue on given that I am *both* hungry and thirsty.

Models of aggregation do exist, however. In those models, different ordinal rankings receive weightings which determine how much they matter to the final aggregate.¹² In the context of pain, those weights might be interpreted as marking individual differences in the importance of pain: individuals who tend to assign low weights to pains in aggregation will, all things considered, discount even relatively intense pains. Hence we might also have a model for individual differences in imperative intensity, one which would further distinguish between different imperative modalities.

¹²So, for example, one might use the *weighted Kendall-tau distance*. The Kendall-tau distance is a measure of how often reversals occur between pairs of ordinal rankings. Aggregated rankings which minimize the weighted Kendall-tau distance have a number of intuitively appealing properties. See Ailon, Charikar, and Newman 2008 for a discussion and review of other ranking methods, and Martínez (forthcoming) for a ranking-aggregation model of pain.

Finally—since it can't be emphasized enough—these differences in the intensity of pain depend entirely on the contents of the imperatives that constitute pain. Hence imperativism remains consistent with intentionalism.

3.5 The Naturalization of Intensity

A few final thoughts on the naturalization of imperative intensity will serve as a conclusion. One might worry that the contents so proposed are unusually florid, particularly the complexities of the ranking function. Note, however, that this is no argument against a model of content itself. One should no more balk at the complexity of the ranking function than at the use of sets of possible worlds. We have not committed ourselves to a story about how the content of imperatives is implemented, but we see no reason why that story faces any principled problems—why the above, for example, could not be cashed out in an approximate sense at least by patterns and frequencies of neural firings.

That said, imperativism is attractive in part because it promises a full naturalization of pain. We conclude, therefore, with two brief reflections on how thinking of pains as imperatives might be especially useful. First, one might look to work on imperatives in natural language, particularly on how the contents of imperatives gets fixed in conversational contexts. In this regard we have found particularly illuminating Paul Portner's work on imperatives. According to Portner, when someone utters an imperative in a conversation, it adds an item to the *To-do list* of the addressee. So, for example, telling me to pass the salt has the effect of appending the property of being such that I pass the salt to my To-do list (2007). The To-do list functions as an action-oriented version of Stalnaker's Common Ground (Stalnaker 2002). It can be affected by other imperatives, as well as by background desires and goals of the participants. A similar model, we believe, might prove fruitful when exploring the effect of pains on the overall mental economy.

Second, imperativism fits naturally with a teleosemantic account of mental content (Millikan 1984; Martínez 2010). The main contention of this kind of approach is that the content of mental states depends on their biological function, or that of appropriately related states. Cashing out that biological function requires a story about the *consumers* of imperative contents. Here again, the imperative view has much to say. Some regions in the motor

cortex, for example, are likely consumers for painful mental states. Intensities here play an important role. A mental mechanism can advocate for the spending of more or less resources in certain behaviors. Take, for example, physiological responses (increasing heart rate, avoidance) in the face of danger: dangers that are assessed as less pressing provoke a lesser response than more urgent ones. A well-adapted biological agent thus needs both information about what to do and about how important it is to do. Our account of imperative content provides a nice framework within which to cash out that idea.

4 Conclusion

The semantics of imperatives which differ in intensity is not a topic that has received as much attention as others in the study of imperatives. Our discussion of imperative intensity, and the formal apparatus we have used to frame it, is proposed in the spirit of a preliminary exploration. In that spirit, we've shown that imperativism has no insurmountable problems dealing with varying degrees of intensity.

References

- Ailon, N, Charikar, M & Newman, A (2008). Aggregating Inconsistent Information: Ranking and Clustering, *Journal of the ACM*, vol. 55, no. 5: 1–27.
- Cutter, B & Tye, M (2011). ‘Tracking Representationalism and the Painfulness of Pain’, *Philosophical Issues*, vol. 21, The Epistemology of Perception: 90–109.
- Ellsberg, D. (1954). Classic and current notions of measurable utility. *The Economic Journal*, 64(255):528–556.
- Hall, R. (2008). If it itches, scratch! *Australasian Journal of Philosophy*, 86(4):525–535.
- Hall, W. (1981). On “Ratio scales of sensory and affective verbal pain descriptors.” *Pain*, 11(1):101.
- Hamblin, C. (1987). *Imperatives*. Basil Blackwell.
- Hilbert, D and Klein, C (2014). No Problem. Forthcoming in *Consciousness Inside and Out*, ed Richard Brown, Oxford University Press: 299–306.
- Klein, C. (2007). An imperative theory of pain. *Journal of Philosophy*, 104(10):517–532.
- Klein, C. (2012). *Imperatives, phantom pains, and hallucination by presupposition*. *Philosophical Psychology*, 25(6):917–928.
- Klein, C. (forthcoming) *What the Body Commands: The Imperative Theory of Pain*. The MIT Press.
- Martínez, M. (2010). Imperative content and the painfulness of pain. *Phenomenology and the Cognitive Sciences*, 10:67–90.
- Martínez, M. (forthcoming). Pains as Reasons. *Philosophical Studies*.

Millikan, R (1984). *Language, Thought and Other Biological Categories*, The MIT Press.

Nunnally, J.C. (1967) *Psychometric Theory*. McGraw-Hill.

Pautz, A. (2010). Do theories of consciousness rest on a mistake? *Philosophical Issues*, 20(1):333–367.

Pautz, A. (2014). The Real Trouble for Phenomenal Externalists. Forthcoming in *Consciousness Inside and Out*, ed Richard Brown, Oxford University Press: 237–298.

Portner, P. (2007). Imperatives and modals. *Natural Language Semantics*, 15(4):351–383.

Price, D. D., McGrath, P. A., Rafii, A., and Buckingham, B. (1983). The validation of visual analogue scales as ratio scale measures for chronic and experimental pain. *Pain*, 17:45–56.

Stalnaker, R. (2002). Common Ground. *Linguistics and Philosophy*, 25: 701–721.

Vranas, P. (2008). New foundations for imperative logic I: Logical connectives, consistency, and quantifiers. *Nouûs*, 42(4):529–572.