

**Critical Notice of Colin Klein's *What the Body Commands:
The Imperative Theory of Pain* (MIT 2015)**

Murat Aydede
Department of Philosophy
University of British Columbia

Text prepared for the *Author-Meets-Critics* session on Colin Klein's book,
What the Body Commands: The Imperative Theory of Pain (MIT 2015).
Eastern APA Division Meeting, Baltimore — January 4, 2017.

(The other participants were David Bain, and of course, Colin Klein)

[[*This is a slightly more polished version of a presentation I wrote for the Eastern APA session. I wasn't expecting to present all of it. Indeed, I ended up going over only the first half. I've decided to post this commentary online pretty much as is, rather than rewrite it to produce a shortened, more structured and polished version as a critical notice suitable for a journal -- I am afraid I won't have time for it. I hope you will find it helpful. At any rate, please treat this piece as a rough, and somewhat incomplete, draft originally intended to be delivered to a live audience. Also, in what follows, I don't do a lot of exposition of Colin's views. It may be advisable for the reader to read this commentary after reading Colin's book.]]*

I wrote a blurb for the back-cover of Colin's book where I said:

“Arguably, the most thought-provoking, resourceful, and imaginative book on pain written by a philosopher in recent decades. Klein develops a novel philosophical account of pain -- imperativism -- that brilliantly manages to be both empirically informed and theoretically insightful. Well-argued and clearly structured, devoid of excessive professional jargon, lucid and often quite witty -- this is a must-read book for any curious mind, not just for philosophers.”

I may have understated the case for this book. After going over it again for my commentary for this APA session, I've realized how comprehensive and rich the book is. Colin doesn't just plunge into stating his account of pain. He covers pretty much every psychological aspect of pain and even manages to touch on masochistic pleasures (indeed he develops a full theory of masochism -- the *Penumbral Theory* in Chapter 13). This unusually long commentary is a firm testimony of how much Colin's book stimulated my own thinking. Even though I disagree with him on many points, there is a way of interpreting his framework that would bring it close to my own views -- I'll touch on this towards the end.

Following AD Craig and Derek Denton, Colin starts by situating pain among homeostatic sensations. (This is a controversial move, a minority view among pain scientists, but Colin

makes the best case for the claim and develops it more convincingly than anybody else I know of). Colin motivates his imperative account against this background. Pain as a species of homeostatic sensations has only one job to do: to protect the threatened part of the body and help bring the out-of-whack operating bodily parameters into an equilibrium. Homeostatic sensations have imperative contents -- their job is not to inform but move the agent.

Colin is a strong intentionalist ('strong' is my term, Colin doesn't use it). He wants the phenomenal character of sensations (indeed, any experience) to be determined solely by their wide intentional content. He pitches his imperative account as a species of strong intentionalism where pains and other homeostatic sensations have solely imperative content. This separates his intentionalism from the orthodoxy where content is taken to be indicative (truth-apt). He is also a *pure* imperativist in the sense that he never makes use of any truth-apt content, not in his account of the pain sensation, nor in his account of the unpleasantness (affect) of normal pains. For him, both of these aspects are to be handled purely in terms of imperative content.

In my talk, after a general set-up, I'll first try to situate Colin's pure imperativism vis-à-vis what I'll call the Received View. I'll then argue that imperativism as a species of strong intentionalism doesn't have the resources to explain pain phenomenology. I'll then take up Colin's account of pain affect (unpleasantness). I'll end with a critical discussion of his account of body as a minimal practical authority and the related notion of body care. (Along the way, I'll also be touching on a few other things like Colin's treatment of pain intensity and temporal pattern qualities of pain.)

GENERAL SETUP:

Six obvious elements of typical pain phenomenology:

- Location (spatial profile), temporal pattern or profile, quality, intensity (call this group of four elements, LTQI).
- In addition, we have: affect (valence/unpleasantness, badness), and motivation -- call these two elements, AM (or, hedomotive component, following Bain)¹

Some of these elements are themselves complex and admit subtypes (especially true of quality and affect -- see below).

¹ We may split A into two: unpleasantness U and badness E (for *Evaluative*, not to confuse with *Body*). The RV would then be LTQI+UEM. The reason why divide A further into two is that Cutter & Tye (2014) for instance seem to have the following picture in mind: LTQIU+EM.

- The LTQI group has been traditionally labeled under ‘*sensory-discriminative*,’ and the AM group ‘*affective-motivational*.’

The **Received View** (RV) in science and philosophy is that pain has both the sensory-discriminative (LTQI) and affective-motivational (AM) aspects or components. There is, however, less clarity whether pains have both aspects essentially. I suspect that most philosophers who hold the RV think that pains have their sensory aspect essentially and their affective aspect only contingently.² But they nevertheless claim that only in certain rare pathological cases the two aspects come apart. Apart from these rare cases, they always go together and this is not accidental (“contingent” shouldn’t be conflated with “accidental”) -- there are biologically good reasons for them to go together. Because of this, the defenders of RV don’t seem bothered by holding the RV as a view true of the nature of pain.

One important feature of the RV is that it is thought that motivation (whether body or sensation directed) belongs to the affective-motivational aspect of pains. On this view, pain sensations merely as sensations don’t, indeed can’t, motivate (Hume’s influence?). David Bain’s uses the term ‘hedomotive’ to capture both the affective and motivational components of pain to be distinguished from the pain’s sensory component.

So, the **Received View** of pain = LTQI + AM (where ‘+’ signifies a metaphysically contingent but biologically necessary or useful connection)

We can get clearer about the elements.

The LT element can be further differentiated thus:

1. arbitrarily different body parts,
2. the size of the pained area (from needle-size spot to different sized patches and volumes),
3. the changes in 1 (shooting, moving) and 2 (spreading, shrinking), and
4. the temporal pattern (frequency) in any or all of them (1, 2, & 3) – throbbing, pulsating, quivering, flickering, etc. or, longer term patterns.

The Q element admits various subtypes:

Types of pain (by sensory quality from the shortened version of MPQ -- modified):
 (punctate pressure) pricking, drilling, stabbing, lancing; (incisive pressure) sharp, cutting, lacerating; (constrictive pressure) pinching, pressing, gnawing, cramping, crushing; (thermal) tugging, pulling, wrenching, burning, cool, cold, icy, metallic; (brightness) tingling, itchy, smarting, stinging, shocking; (dullness) dull, sore, heavy, sneaking. (Etc. -- there are many more to be added especially if we cover chronic pain.)

² There is strong evidence that pain scientists and clinicians might be thinking of both aspects as essential to pain (see the IASP definition of ‘pain’ and my discussion in XX).

It's not that to each of these adjectives there corresponds a distinct sensory quality -- one would expect to find substantial overlap among them. Nevertheless, the intuitive idea is that there are many distinct sensory qualities peculiar to pains qua pains.

The A element: Primary versus secondary affect:

There are good grounds to distinguish the immediate moment-to-moment hurt of painful pains from their more distant general emotional impact that is more sensitive to the agent's doxastic and conative background and the context. Consider that you discover through a slightly painful tactile encounter a developing lipoma under your arm. Your lymph cancer has been in remission for three years. There are two clearly distinguishable affective aspects of this encounter: one is the brief hurt you experienced when you pressed on your bump, and the other is the sudden, deeper, sickening dread that overcomes you for fear of the cancer coming back. Several emotions are caused here: fear, anxiety, panic, and more. The first brief hurt is clearly not properly described as *suffering*. In the clinical literature, the former is sometimes called the "primary (or immediate) affect" and the latter "secondary affect" (Fields 1999, Price 2000). This distinction can be applied in many kinds of pain including chronic pains. (Colin doesn't mark the distinction.)

The M element: motivation can be direct (immediate) or indirect (deliberative, inferential, requiring more or less background information); more or less under control; body-directed or sensation-directed.

I will take up the **intensity** element later on.

Start with Colin's two "background commitments" (§1.3)

- Pure imperativism
- Strong Intentionalism

PURE IMPERATIVISM:

Colin distinguishes primary motivational force of pain from its secondary motivation:

"... secondary motivation arises in virtue of the way pains *hurt* or cause us to *suffer* or are *painful* or *feel bad*. I'll use these terms to pick out the same phenomenal character felt by many pains, a character that is disagreeable and dislikable. When a pain hurts, we gain a motivation to end that very pain." (p46)

It's somewhat unfortunate that Colin lumps all these terms together to mark the negative valence or affective aspect of typical pains. Colin doesn't make the distinction between primary and secondary affect; not making it, he claims that neither is an intrinsic, necessary, or essential feature of pains. A lot of arguments and examples he uses to

motivate the contingent connection between pain and its negative affect lose their grip once we distinguish the primary from the secondary affect -- more on this later.


For Colin, pains' primary motivational force doesn't come from any sort of affect or valence. Pains are motivational irrespective of their affect or valence. They essentially motivate you to protect body parts where you feel the pain irrespective of whether the pain feels bad or hurts or unpleasant.

Pure imperativism is the claim that pains are essentially constituted by LTQI where this is to be understood reductively in terms of imperatives to protect body parts with a certain urgency. According to Colin:

“The content of each pain can be expressed as an instance of the following schema:
PS: Keep B from E (with priority P)!
B stands for a particular body part, *E* a nominalized passive gerund phrase [designating the sort of protection required], and *P* a ranking function.” (p57)

So, for Colin, each pain sensation (understood as only LTQI) is a command issued by the body to the agent to act in a *certain* way to protect a body part with a certain priority. Commands are intrinsically and essentially motivational, according to Colin (*modulo* the presence of body-care introduced in Chp 12 -- more on this later).

So, Colin's picture is something like this:

LTQI + AM

M₁ + AM₂

Colin doesn't of course deny that pains have LTQI -- he just thinks that these are to be reduced to protective commands with intensity. He insists that AM₂ does not belong to the nature of pain, and he puts it aside for the most part as not essential to his main project -- although he still devotes two chapters to it.

Colin says:

“Imperativism is a species of what I'll call *motivationalism* about pain.
Motivationalists claim that motivational force is an intrinsic property of pains.
Many have denied motivationalism” (p141).

What Colin has in mind are those philosophers who rejected the AM (hedomotive) element of the RV as intrinsic or essential to pains. As far as I know no one has ever denied that pains consist of at least LTQI. (There are a few people who seem to have thought that

pains are perhaps better characterized in terms of the following combination: LTIAM — HR Marshall? D Gustafson? B Helm? J Corns?)

All this may look a bit disorienting. But what is happening is that Colin identifies LTQI with a certain command (issued in sensory code by the body) with an imperative content, and claims that pains are essentially these commands. He doesn't reject that most normal pains have a hedomotive component. He is fine with the RV except that the LTQI is intrinsically motivational in virtue of being commands with imperative content, and that the '+AM' should be interpreted as only expressing a merely contingent feature of pains that may or may not have a further biological role: He claims "that suffering [AM] is not a feature of pain: it is a response *to* pain. This means that suffering is only contingently connected to pain, and hence that pains only contingently hurt and feel bad." (pp46–47)

Strictly speaking "pure imperativism" about pain is indeed pure: it is a single component view about *the essential nature pain*, where indicative information has no role.

So, on this way of counting, what matters is not the outright rejection of hedomotive component of pain but rather its connection to LTQI. If those who promote a "dual-aspect" or composite view of pain had been explicit about the nature of +, it would have been easier to see whether they would count as pure or single-aspect theorist. So, for instance, if strong representationalists such as David Bain and Michael Tye think that the hedomotive component is not metaphysically essential to pain -- they will also count as pure intentionalists except that their analysis of LTQI is indicative and informational: pain in a body part is nothing but an experiential representation of that part as damaged (or the like). What makes them evaluativist is the additional claim that the hedomotive component is also representational: painful pains are those that represent the damage as bad for the agent to varying degrees. They may or may not hold that the hedomotive component is essential to pains. I doubt whether they hold this additional claim about the nature of '+'.

When understood this way, it seems to me that Colin's pure imperativism isn't much different than the dual-aspect views in terms of the general structure of normal pains: they all endorse the RV but may differ in terms of how strongly one interprets '+'.

The real difference then lies in how to analyze the LTQI part.

Colin says the pain sensations (LTQI) are exhausted by certain sorts of commands, and they don't have *any* informational (indicative) content.

Strong representationalists like Bain and Tye disagree: for them, pain sensations are indicative representations about what is physically happening in the part of the body where the pain is felt.³

So, then, the philosophical theories of pain may be usefully classified as theories of LTQI depending on whether the philosopher thinks that the connection is contingent or not. If contingent, then the different theories reduces to disagreements about how to account for LTQI.

Strong intentionalists want to handle this purely in terms of (wide) content. They differ over what kind of content is right: indicative or imperative?

Weak intentionalists (or, the old-fashion pain theorists who happen to be naturalists) might grant that LTQI has informational content about happenings in body parts but deny that pain phenomenology can be explained solely in terms of this content. I consider myself to be an old-fashion pain theorist who takes naturalism/physicalism to be non-negotiable. I argued against strong representationalism in a number of places and offered the beginnings of a naturalist qualia-friendly adverbialist alternative (see my XX). For me the phenomenology of LTQI is partly determined by vehicle properties realizing the pain experiences. The hedomotive component is then functionally defined over these. Long story... totally old-fashioned, but naturalist.

One reason why I prefer my story over strong intentionalism (of both kinds) is because strong intentionalism is committed to a strong form of transparency of experiences; but pain experiences lack this sort of transparency. Colin thinks that his imperativism escapes my criticism -- I'll try to show below that it doesn't.

But let me wrap up this part of the discussion.

The **Received View of Pain** (RV) = LTQI + AM

(where '+' signifies a metaphysically contingent but biologically necessary or useful connection).

The RV accepts the following conditional along with its antecedent:

- if AM is only contingently connected to LTQI and the true essence of pain is exhausted by LTQI, then pain is only contingently unpleasant and only contingently motivating.

³ I'll ignore the complications created by Cutter & Tye (2014) response to Jacobson in seeming to endorse LTQIU+ EM, where U is still the truth-apt [*that* damage is bad for one] kind of representation. Hard to pin them down...

Colin rejects the conditional but accepts the antecedent. But the LTQI element is itself intrinsically motivational for Colin and has no negative affect.

If Colin thinks that AM is not biologically adaptive or useful, then this may be another place whether he and the rest of us might be disagreeing. But otherwise he is a Received Theorist, where he differs is his view of the nature of LTQI.

STRONG INTENTIONALISM

Colin writes:

“I will assume that some form of *intentionalism* about phenomenal character is true. Intentionalism is, minimally, the claim that the phenomenal character of an experience supervenes on its intentional content (Byrne, 2001). This means that there can be no change in phenomenal character without a change in intentional content.” (p7)

And later he says: according to strong intentionalism,

“there are no additional noncontent-based intrinsic properties of experience of which we are aware” (p123)

This latter thesis is what I’d like to call the *Strong Transparency Thesis* (ST) and is implied by Strong Intentionalism (SI). Colin thinks that standard objections to strong intentionalism from the lack of transparency of pain don’t apply to pure imperativism.

“[such objections] — and especially the ones that lean on an appearance-reality gap—presuppose that pain’s content must be indicative. Such arguments lose their grip if we think of pains as having imperative content. What am I aware of when I introspect on a pain? Surely at least three things: where the pain is, what I should be doing with that body part, and how urgently I should be doing it. Those are all features of the *content* of pains. They are not experienced as features of the world. Well and good: they *aren’t* features of the world! Commands reveal ways I must make the world be, not ways it already is. When I look inward, those features of the command are precisely what I am aware of.” (pp124–125)

Let me recast the structure of the required pain commands in my own terms to highlight why and how transparency worries still apply.

Colin proposes to explain the elements of LTQI completely in terms of wide content. We feel pains in body parts. That, for Colin, is the body part that the agent is commanded to protect. We feel the pain with varying degrees of intensity, that is to be handled, according to Colin, in terms of a ranking function among all the possible worlds (not only the satisfaction worlds). But intensity so understood is still purely a matter of content,

according to Colin. The temporal features of pain turn out to be not in the content but are defined over content by an introspective awareness function. This is puzzling and I will come to this later if we have time. I'll for the moment leave the T and I in LTQI aside to come back to it if we have time. Let's concentrate on the location and the sensory qualities that pains have. Differences in the sensory qualities of pains are to be reduced to differences in what sorts of protective actions are commended.

Given all this, why is the command not simply this?:

<protect *that* body part!>

Well, because this would imply that all pains are qualitatively the same. Of course, this command (implied by the more determinate commands) may be what explains what makes all these different types of pains *pains*. But we need to be able to capture every sensory difference there is by specifying a more determinate command that is different from all the other commands corresponding to different sensory qualities. Remember, according to strong intentionalism, there can be no change in phenomenal character without a change in intentional content.

According to Colin, "imperatives really enjoin *actions* not states" (p59). But of course, these actions must be protective, not just any old actions. They must be actions done intentionally aiming protection -- however automatic these actions may be many times. So, we need another parameter to stand for more determinate forms of protective actions:

<protect *that* body part in manner M!>

where 'M' stands for a set of more determinate action types that are either immediate instances of protective action or are such that when executed they will likely result in your body part being protected (e.g., kicking in the groins of the guy who is strangling your neck, versus moving and withdrawing your head or upper torso.)

Colin also distinguishes between passive and active protection commands:

"... active protection commands are satisfied when you do something, whereas passive ones are satisfied if you *refrain* from doing something. Most pains command a mix of active and passive protection. Depending on the situation, one or the other might dominate." (p60)

Colin models the content of protective commands by collecting the set of possible worlds in which the command is satisfied, plus a ranking function. These worlds are presumably centered worlds where we have the body of the agent (perhaps the agent too?) and the time of the command in the center. (Perhaps also in the center is the demonstrated body part-- this depends how Colin want to secure reference to body part -- I see technical difficulties in any choice in terms of possible world semantics).

[[Note that if the imperative content is to be modeled by centered possible world semantics, each token pain is assigned a set of satisfaction worlds where the time is either identical to or later than the time of the command. Either way, there is a proposition fixed by this set -- it's the proposition that is true in each of these worlds, and false in others. Indeed, we can even informally state what that proposition is:

<that body part is being protected by (the agent) in manner M>,

or something along these lines. The set of satisfaction worlds determines a proposition. Then the issue may simply be that the vehicle realized in sensory code has a certain kind of functional role corresponding to imperative mood. I prefer this more traditional story to positing a different kind of *content* -- imperative content -- as Colin does...]]

Each *token* pain has a certain determinate feel to it. According to imperativism, this is completely determined by the imperative content of that token pain. Putting the location and intensity aside, each token pain therefore belongs to a determinate sensory quality type (maybe a profile of complex sensory types -- leave that aside). The question, then, is whether we can pair each qualitatively different token pain with a different imperative content, i.e., with a different set of determinate protective action types (distinct values of M above in my schema).

The answer seems to me to be: No. Let me try to show why. Take for instance all the different qualities of different insect stings or bites described by Schmidt's Sting Pain Index ([SSPI -- see the link](#)).⁴ Whether or not the descriptions are apt, the point is that most pains due to a sting or a bite differ in quality, not just in intensity. With some luck, we may even manage to define a rough (probably partial) quality similarity matrix. Keeping the intensity more or less fixed, consider (counterfactually) the stings/bites by different insects occurring in exactly the same bodily part under the exact same circumstances (where, let's suppose, you don't have a visual/auditory input). Consider the first few seconds right after the sting, say, somewhere in your right forearm. The difficulty for Colin's imperativism is to specify sufficiently fine-grain distinct protective action types for each qualitatively distinct token sting pain during these few seconds.

Keeping the exact physical and psychological context the same, if the bodily spot and the intensity are the same but the sensory quality different, how plausible is it to claim that each such token pain has a different imperative content (i.e., assigned a different set of determinate action types)? In fact, this sort of problem can be generated by most well localized acute injury pains (and even with some warning pains). To generalize, the set of determinate protective action types that can be executed in any given circumstance involving a particular bodily location is fairly fixed, but it seems like the sky is the limit in imagining pains in these exact locations (with the same intensity) differing in quality --

⁴ Colin mentions this briefly but puts it aside by just declaring that his story can handle the sensory complexities with different action types.

indeed many injury types resulting in qualitatively different pains may demand the same set of protective action types. [[Give some further (non-insect) concrete examples here...]]

I can imagine Colin making the following kind of move. Look, he might say, for each qualitatively different token pain (same intensity, same locale, same circumstances), there *is* a different set of protective action types. It all depends on how fine grained we want to individuate the actions. For all we know, for each qualitatively different token pains, there exists a different (although hugely overlapping) set of action types. We may not be able to tell them apart, or tell them apart very easily. In other words, just as the real nature of a red₁₆ surface may not be transparent to us in perception, what different protective actions we may be commanded by qualitatively different pains may not be transparent to us.

But this move is problematic for an imperativist. An indicative representationalist may be happy to say “all I need to know is that this surface color property is different from this or that, or similar to this or that” without compromising the transparency of perceptual experience. But an imperativist cannot say “all I know is that I am being commanded to protect this body part more or less this way or that way but I can’t tell the differences”. The body cannot give two distinct commands to an agent without that agent being able to tell the difference between the different compliance conditions *so that she can act differently for appropriate protection*. In order for the pain system to do its biological job, the protection commands issued by the body need to be transparent to the agent.

Here is a more concrete and vivid way to cast the problem with two stings at the same time. Suppose I am stung by a yellow fire wasp on the right forearm and by a fierce black polybia wasp on my left forearm on exactly the same corresponding spot at the same time. Schmidt describes the quality of the sting of the yellow fire wasp thus: “An odd, distressing pain. Tiny blowtorches kiss your arms and legs” (intensity 3/4). The polybia wasp’s sting feels this way: “A ritual gone wrong. Satanic. The gas lamp in the old church explodes in your face when you light it” (intensity: 3/4). I have two qualitatively very different pains (with the same intensity) on two symmetrical points on in my forearms. Suppose further that the available protective actions under the circumstances for both are symmetrically the same -- as far as anybody can tell. Suppose like Justin Schmidt I am devoted to getting the sting index right or confirming it, or whatever. I am pained, but I soldier on to reflectively attend to the pains right away. I am attending to the huge qualitative difference (see Schmidt’s descriptions) between the two, but I have absolutely no idea in what way they differ in terms of making protective demands on me. In other words, we need to find two different instances of PS that differ in the values ‘M’ takes:

<Protect *that* spot (on your forearm) in manner M with priority P!>

I have no clue what may be plugged into ‘M’ that are different for my both pains.

- (I think:) *that* feels different than *this* one,

referring demonstratively to each sensation. But sensations, according to any strong intentionalism, have no “noncontent-based intrinsic properties” that are introspectively available to us. Given strong transparency, the sensations should be available to me as mere imperative contents directing me to the two spots on my forearms asking me to protect one differently than the other -- by telling me what these protective actions are that are in fact different. But in this case I have absolutely no idea about what the differences are *so I act differently*. In other words, thinking of the two sensations as qualitatively different should directly and immediately draw my attention to two different protective action types in a way that are intelligible to me as such.

In the case of descriptive property attributions to extra-mental objects in perception -- mere differences without being able to tell what the difference consists in -- may be compatible with transparency. Compare:

- (I think:) *this* (pointing to a red₁₆ patch on the wall) looks different to me than *that* one (pointing to a red₁₈ patch just next to the red₁₆ patch)

I am reporting the perceptual registration of two different color properties (instantiated on the wall) as different. My perception may be veridical or not.

But when it comes to being commanded in sensory code to act in two different ways, I need to know what these ways are to be able to comply. If I don't understand how I am being commanded differently, the issue is about the intelligibility of the content. The difference in the contents is not intelligible to me, yet I can tell the difference because they feel differently. This violates transparency.

[[Compare this to the following situation. I am hiking on a mountain trail. I come across two other hikers, one shouting at me to do something in Japanese with a horrified tone and look and proper intonation with relevant gestures. The other is shouting at me in German with a similar gesture/intonation profile. I speak neither German nor Japanese. But I can tell both are issuing a command, even the same command (say, they are in fact asking me to immediately step aside, to avoid a falling rock). I can tell, let's say, that they are issuing protective commands. In this case, I am aware of the qualitative differences between the two commands because I am aware of the differences in the vehicles that carry the imperative content but I don't know what the precise content is. Note in fact they the hikers may even be asking me to do slightly different things. This can't be the proper model to understand what is going on in the pain case above, for Colin, since it straightforwardly violates the transparency. For it involves being aware of the vehicles as different without being aware of the contents. But it seems to me to be a better model all the same.]]

There are two more similar troubling cases for Colin that make related but different points.

Consider the Florida Harvester Ant's bite on the skin. The pain quality and intensity is

described in these terms (from Schmidt's stinging pain scale): "Bold and unrelenting; somebody is using a power drill to excavate your ingrown toenail." In terms of intensity it gets a 3 out of 4 being the most intense pain. Whatever quality is described, the idea is that it will be the same or arbitrarily similar quality you would feel on many different parts of the body. But clearly the types of determinate protective actions (whatever they are) will be (possibly radically) different depending on what body part you get the bite (and even depending on the circumstances). Same quality in different parts (perhaps in different circumstance) -- what explains the *sameness of the quality*?

Here is the second case. Suppose again a FL Harvester ant just bit you on the arm and you felt a sharp intense pain there as a result. Under the circumstances, some available protective action types that this token pain is commanding:

1. Immediately withdraw your arm in the direction opposite to where you got bitten.
2. Shake your forearm vigorously.
3. Move your arm in the same direction to kill the ant by bashing your arm against the wall you were just leaning on.
4. Forcefully use your breath to blow it away.
5. Use your left hand to (a) smack it, or (b) to dislodge and throw it away.

These are a mix of passive and active commands. Consider now changing the circumstances and see what other determinate action types may become available or unavailable. For instance, you may be holding a folded newspaper in your left hand and may use it to smack it instead of your bare hand, or you are not leaning against a wall so #3 is unavailable, or you are standing next to a swimming pool and decide to jump in the water, or you cry out thinking that loud noise would scare away the ant, or your right ankle is recovering from a previous injury so you cannot in fact shake it, etc. In fact, come to think of the range of determinate protective actions you can take depending on different physical as well as psychological circumstances -- it's huge!

So, let us keep the same spot on my body *and* the same quality of its pain (*and* its intensity), but let us change the circumstances (physical as well as psychological). Each such token pain in *different circumstances* will likely bring in different imperative contents -- indeed some may be radically different than others. Again, we have imperative content coming apart from pain quality in a third way.

Colin may be tempted to abstract away from the fine-grained descriptions of the protective actions. There is theoretical pressure to give a more abstract but unifying description of the actions. For instance, what five action types above seem to have in common is something like <remove the ant from the skin>, or <remove *that* from touching the skin>, or something along these lines. If Colin heeds this pressure, he will lose his ability to explain the differentiation in the sensory qualities (or perhaps even bodily locations) of pains. But on the other hand, if he goes more fine-grained in his descriptions of protective actions, he

will lose his ability to explain the similarities or sameness of the pain qualities.

Add to these troubles, the fact that we have been talking *only* about acute injury pains. When you add the *warning* pains, *recovery* pains, and *chronic* pains to this list, I see no reason to think that the qualitative differentiations of pains can be adequately captured in terms of differences in imperative content in a way that will preserve strong intentionalism/transparency.

So far the cases we've looked at are cases where we have sensory quality differentiation without a corresponding content differentiation. But we can also have a somewhat converse problem for imperativism.

Protective command without pain:

Suppose you are seeing and feeling and smelling and hearing (what have you) a light candle approaching your skin when you are tied up firmly. You are feeling a strong urge to protect the relevant skin area issued by your body. Indeed, you are trying to blow it out and trying hard to get your relevant bodily part away from it, to no avail, etc. Imagine this instinctive strong protective urge (command) and your attempt to obey it. It is issued by your body (what else?), directed towards a bodily part, with a certain priority. But there is yet no pain at all (let's assume the candle isn't yet close enough to your skin). Why is this imperative (issued by your body upon reception of sensory information) not pain? You have every element of the mere imperative content firmly impressed upon you. You introspect to find pain, but there is no pain that you are introspectively aware of. Why not? There is the right protective imperative content issued and you are aware of *that* -- yet what you are aware is not pain. Why is this not a counterexample to pure imperativism? The answer can't be because you are not getting the information through appropriate nociceptive/somatosensory afferent channels. Any trace of these, according to Colin, must have been completely lost by the time you are aware of a pain (a mere imperative content).

Generalizing the point: we can easily imagine developing quite specific obsessive behavioral tendencies/urges to protect certain body parts in the absence of non-nociceptive but threatening stimuli. As a result, one might suffer without there being any pain proper in the body parts one tries to protect. But no pain proper. How do you rule out apparent counterexamples like these without giving up pure imperativism?

Itches:

Similarly, why is an itch not a pain? At least some itches? It seems to have the same imperative structure: <Protect B in manner M with priority P!>. I take it that scratching is a very specific and non-flexible form of protective action that can be easily plugged in place of 'M' with appropriate grammatical changes. At least protection seems to be the biological story behind the scratching behavior. If not protection, what is its homeostetic function in the way required by Colin's account of homeostasis? By scratching we are commanded to bring about what bodily parameters in equilibrium if not protecting bodily integrity? It seems to me that we are commanded to scratch because we want to stop the

unpleasant itch sensation itself -- Hall for instance seems to have thought so... But this is a different issue. On Colin's view, itches should count as species of pain, but they are not pains.

I have just argued that, at least with respect to sensory quality, there are qualities without corresponding contents, and that there are also imperative contents of the required sort without corresponding sensory qualities.

But the bigger issue with pure imperativism when it is promoted as a version of strong intentionalism and transparency is this: when I am aware of my pain sensation, that I have been commanded to protect my body in *determinate* ways should be transparent to me irrespective of its hurt. But this is simply not the case. Let me go over a very small sample.

Chronic pains:

Many chronic pain types seem to "command" the sufferer to do things which have nothing to do with protection. It isn't necessarily like you are commanded to protect some part but you can't comply (like phantom limb pains) or what you are commanded is maladaptive -- these entail understanding the command. Many chronic pains, from day to day, if they are commands, are not intelligible in terms of their sensory qualities directing the agent to take certain forms of action independently of their hurting aspect. Even when they can be taken as intelligible command, they have nothing to do with protection.

I sit in a certain position; within minutes my buttocks and back thighs start to have a deep burning slowly pulsating pain. I am all too familiar with this pain -- combination of three nerve problems in my lumbar spine. It's started with no apparent reason a number of years ago progressively getting worse and the pain is spreading. It is by no means clear to me that I am being commanded to do anything. If I may be said to be commanded, it's not clear at all what it is that I am being commanded. Even if I am being commanded to act or avoid acting in a certain way that I understand, there is no guarantee that the action type is protective of tissue integrity. Indeed, not sitting is the only thing that helps in my case. Here this action type is clearly not protective of bodily integrity. If this is a command, it is necessitated in order to avoid the sensation of pain. I just want *that* to stop (where the referent of 'that' is my pain -- not intelligible to me as a protective command). But this is the negative affect of pain, it's a reaction to pain, not the pain itself. On Klein's view, this is to be accounted by a second order imperative <stop that pain!>. Furthermore, the sensory qualities of my pain have periodic complex sensory variations that seem to me not correlated with anything that happens to me or what I do... If strong intentionalism had been true, what I have just said would have been *obviously* false! But just on phenomenological grounds, it seems to me eminently true.

Sciatica:

Colin says: “The pain caused by sciatica is (more or less) identical to the pain caused by some disorders of the foot itself. Both pains command you to protect your foot.” (p115). This is certainly true for some sciatic pains. But by no means true for all kinds of sciatic pains. Just ask any clinicians or pain scientists who have also clinical experience. They will tell you that the qualities of prolonged sciatic pains undergo significant qualitative changes that have very little resemblance to the initial injury phase or recovery phase. In fact, many central as well as neuropathological pains have qualities that are not intelligible as protective commands. [[Give some concrete examples and case histories here...]] If there is any command, it is the stop-that-sensation command and the accompanying behavior aimed to bring this out. But this is a completely different issue for Colin.

Chronic headaches.

There are qualitatively very different chronic headaches. Colin writes:

“I’ll assume that headaches are just a classic case of a misfire in the most adaptive system. The most adaptive system might have reason to cause pain due to transient ischemia (perhaps to regulate exertion), and it might be unusually sensitive to transient ischemia even though the head muscles aren’t involved in vigorous exercise. One has only so many options for building a body when you start from a single cell each time. That process apparently spreads receptors far and wide, including places where they’re less useful. So it goes. Even if headaches are functionally mysterious, however, their content isn’t, “Protect your head!” (pp114–115)

However, even though Colin’s explanation of why we have these chronic pain syndromes may be true, it doesn’t address the problem of explaining qualitatively different kinds of headaches even if we agree that they all command <protect your head>. What are the *specific* ways of protecting my head in having these different sorts of migraines? I have no clue, I am being “commanded” in different ways, again, without being given any clue as to the differences... But even more importantly, what *concrete* protective action am I being commanded with that command <protect your head> so that I *can* act on it -- with the kinds of chronic headaches relevant here? The issue is not satisfiability but intelligibility. I don’t understand the command -- if it is a command.

Menstrual pains.

Similar remarks apply to Colin’s discussion of menstrual pains and dysmenorrhea. He traces the source of these to sharing the same peripheral mechanisms with the pain of child birth. Then he writes:

“The imperative that constitutes them is of roughly the same sort as the imperatives of childbirth— something about protecting the uterus, say.” (p115)

Even if we forget the qualitative differences between these kinds of pains which have, let’s

grant, something to do with uterus, I am not sure <protect your uterus> is an intelligible command to be issued by the body to an agent who could use some help in figuring out, even roughly, what to do with or to one's uterus to protect it. Again, the issue here is not just satisfiability, but rather the lack of sense about what it is to satisfy it -- I don't know what the members of the set of satisfaction worlds are that make up the content.

Visceral pains.

This sort of complaint can very easily be multiplied with all sorts of visceral pains. What are the more *specific* protective actions that are being commanded directed towards a particular region in the viscera and how to explain the qualitative differentiations among different pains roughly in the same area? Is this a pain in my pancreas, my stomach, my intestines, my kidney? What should I protect? How exactly should I protect it? As far as strong transparency is in force, mere differences qua differences will not do. The agent need to know what these different protective actions are that are being commanded so that she can act differently on this basis. But again, more importantly, I have no clue what it is to protect what organ?

Keep in mind that Colin is not talking about secondary motivation due to hurting aspect of pains. This is a different story. I will get to it in a moment. Let me remind you again what Colin says:

“Such arguments [from lack of transparency and appearance-reality distinction] lose their grip if we think of pains as having imperative content. What am I aware of when I introspect on a pain? Surely at least three things: where the pain is, what I should be doing with that body part, and how urgently I should be doing it. Those are all features of the *content* of pains. ... When I look inward, those features of the command are precisely what I am aware of.” (pp124–125)

I've argued that transparency arguments, when properly understood, don't lose their grips when it comes to imperative content. In fact, they have a stronger grip. I suppose Colin might still insist that my body does in fact issue two distinct protective commands in the above two sting pains case, but it's OK for imperativism that I can't decipher what the difference is, when I introspect my pains. But with this move, we are no longer talking about imperativism as a species of strong intentionalism.

If pain sensations in bodily parts were nothing but commands to protect those parts in a certain determinate manner, then, given strong intentionalism, one would expect to see pain reports to be regularly reporting these imperative contents as such. But we rarely, if ever, see that. Compare this to reporting urges or strong desires to act in certain ways. You can't just report an urge or a desire, period. It is always an urge or desire to do something. You report necessarily *what* the desired or urged action is. If pure imperativism as a version of strong intentionalism were true, we would expect to see pain reports to be at least partly but essentially reports of commanded protective action. We don't see this at all. We simply report the pain as having a location, intensity, quality, and especially, its badness. What

specific action it “commands” is not specified, and when asked, many times it’s anybody’s guess...

I conclude that Colin don’t have the resources to handle the myriad ways in which pains can differ in their qualities in a way that will preserve strong intentionalism, and that many times pains are not even intelligible as protective commands.

The “T” in LTQI: Temporal properties

Let me now take up how Colin wants to handle some of the temporal features of pains. Colin says:

“Some of the sensory descriptors are qualities that a pain can have at any particular instant. Call these the *sensory qualities* proper. Others refer to spatio-temporal patterns that pains exhibit. Call these *pattern qualities*.

Pattern qualities are had by pains in virtue of their variation in intensity and location over time. Pattern qualities are picked out by descriptors such as “flickering” or “shooting.” The former picks out a pain with a certain kind of temporal variability: it makes relatively minor changes in amplitude over time at a relatively high frequency. The second picks out a pain that moves location quickly...

Pattern qualities are not part of the content of pain. Nor need they be. We can account for them by postulating certain more basic qualities of pain plus a suitably introspective agent. The reason that you are aware of a shooting pain is because you (a) have a pain that varies, and (b) are aware of that variation. So being aware of a shooting pain is being aware of a fact *about* a pain, not being aware of a quality of the pain. This might appear to be in tension with my intentionalism. But intentionalism is perfectly compatible with introspective facts about mental states. Every intentionalist thinks that you could be aware that one sensation occurred before another, and that this fact needn’t be traceable to the content of either sensation. All you need is two sensations, plus appropriate introspective access. The present account of pattern qualities simply extends that plausible idea.” (pp96–97)

I do sometimes get a shooting pain in my legs, especially in my right leg. Let me try to describe my experience (its phenomenology). The pain is sharp, bright, almost as if something is cutting from inside, and burning, it starts deep in my right buttock, and quickly moves (within milliseconds) down to my thigh and ends in about the back of my knee. It jumps, shoots down. I feel the pain moving rapidly. Sometimes, this repeats through several seconds...

According to Colin, the model is something like this. At the instant at which my pain starts deep in my right buttock, I have a pain with a certain quality and intensity. Suppose it commands at t_1 :

<protect *that* buttock in manner M>

Never mind what fills the 'M' position here (I have no idea whatsoever!). At t2 (say a few milliseconds) after, I have another pain of the same quality but slightly down my buttock (call this spot) b2, which commands me:

<protect *b2* in manner M>

At t3 (a few milliseconds later), I have another pain of more or less the same quality felt deep in my upper right thigh, which commands me:

<protect *that* upper thigh in manner M+>

As the location is different, slightly different set of protective actions is in place -- this is what notation M+ is meant to capture (without changing the sensory quality). Again, I have no idea what M+ demands me to do (especially, say, in a way different than the previous command). Never mind. We can keep going: t4 ..., and t5..., and so on. Se we have got:

t1 <protect *that* buttock in manner M>

t2 <protect *b2* in manner M>

t3 <protect *that* upper thigh in manner M+>

t4 <protect *that* middle thigh in manner M++>

t5 <protect the back of your knee in manner M+++>

t1–t5 stretches about at most 100-150 milliseconds, maybe less. According to Colin, I have got five distinct sensations, and I am introspectively aware of these five sensations as occurring one after the other. It is this model that is supposed to explain my phenomenology of feeling the same pain shooting down my leg.

First problem is, as mentioned above, how do you keep the sensory quality to be the same if the values of 'M' keep changing due to different locations?

Second, the phenomenology. Compare this to your seeing a fly moving across your visual field from right to left in good light within milliseconds. A fly looks to you to be moving across. The movement *of* the fly is what your experience informs you -- it seems properly to belong to the content of your visual experience. It is a dynamic property *of* the fly that your visual experience carries information about. But we can imagine someone insisting that this is an illusion, motion properties are not proper part of the content of visual experiences. Rather, they are to be explained by the subject's awarenesses of a series of distinct visual sensations with the different position information in them. It may be that the underlying subpersonal mechanisms may contain processes that are computations over

static snapshots of scenes -- this is an empirical question. This is fine -- we can still claim that this a subpersonal account of how the visual experience manages to carry information about the motion *of* the fly. But we are not *introspectively* aware of our visual experience as being constituted by a series of distinct sensations with different position contents.

Colin's model of the shooting pain dictates that our awareness of the shooting pain consists of our introspective awarenesses of a series of distinct sensations/commands. This is simply not believable.

Again, consider I have a throbbing toothache -- consider the frequency to be in the milliseconds. According to Colin, my awareness of my throbbing toothache is the awareness of two or three rapidly repeating commands:

- t1 <protect your tooth in manner M >
- t2 <protect your tooth in manner M+ >
- t3 <protect your tooth in manner M >
- t4 <protect your tooth in manner M+ >
- and so forth...

M and M+ are place holders for different protective action types that are meant to explain more basic sensory differentiation within the same ongoing toothache. Does anybody have any idea what M or M+ might be? (Also how plausible is that two different commands are given within milliseconds alternating? Clearly if these are commands, they can't be obeyed differently by the agent...)

But the basic problem with this picture is that my awareness of the throbbing pattern quality is explained in a different way than my awareness of the different sensory qualities, call them M* and M+*, corresponding to M and M+. When I am aware of M*, I am aware of M (as being commanded). When I am aware of M+*, I am aware of M+ (as being commanded). Then, my awareness of the throbbing *pattern* of my pain is my awarenesses of these two continually alternating commands -- that *they* alternate. So, my awareness of pattern qualities is always introspective -- unlike, presumably, my awareness of the sensory qualities, which isn't supposed to be introspective: recall that when you are aware of your pain, you are not aware of any noncontent-based intrinsic feature of your sensation. In other words, your access to basic sensory qualities is access to pure imperative content defined over centered satisfaction worlds without anything psychological in the center. You are aware of only the content of the specific command: what specific protective action you should take towards which part of your body with what priority. Nothing else. But if Colin's account of pattern qualities is right, my epistemic access to the throbbing pain quality itself is quite different than the epistemic access to the basic sensory qualities I have. This is, quite simply put, incredible.

At this point, why not just say that pattern qualities like throbbing etc. can be explained by positing yet another set of distinct protective action types. It is not like we don't want to posit action types that are different without knowing what the differences consist in. It looks like Colin is committed to positing these anyway... As complained before, the view with these sorts of commitments is no longer a strong intentionalist story -- it looks like our access to pain sensations isn't access to mere content...

Again, so much for strong intentionalism and the strong transparency it implies.

SECOND-ORDER IMPERATIVES (AFFECT)

Let me now get back to the affect and secondary motivation. According to Colin, the painfulness of pains is a response to our pains. These are higher-order imperatives directed towards the pain sensation: <stop *that* sensation!>. One immediate worry is that the structure of these imperatives will not allow qualitative variations among the different ways a sensation feels bad/unpleasant/etc. Colin may be thinking that the unpleasantness or suffering doesn't have any internal differentiation -- it's just a negative valence with an intensity scale. This may be true for primary affect (not sure), but it doesn't seem to be true for secondary affect or suffering properly so-called.

More worries about these second-order imperatives.

Do we deal with ourselves in terms of commands? These commands are issued by us to ourselves. We command ourselves to stop certain sensations. These are agential commands. Colin writes:

“I suggest, second-order imperatives might be issued by the *agent*. That is, suffering is a second-order command that *we* issue to *ourselves* because we have judged a first-order sensation to be bad in some way. Agential-level commands, in this sense, would still have to be relatively quick, encapsulated, involuntary sorts of states—for again, we can't talk ourselves out of suffering. Note too that this wouldn't be a pure form of imperativism: suffering would consist of both an evaluation of a first-order state and a command with respect to that state” (p187)

It is a little hard to follow what is going on here. If these agential level stop commands are “relatively quick, encapsulated, involuntary sorts of states”, why do we need to base them on prior agential level evaluative judgments? A more basic concern: if these commands are “quick, encapsulated, involuntary sorts of states”, in what sense are they truly agential? The painfulness of pains attach to pains, they qualify pains (something Colin seems to agree). They are presented to us as qualities of pains. But pains are peripheral first order commands.

[[Cf. Colin elsewhere: “Nor does it seem as if *I* issue imperatives to myself. More precisely, pains are not personal-level self-commands. If so, they would be either voluntary or

involuntary. If I voluntarily issued imperatives to myself, then I could presumably stop doing so when imperatives were inconvenient. Pain is not so easy to overcome. Pains come and go of their own accord, and they persist even when I would like them not to. So they aren't like self-commands. Nor do they feel like a kind of involuntary self-exhortation (say, as compulsions do). The perceptual [sic.] nature of bodily imperatives makes them feel much more akin to something done to us, rather than something we do to ourselves." (p75). *But this is precisely how the hurtfulness of pains strikes us, no? ...]]*

What makes pain sensations bad, so that we can correctly judge them to be bad and command ourselves to stop them?... I am not sure I am in good control of the material in this last chapter. But Colin seems to claim that what makes pain sensations bad is not how they intrinsically feel, but rather the badness of their intrusive consequences on the free expression of one's autonomous agential powers. But this is confusing. Remember that pain commands have an intensity function built in. What this does is to mess with one's preference functions from inside by forcefully reordering one's priorities. If this is what intensity of pain is, and if this in turn just consists in the intrusive messing of one's priorities, and if *this* is what *is* bad; then it looks like Colin has already built the negative affect (painfulness) into the pain sensation itself. He doesn't need the agent independently confirming it and judging it to be bad so that the agent can start issuing stop-commands. But then it is simply not true that painfulness is only contingently connected to pain sensations, on Colin's way of carving up the theoretical space. You can't have pains without negative affect in so far as pains come with an intensity built into the content of pains.

I myself welcome this way of looking at things. I have argued in the past (with Matt Fulkerson) that the sensory affect is a matter of more or less mandatory functional processing of incoming sensory information in relation to one's built-in (or otherwise) priorities, reordering of preferences, pre-motor biasing and planning, etc. It isn't that pain's painfulness is bad *because* these agency-limiting *consequences* are bad, rather: these sorts of effects of sensory processing *are* what metaphysically *constitutes* the pain's feeling bad/unpleasant. But this is a weak-intentionalist psychofunctionalist story... Describing the system in terms of imperatives is useful but not necessary...

But perhaps Colin wants to say that the intrusiveness of pains is just bad in itself but pains don't become unpleasant/painful until the agent recognizes pains' badness and start issuing stop-commands. But this has the weird consequence that pains don't hurt because they are bad, but they hurt because they are recognized as bad or evaluated as bad. If this is the picture, it makes Colin a very close ally to second-order desire theorists or second-order evaluativist. One of the main troubles on this picture is that pains' being bad doesn't start to *feel* bad (feel unpleasant) until evaluated as such. But this is implausible unless it is a metaphorical description of mostly subpersonal functional-causal affect processing systems. I have criticized such views elsewhere but let me give you the gist of it here by asking:

What is exactly the referent of ‘*that* sensation’ in the second-order stop command?

If we are literal about this command talk directed toward a sensation, we need more details. These commands are agential imperatives directed towards sensations. But this means that for an individual to experience a painful pain, that individual needs to be able to refer to his/her sensation (this is an introspective command). Colin writes:

“Second-order imperativism is flexible enough to include babies, animals, and other cognitively limited creatures. The general form of the second-order imperative can be shared between them and us: when I suffer and when a dog suffers from pain, we both have a first-order imperative and a second-order imperative directed at the elimination of the first-order one. The capacities I can use to satisfy the second-order imperative are more flexible than those available to the dog, however (I know where the ibuprofen is, and he doesn’t).” (p187)

So, dogs and babies can refer to their own pain sensations and command themselves to stop *them*! In order to make this story plausible, Colin will probably make the same move that many other attitudinal theorists before him have made: make the commanding involuntary, automatic, largely hard-wired, primitive, etc. The more you elaborate this sort of story the more it will start to sound like the description of a causal/functional mechanism where the description of its states as commanding is mostly metaphorical. (This is not to say that the states of some such system can never be non-metaphorically described in intentional or information-theoretic terms -- I am not arguing against any sort of intentionalism here.)

Similarly, on a strong intentionalist story, commanding oneself to stop a sensation requires being aware of that sensation, and that in turn requires being aware of the content of the imperative content *only*, according to Colin -- no features of the vehicle of the command are supposed to be involved in one’s awareness. But it is really weird to think of the second order stop command being directed at another command we are aware of only as a mere content (a mere set of satisfaction conditions). Note that there is no parallel of this sort of commanding in ordinary commands. In those rare occasions where we might find ourselves commanding to stop another command issued to us, we are always aware of the vehicle of the first order command even when we know what the satisfaction conditions are for stopping *that* command. But to stop *it*, you have to act to deal with the vehicle -- there isn’t much else you can do to deal with the content directly even if you know what that content is (indeed it doesn’t even make sense to act towards the content bypassing the vehicle). But if we transfer this to introspective commands, it must be the case that we are aware of vehicle properties of the first-order command so we want *that* act of commanding to stop. If so, transparency is violated again as we must be aware of some noncontent-based intrinsic properties of the first order command in issuing the second order command to stop *it* (even if we are also aware of the content).

“I” of LTQI — INTENSITY

How intense a pain is is a matter of how that pain ranks its own satisfaction relative to other options (satisfaction of other conative states of the agent). Two things: 1) coming up with the ordering itself; 2) securing compliance with the ordering in execution.

(1) requires that the pain has access to what other conative states the agent has: for a token pain *p* to have intensity *i* is just for *p* to have a certain ranking of various priorities of the agent done by the pain (?). This seems to preclude agential justification of this sort: “I prioritize doing *x* over attending to pain because the pain is not that intense” — this sort of explanations/justifications would not be informative under Klein’s proposal. On the face of it, pain intensity enters our practical deliberations as an independent factor. But on Klein’s view, pain intensity seems to be already an output of some sort of practical deliberation. Furthermore, if intensity of a token pain essentially requires access to other (some/all?) conative states, then intensity is a holistic matter that can be changed by adding/deleting other conative states. Is pain intensity *that* pliable? A lot of pain scientists talk about pain intensity as a function of the stimulus intensity. Colin talks about pain’s ranking their own satisfaction conditions relative to others. But pain sensations are supposed to be peripheral commands issued to the agent’s deliberative processes. It is odd to think of pains doing the rankings themselves as opposed to the agent’s doing the ranking. [[There are also oddities about the interaction between the intensities of first and second order commands -- elaborate later...]]

Pains in having different intensities *force* compliance to varying degrees. It’s one thing to have the ordering, another to secure compliance. Securing compliance doesn’t seem to be a matter of content, though. It seems to require having causal influence on execution of behavior and deliberation. This is causal/functional structuring of ongoing activities. (Consider the difference between a mere proposal and an urge.) This is in fact a general complaint about imperativism -- merely issuing protective commands will not motivate unless it is bound to some executive mechanisms that will force compliance or meddle with execution/deliberation to varying degrees.

ACCEPTING BODY AS A MINIMAL PRACTIAL AUTHORITY

Indeed, Colin does appeal to accepting the body as a minimal practical authority to give the commands motivating and justifying powers. According to Colin, we have good motivating and justifying reasons to act in ways commended by our body because we accept our body as a legitimate minimal practical authority. He writes:

“We accept our bodies, I suggest, as minimal practical authorities. Our body commands us to protect a certain part. Because we accept our body as a practical authority, that command gives us reason to act—regardless of what else we’d want to do and regardless of what else we know. The reasons that pains give are ones that we continue taking seriously even when we want to do otherwise and even if we

know that the body has made some sort of mistake. So they give us both a first-order reason to act so as to satisfy the command and a second-order reason to make that first order reason a mandatory part of our deliberations...” (p80)

The attitude of accepting the body as a practical authority is not a personal level attitude that one can change. We are hardwired to accept our bodies as such authorities. Many times, Colin talks as if it is really a personal level attitude:

“we accept our bodies as authorities for good reason. Our body is important to us. We care about it. We are in bad shape if it doesn’t work. We cease to exist when it does. ... there are many good reasons for accepting the authority of the body” (p81)

But he immediately adds:

“Note, by the way, that “acceptance” here shouldn’t be construed as something like voluntary acceptance. It’s not as if we first deliberate and then decide, all things being equal, that our body is worth listening to. There are powerful evolutionary reasons why the attitude of acceptance should be innate and difficult to overcome” (p81)

So, talking about acceptance is a *façon de parler*, not a lot more meaningful than talking about acceptance of the effects of gravitational forces on our bodies. Similarly, with the talk about body care -- to which I will get shortly. On evolutionary good reasons, we are simply built to respond to the various demands of our body. So, pain commands motivate because we are built that way -- we can’t help but respond to these commands.

But pain commands, even if they may provide motivating reasons to act in certain ways in particular situations (I doubt this -- but put it aside), I doubt whether they can provide *justifying* reasons for acting in the ways commanded by them in certain *particular* situations. Take the allodynia case Matt and I discuss in our paper on reasons and theories of sensory affect. Sally is suffering from allodynia in her arm due to a burn, a gentle doctor’s touch to remove the scar tissue causes her to experience extreme pain. She withdraws her arm vigorously. I claim that she has absolutely no good justifying reason to act in this protective way. This sounds bizarre simply because we know why she does withdraw her arm -- to stop the painful pain sensation. But this is different. Talking about pain commands providing reasons is therefore ambiguous. Evolutionary accounts that explain why a system came to exist in a species cannot provide normative reasons in the way Colin wants for *particular* actions. If this were not so, the Planned Parenthood would by default be an intellectually confused and morally bankrupt organization.

BODY CARE

It turns out that accepting the body as a minimal practical authority is based on our caring

our bodies. For Colin, body care turns out to be an enabling condition for pain commands to actually motivate agents. Colin's discussion of pain asymbolia leads him to postulate intact body care as an enabling condition for the protective commands to motivate the agent who cares about her body. The attitude of body care turns out to be a fairly complex one. It is not a personal level attitude that we can directly control. Colin doesn't say much about this in his book. But he says a bit more in his response to Frederique de Vignemont:

“Basic [body] care manifests as a set of appropriate *dispositions* towards the felt body which are sensitive to goings-on in the surrounding peripersonal space..., current homeostatic demands ..., facts about which body parts belong to the organism, and facts about what can be done with those parts.” (2016, p2).

This is fairly brief and hugely incomplete. The most important thing that is missing here is the articulation of what makes these dispositions appropriate to body *care*. What is it about this complex but mostly hard-wired system that makes it a *care* system whose object is the body? This system must have built-in routines to interpret certain incoming sensory signals as a *threat* to the integrity of the body and prepare to do whatever it can do to *ward it off*. But note that this implies having certain built-in routines or states that are *conative*. You ward off threats only if you *want* threats to disappear, *ceteris paribus*. In other words, some of the hard-wired dispositions to act or process information must have a conative direction of fit to the world. Suppose your body registers the occurrence of potentially injury-causing stimuli on location L, your body in its own basic/primitive way doesn't want to sustain injury -- this is part of basic body *care*. This translates into *your* not wanting to sustain injury as an agent -- you care about your own body in this basic way. According to Colin, your body then issues a command to act in a certain way to protect L -- and you obey because you care about your own body in this basic sort of way.

My complaint here is that the picture is unnecessarily complex. Suppose we have this basic body care -- this is mostly (but not entirely) a subpersonal system with built-in information processing routines some of which are the expression of built-in conative functions. That is, the system has a built-in teleology manifested in the complexity of its information processing and its interactions with deliberative and motor systems. We can simply express the system's conative states or processing by saying, for instance, that the system “desires” not to sustain injury. Simply put, if I care about my body, in whatever basic sense relevant here, I certainly *desire* not to sustain injury in the same basic sense. If this is correct, then we have the basic framework to explain how I am motivated to act in ways to protect myself from injury.

1. I (my body) have incoming sensory information about injury happening in L.
2. I desire (in the basic body care sense) not to sustain injury.
3. My basic body care system (the set of dispositions Colin mentions) provides the wherewithals of how to ward it off.
4. I act to remove the threat (i.e., I engage in protective behavior directed at L)

Where must, in all this, a need to issue an imperative arise? This pattern fits well with the general belief-desire explanations (except it's at the more experiential/peripheral level). In (1) we have the sensory/perceptual registration of injury (pain). In (2) we have the negative valence revealed when the sensory information enters into the system with built-in anti-injury conative states. In (3) we have the body-matrix providing the fine-tuned resources about how to efficiently deal with the current threat. Finally, we are ready to act, and most often we act right away...

Although its defenders have more or less ambitious metaphysical agendas, the above is more or less the kind of picture that lies behind the Received View. The insistence on keeping the protective commands not even connected to affect, especially after postulating basic body care to make imperatives work, seems to me not theoretically well motivated. [[Btw, I have the same basic complaint about David Bain's evaluativism -- once you postulate a basic body care to explain the motivating power of evaluative representations, you make representationalism otiose...]]

I myself don't quite understand why Colin wants to run his imperativism on LTQI insisting that pain sensations don't carry *any* truth-apt information about what is going on in the periphery where the protective action is directed. There is plenty of empirical evidence to think that there are relatively specialized peripheral and spinal sensory afferent systems in processing noxious stimuli. Yes, the debate between the Specificity and Pattern Theory is still going on. But as far as I can tell, there is a consensus that the truth seems to be somewhere in the middle. But specialized or not, Colin needs the body-matrix to have the necessary sensory information to direct its imperative activities -- after all, one of its jobs is to collect sensory information from multiple resources. You can't expect the body to make guesses about where the help is needed. Indeed, throughout the book, Colin himself makes plenty of reference to peripheral and spinal *sensory* mechanisms to explain facts about pain sensations/commands. The analogy with fire alarms is nice, but has limited use when applied to flexible ways in which an agent can deal with pain causing stimuli. Colin's argument about when signals carry imperative as opposed to indicative content is not convincing when one realizes that protective action needs all sorts of more determinate manners specified in the imperative content. This creates exactly the same sort of degenerate relationship between the signal and its upstream causes on the one hand that Colin complains about the indicative representationalists, and the signal and its downstream effects on the other. Given the flexibility of pain behavior, pure imperativism doesn't have any significant advantage over indicative forms of representationalism.

The most natural home for imperativism, as far as I am concerned, is in addressing the affective dimension of sensory experiences including pain. Indeed, this is how Manolo Martinez runs his imperativism. He introduces commands to explain the affective-motivational dimension (+AM). He is in fairly good agreement with representationalists, such as David Bain and Michael Tye, in his account of the sensory-discriminative component of pain (LTQI). In fact, when Colin published his 2007 paper, most of us party

to the debate took Colin to be giving an account of the affective-motivational aspect of pain and claiming that there is no sensory discriminative component to pains. The apparent change in position in his book came to me as a bit of surprise.

On the other hand, I am inclined to interpret Colin's particular take on the intensity of pain as part of imperative content as a particularly apt way of bringing the affective-motivational (+AM) component back into the experience of pain. Intensity turns out to be nothing but a relatively hard-wired way in which the pain information is processed vis-à-vis the agent's other occurrent and standing conative states. This is affective processing that is best given a functional-casual description of the way nociceptive information is processed to set motivational and learning parameters. Indeed, I myself developed a view of this kind about sensory affect in general (with Matt Fulkerson). So, I welcome Colin's framework for the intensity of pain except that it's not really about intensity but about affect.

Let me stop here. I enjoyed reading, indeed wrestling with, Colin's book immensely. It's a tremendously rich and stimulating book that we'll keep talking about, I am sure, for years to come.

SOME FURTHER CONCERNS:

[[Maybe to be elaborated later...]]

- Colin tries hard to find examples of motivating pains without hurt in §12.4.2. The cases he lists seem to me to be all cases where if there is no hurt there is also no motivation, or if there is no motivation there is also no hurt.
- The arguments Colin marshals for the contingent connection of pain to hurt in §4.3 are not persuasive...
- The arguments against indicative content in §3.4 can be resisted -- here I would argue that there are difficulties about Colin's own reasons for thinking pain signals have imperative content that exactly parallel his own criticisms of views that take pain signal to have indicative content.