

The impossibility of blaming token people fairly: the problem of demands

A impossibilidade de culpar pessoas particulares de forma justa: o problema das demandas

Eduardo Estevão Quirino¹
 Universität Vechta
 eduardoquirinio@gmail.com

Abstract: This essay attempts to do two things. First, to problematize the relation between obligations and demands. Second, to show that the popular principle of Ought Implies Can and a plausible reading of what it is for blaming to be fair are incompatible with some cherished assumptions to the point of being impossible to blame concrete people, those with flesh and bones, fairly. The argument can be summarized as follows: For a person to fairly blame another subject, they need to be justified in believing both that a) the subject was obliged to act in accordance with the demand associated to blame; and b) the subject was capable of acting in this way. Unfortunately, there are reasons to think that b) is never justified, leading to the blaming itself to never be justified. I try to show that this argument is almost entirely independent of positions on free will, making it the only overall skeptical argument (that I know of) that delivers this conclusion about fair blame and, consequently, moral responsibility to an extent, without involving substantial debates on free will. The essay connects some previously unassociated literatures on Ought Implies Can, blame, the nature of normative demands, Objectivism/Subjectivism about moral obligation, and moral psychology. The conclusion of the piece is not to be an endorsement that there is no fair blame, rather, it claims that these arguments should be taken as a *reductio ad absurdum*.

Keywords: blame; free will; moral epistemology; moral responsibility; normative demands; ought implies can.

Resumo: Esse ensaio busca fazer duas coisas. Primeira, problematizar a relação entre obrigações e demandas. Segunda, mostrar que o famoso princípio que Deve Implica Pode e uma leitura plausível sobre o que culpar de forma justa requer, são conjuntamente incompatíveis com algumas suposições muito aceitas, incompatibilidade essa que chega a ser impossível culpar pessoas concretas, de carne e osso, de forma justa. O argumento pode ser resumido assim: para uma pessoa poder culpar outra de forma justa, ela precisa estar justificada em acreditar tanto em a) o sujeito do julgamento era obrigado a aceitar as demandas associadas com a culpa, e b) o sujeito era capaz de agir de acordo com a demanda. Infelizmente, há motivos para crer que b) nunca é satisfeito, assim tampouco o é a culpa em si. Eu tento mostrar que esse argumento é quase inteiramente independente de qualquer posição sobre o debate sobre livre-arbítrio. Sendo dessa forma (que eu saiba) o único argumento para um ceticismo geral sobre responsabilidade moral que não depende de debates sobre livre-arbítrio. O ensaio conecta áreas antes não relacionadas que discutem Deve Implica Pode, culpa, a natureza das demandas normativas, Objetivismo/Subjetivismo sobre obrigação moral e psicologia moral. A conclusão cética, no entanto, não deveria ser aceita como tal, e sim como uma redução ao absurdo.

Palavras-chave: culpa; livre-arbítrio; epistemologia moral; responsabilidade moral; demandas normativas; deve implica pode.

¹I thank all the people who discussed some of the contents of the present article with me in the presentations I made on it both at Vechta and at Curitiba. Prof. Jean-Christophe Merle, Frank Rettweiler, Tania Eden, Tales Yamamoto deserve special mention.

Recebido em 31 de julho de 2025. Aceito em 10 de novembro de 2025.

Introduction

There have been many controversies about the relationship between most kinds of actions and their respective morality. However, one such kind of action has seen scant philosophical attention, namely, the acts of normatively demanding that others obey, or satisfy, some duties, norms, or expectations. I wish to add something about the moral significance of demanding itself, but because that would be a whole book (at least!) I will narrow down to a single kind of demand i.e. “blaming”, that is, the demand agents make upon others (or themselves) to recognize and act upon their blameworthiness for a given misdeed.

If a person blames another they are making both a claim of fact (that bears justification); and a normative demand (that requires legitimacy). The claim of fact is usually something like: “you are blameworthy for the misdeed”. The normative demand is usually (similar to): “you should recognize your error and act upon it accordingly to redeem yourself”², or “you should be morally reprehended”. This second, normative part explains why I consider blaming an act of demanding.

In this paper, I argue that when it comes to the morality of demands, some unexpected tensions arise from our intuitions about fair blaming and the widely accepted principle ‘Ought Implies Can (OIC)’. The tension is quite strong, in fact, for given some plausible empirical and philosophical assumptions, I show how we can conclude (absurdly, no doubt) that we cannot blame existent, flesh-and-bones people fairly. This conclusion is rather drastic, and should better be seen as a *reductio ad absurdum* against some views—exactly which ones is not clear. I will present a few options of how we could avoid this drastic outcome, and while I leave it opened which is the best, it is important to highlight that all options seem to come at a price. In other words, this argument is fun for the whole community of metaethicists.

For the sake of clarity, in Section 1, I will build my argument over SHER’s (2006) account of blame, however, I do not think that much hangs on this account’s details, for what is problematic about fair blaming seems to be core intuitions any plausible account of blame would have to endorse, or so I argue. In Section 2, I briefly present the discussion over whether the nature of obligations is subjective or objective and whether the answers are compatible with OIC. I also argue for a particular interpretation of OIC based on capability rather than nomological possibilities. Section 3 takes a detour and enters into some empirical claims about the biology of moral human agents. These empirical claims are more illustrative than actually crucial for the argument that follows, and are very general. Much the same argument could be advanced with merely philosophical thought experiments, but I think the more down-to-earth construction adds plausibility and shows that I am not just running after philosophical drivels.

These three first sections are rather brief, being here only to set the stage. Section 4 starts with an account of normative demands and their relation to obligations. Furthermore, I propose a definitions of Fair Blaming. Additionally, it presents the new concepts I introduce: “token/type people” i.e. flesh-and-bones, spatiotemporal people versus hypothetical, generalized people. The crux of the difference lies in the epistemic availability of morally relevant properties in the scenarios under consideration. When we are talking about type people, we can stipulate *all the morally relevant properties*, thus philosophers frequently build very precise examples to capture exactly some moral properties ignoring others. For token people, this is not so. It seems that

²Different theorists will have dispute which particular demands are being made by blaming. The one I chose is my own and a similar one, by SHER’s (2006), will be discussed in section I.

the morally relevant properties in situations involving token people are a contingent, partially empirical matter. This section also explains how OIC is connected to the argument. Section 5 is where it all come together in the two *Epistemic Arguments Against Token Blaming*. Section 5 is the core of this paper. Finally, section 6 briefly explores what can be done about those arguments, showing some ways out the skeptical conclusion and their costs.

These arguments against fair blaming are the only arguments I know of for generalized skepticism against normative responsibility (epistemic and moral included) that has no dialectically relevant connections to free will or determinism. Other arguments are either locally skeptical, or depend on certain views on free will/determinism. My argument does depend on a certain view of free will, but it is a negative dependence (it depends on X view not being true) and X in question is such a wild belief about free will that not even Cartesian Dualists would think plausible. Hence, if not for anything else, the arguments might be interesting for that reason.

1. Blame and Blame realism

We start our discussion with a brief sketch of SHER's (2006) account of Blame. His account is very influential and is still well-taken, which partially justifies the narrow scope of the discussion I wish to advance. A methodological problem for me arises from one of the reasons his account is so distinctive: Sher is a masterful structural writer, such that each point of his is carved by the previous and lends itself naturally to the next. Although a beautiful piece of philosophy, this makes my work as a reconstructor of his account somewhat lacking, unless I devote myself to a long piece, which I cannot. Hence, the schematic version of his views I wish to advance should be taken more like a way to focus this discussion, rather than a proper reconstruction of his arguments.

With these remarks behind us, we can start exposing Sher's account. His work is divided in two parts. The first, chapters 2-4, are dedicated to answering the question of what justifies blamers (people ascribing blame) to blame subjects based on their past wrongdoing. The second part, chapters 5-7, hopes to answer more directly the question of what blame is. I will divide my presentation accordingly dealing with those early chapters in 1.1 and the later ones in 1.2.

1.1. Justifying Blame

Sher's arguments in part one start with a discussion about what justifies the transfer of opprobrium from a bad action to a bad actor. In Sher's view, genuine (moral) blame is only directed towards agents; whereas the reasons from which we ascribe blame are based on our reprobation of bad actions. Because we usually judge others based on their *past* actions, it is not immediately obvious why we should take action's badness to impinge upon the *current*, existing agent. If blame is to be justified, we must account for this apparent asymmetry (SHER, 2006, p. 7). He then follows up suggesting that one of the best extent proposals to deal with this problem is the Humean thesis, which is roughly as follows: we can infer from bad actions to the defect in an agent's character; and, because of this, we can blame the agent bearing the defective character for such. In other words, bad actions justify the blame only inasmuch as they offer good reasons to suspect bad character, and bad character is the main problem all along.

Sher's second chapter engages with this proposal and argues forcefully to its dismissal. He offers three examples, each of which brings forth a problematic objection. I will focus on the one I find most compelling. The account above is committed to the source of blame to be the badness of

character evidenced by bad actions. However, assuming a thoroughly evil person who has an opportunity to enact their evilness— cruelly taunting a child — the Humean account would suggest that the taunting is a legitimate reason to think that the aggressor deserves blame. All good. However, Sher points out that this is a problem, because the taunting only produced evidence for the claim that the aggressor was evil. Given that the person was already evil before the taunting, even if he had not acted in that way, he would already be blameworthy – his character being prior to his actions. But, Sher argues – correctly to my view –, that this is not acceptable. We should only be allowed to blame people after, and *because of*, their bad actions. (SHER, 2006, p. 30-31).

The Humean position being defeated, Sher rushes to point out which parts of it made it so plausible. He highlights two broad claims he thinks we should bring to bear in a plausible answer to the initial question. Claim (1) is that a person's *character* is uniquely well-situated to link the responsibility for the bad action to the agent's desert for blame. Claim (2) is that someone's bad behavior need not exhaust all the things for which one can be blamed (SHER, 2006, p. 7). The base reason for the first claim is that character, seen as a complex mixture of beliefs, desires, emotions and dispositions, is both co-occurrent with the agent, and causally responsible for their actions. This combination is rare enough, but his account will also benefit from another trait of characters: they seem to be especially connected to our self-identity: the thing we refer to when we say first-person pronouns like "I", "me" and "myself"³. More on this shortly.

In what comes to the claim (2), Sher's most sophisticated, and difficult to defend, position in the book is to maintain that we might be blamed for things we have no control over. Not actions, of course. If we have no control over an action, we have no blame for it (with plausible qualifications). But character traits are things for which, Sher argues, we can be responsible for even if we never have power to choose them. The argument is long and controversial (abducted in SHER (2006, chapter 4)). In particular, it is based on a reinterpretation of what it means for blame to be someone's.

He starts arguing that our self-identity is conceptually connected to our character. After that, he contains that our relationship to our actions is not conceptual, but causal. Thinking diagrammatically: we have a conceptual arrow connecting the self to the character, and a causal arrow connecting the character to the action and the follow-up outcome of this action.

The very popular condition of control for responsibility ensues regarding action, but not regarding our own character traits, because the ways in which actions are ours; and the way in which our characters are ours are different. That is, we need control over whether or not an action (or omission) happens because the connection between us and the action is only causal. If the causal link is broken, the action is no longer ours: it no longer bears on us. However, the link between us and our character is much more intimate. Our lack of control over how our character ended up being as it is does not touch upon the fact that it is *still our character*. In a nutshell: If it

³It is important to differentiate this from logical notions of identity, about which one may find PARFIT's (1971), and WILLIAMS' (1970) papers very insightful and still influential.

can be traced to us, we can be blamed for it (SHER, 2006, p. 58)⁴. This is clearly very polemical, yet so, I will grant it for the sake of argument⁵.

To sum up the answer about justification: we wanted an explanation of what, if anything, can tie up the agent and their action such that their bad actions can lead us to blame them. The answer is that, differently from Humeans, Sher's account goes from the badness of the action (thus requiring it to have occurred, blocking above's objection) to the cause of the bad action, namely the vastly complex array of (among other things) beliefs, desires and dispositions that constitute the character of the person, which by its turn constitute the person. Hence, the bad action allows for the blaming of a given subject, because it was that subject's character that caused it.

1.2. The Nature of Blame

This subsection is more directly relevant to our discussion. Sher's answer here depends on the truth of the previous claim, to the effect that it is that connection that a person trying to blame must establish: if I want to blame you, I must believe that you (your character) is connected causally with the action for which I blame you. Hence, the first part of his account of blame is that blame involves *a belief* that the person being blamed acted contrary to the (presumed adequate/legitimate) standards, or that they have a character inclined to do so (SHER, 2006, p. 95). To narrow our attention in this discussion, we will only talk about moral standards, consequently, blame includes the belief of moral failure. However, this is not sufficient for a number of reasons associated to our everyday experiences. People might believe that the other person acted contrary to standards, but not blame someone, instead, the person might praise them for it, because the blamer might be against some widely accepted moral standard. To complement the account, then, Sher suggests that blame is a pair of dispositions, not any single one. The pair is the belief above added the desire that the blamed person had not acted in that way (or had not had that character) (SHER, 2006, p.103).

After suggesting his favored pair of dispositions, Sher shows that they can account for some desiderata for accounts of blame. His desiderata are basically that which we associate with blame in everyday life. Given that this whole discussion emerges from STRAWSON's (1962/2008) *Freedom and Resentment* paper, it is unsurprising that "reactive attitudes" take a major part in establishing what matters about blame⁶. I think it is rather straightforward to see that beliefs and desires suffice for most reactive attitudes, but two common attitudes we have associated to blame are less obviously connected to that pair i.e. apologies and reprimands.

Hence, to conclude this brief account, I wish to explain how Sher deals with these seemingly problematic cases, and make some claims about the metaphysical and epistemological implications of his account. So, regarding the first point, Sher argues that moral desires of the form above are distinct from the average desire one has because, given that morality is a public affair, externalizing our desires that morality should have been uplifted we present ourselves as moral, while, at the

⁴In fact, this formulation is similar to Sher's truism: "that people can only be blamed for what reflects badly on them" (SHER, 2006, p.58).

⁵But see SMITH (2013) for a criticism of his whole account, and SCANLON (2008, 2013) for a different account (also criticized by Smith).

⁶COATES & TOGNAZZINI (2013) explain the relation with Strawson's work and explore it in a helpful way. HIERONYMI (2020) offers a robust introduction and (re) interpretation of that complex piece.

same time, we avoid taking part on (or seeming complacent with) the transgression. The belief part is necessary for those two attitudes because on one hand, it explains why we apologize or reprimand. On the other, the belief that the person has acted poorly is a necessary part of a genuine apology or reprimand. If I apologize without believing that I have acted badly, I am faking my apology. If I reprimand you without believing that what you did was really wrong, I am merely scolding you. Sher's discussion goes on a while after this, but for the sake of brevity, I think this explanation suffices to show the gist of the solution.

Now, let us see the metaphysical and epistemological implications of his account. To do that, I must bring his general slogan for what a bad action is supposed to mean. A bad action is an action that acts against or ignores good moral reasons. If this account holds, then the belief associated with blame is the belief that the agent under evaluation has acted in ways that are contrary to good moral reasons either for ignorance or for inobservance of such reasons. The belief, therefore, has two main components: "X acted in a bad way" and "X ignored (culpably), or went against good moral reasons". This means that reasons against blame could operate by undermining either of the conjuncts. The first strategy would be the "Alibi claim" and the second the "Excusing claim", undermining the first and second conjuncts, respectively. The Alibi claim, naturally, is the strategy of arguing that it wasn't this person who did it. The Excusing claim, on the other hand, accepts that it was the person in question but wishes to point out that the bad action does not suggest moral unconformity. Because of this, the principle that Ought Implies Can (OIC) is a kind of Excusing claim. If it is invoked, it works by claiming that the person who did the bad thing, could not have obeyed the moral rule in question. This will be important later, for the excusing happens objectively and thus we might not know that someone was excused.

Insofar as the epistemological implications are concerned, the issue is more connected to the belief component. This epistemology of blame has two connected parts. The first regards the truth of the belief that the person being blamed is truly causally connected to the bad action (in adequate ways). Additionally, one must believe *justifiedly* that there are no sufficient excusing factors that mitigate, or fully exempt, the actor from blame. Fairly straightforward.

Now, for the metaphysical implications. As for the desire component, Sher argues that the desire we have that the person who broke a moral commitment had not done so (or that their character should have been better inasmuch as this moral commitment is concerned) is *constitutive* of fully accepting a moral stance. He argues, compellingly, that one cannot fully accept a moral commitment, stance, law, or etc. without at the same time desiring others (barring excuses) to follow it.

Regardless, I think the desire is associated to a normative demand too, because the desire Sher evokes has this constitutive role in making people part of the moral world. Desires, by themselves, seem to be the wrong kind of entity to explain this constitutive role entirely for they seem to be overly action-commanding, whilst moral acceptance is passive in a way: most of the time, the person just move along with it as if the moral norm was a piece of furniture.

Whereas the desire *plus* the demand could accommodate better the passive parts of fully belonging to the moral realm in a relatively minor fix, I think, moreover, that the demands for others and ourselves to follow the norms is *constitutively prior* to the relevant desire. The desire is only evoked *after* the moral norm was violated, and therefore cannot explain why the moral norm

exists in the first place. Hence, I think that the correct pair of intentions associated to blame is a belief and a demand, not a desire.

There is another metaphysical component here, that is tied to the beliefs, namely, their truth-conditions. Sher claims that our beliefs about the person have to be *true* about them if we are to blame appropriately, which suggests that blame has a fact of the matter. A very complicated one at that, given that the truthmaker associated to blame attributions corresponds to the existence of complex mental states causally connected to the bad action. Complex it may be, but real nonetheless. This realist position says that when a person believes others to be blameworthy, the type of belief is like any ordinary propositional belief. There is a fact of the sort “the person is blameworthy” that, if actual, would make my believing as much true.

In other words, anyone who posits a relationship between real states and blame, or between real obligations and blame, will be ascribing to a view that is committed to there being a fact of the matter that makes the claim “S is to be blamed for ϕ in t ” to be either true or false. Adding OIC, we have the following:

BO (Blame for failing Obligations) If S is to be blamed for ϕ in t , then they ought to not have ϕ -ed in t .
Which implies (conceptually) that:

CC (Core Condition): If S is to be blamed for ϕ in t , then S could had not- ϕ -ed in t .

Focusing on Sher’s account helps us see that CC is very plausible, but a glance to the very proposition should suffice to make much the same point. CC is the first assumption of the argument for generalized skepticism that I will advance. The next two are defended in the following sections.

2. Capability and Objective Obligations

Most accounts of blame are realists in the sense above. Now, it would be beneficial to think about what that entails for the nature of obligations. The reason being that CC only works in this realist way because the bridging principle appealed to, BO, applied an objective account of blame and obligations, such that whatever contradicts an obligation is blameworthy. This is seen as controversial for many. As I hope to build my argument on premises many hold, I should clarify what is the difference between subjective and objective ought, as well as point out why the later is to be preferred.

According to VRANAS (2007) and GRAHAM (2010), subjective obligations are those that (somehow) depend upon one’s epistemic position regarding them. Objective obligations are those that do not depend thusly. Vranas articulates his influential account of OIC in terms of objective ought. I am objectively freed from an obligation if I (really) cannot dispatch it.

2.1. Possibility and Capability in OIC

‘Cannot’ here means the lack of capability, not the absence of possibility⁷. Although capability implies possibility, the converse is not true. Frequently, it is possible for a subject to do something while, nonetheless, they are incapable of doing it. For instance, a heavy metal door is locked from the inside and someone asks their friend to come open it. The friend would definitely be facetious

⁷ For people interested in free will, this distinction is also why CC is not immediately a statement of the famous Principle of Alternative Possibilities criticized by FRANKFURT (1969).

if their answer was: “well, I opened the door earlier today, so it is certainly possible to open it. There is no reason for me to open it for you, you surely *can*(!) do it!”. Yes, we can (in the sense of possibility) open the door. However, we cannot (in the sense of capability) open it from the outside, say.

Capability is, therefore, different than possibility. So, why should one prefer capability to possibility when formulating OIC⁸? An answer might benefit from definitions of those terms, and I will offer them later. However, we can run with an intuitive grasp for now, taking capability to be “something agents can do”, and possibility to be “a way reality could be”. I think the answer could be made in four short arguments. Those arguments are mine: Vranas only stipulates that capability is the best way to go, and KING (2019, ch.1) only alludes to the difficulties of separating the relevant kinds of possibility. A ‘capability’ understanding of OIC will be very important to my arguments in Section 5.

First, OIC is usually defended with an appeal to the unfairness of requiring people to do what they truly cannot; consequently, a good version of OIC would have to be related to this requirement. Unfortunately, possibility, seen as above, would be so broad as to not sustain this intuition at all. For instance, assuming only nomological possibilities (possible in accordance with the laws of physics), every invention and piece of knowledge humanity now has is nomologically possible to be had. So, assuming that having some knowledge might be a moral duty— to know how much medicine to give to a child is a moral duty for a parent needing to treat their child, say— we can conclude that everyone of those who were doctors in the past and who did not know about the importance of sanitation are blameworthy, because they ought to know about sanitation and it was *possible* for them to know it (the laws of physics do not prevent it, given that the laws are the same for us now, and we do know that). Therefore, OIC-possibility fails fairness: capability is to be preferred.

Second, duties are attached to agents; however, possibilities are attached to worlds. It is unclear why a duty that commands a person to do something would be undermined by things utterly unrelated to the agent. That said, this is reasonable for metaphysical and logical possibilities, but for reasons quite apart from OIC. Duties cannot be metaphysically or logically impossible, but that is not because OIC is true; rather, it is because duties are connected to the space of possible actions: they separate the actions one ought to do from those they not-ought (sic)⁹. Obviously, metaphysically impossible things are not part of the set of possible actions, likewise for logically impossible actions. This is a trivial claim¹⁰; and the relevant possibility for OIC must be nomological or natural possibilities. Now, it is not clear why the duties one must have would be connected, in any immediate way, to the laws of nature¹¹. Contrarily, it seems rather plausible that the duty, being already connected to a given agent, would also be connected to other properties of that agent. This favors OIC-capability’ accounts.

⁸ Some people do not prefer capability to possibility, see BASSFORD (2022). For additional discussion on the formulation of an OIC thesis see: KING (2019, ch.1).

⁹ For reasons related to precision, I had to follow Logic’s (rather than English’s) grammar here.

¹⁰ Although see KING (2019, p.73), who argues that there could exist logically impossible ought. My argument is a simplified version of the one derived from WEDGWOOD’s (2016, 2018) papers on the semantics of ‘ought’.

¹¹ This connection is so problematic that GOLDWATER (2020) uses it in an argument against OIC.

Third, the epistemology of capabilities is much easier to apply than the epistemology of nomological possibility. Although, it is not necessarily easy, in fact, much of my argument below is based on there being practically impossible cases. My contention is relative: what is bad for ‘capability’ is worst for ‘possibility’. Modal epistemology is a field in and of itself, a particularly difficult one and most accounts it offers are, at best, unclear about how to deal with modality in general (they do tend to work better for logical possibilities, though). I do not wish to dwell too much on this¹², the point, however, stands: modal epistemology is a hugely controversial area; making any applications of OIC to be dependent upon it undermines its application considerably.

Finally, duties can be stronger and weaker, and one of the issues that surround how strong the duty is concerns how easy the person can dispatch the obligation. For instance, a person passing by a child who is in risk of being attacked by a small dog is strongly obligated to help the child. A person passing by a child in risk of being attacked by a bear is still obliged to help, but less so. OIC-capability helps explain this intuition. OIC-possibility does not, because possibility is a categorical thing. Either something is possible or it is not. This is not to say, however, that obligations are gradual all the way down. There may very well be a point from which a thing is, or is not obligatory, with no middle ground. But the strength of duties is a fairly intuitive notion that OIC-capability, but not OIC-possibility can make sense of.

I conclude that OIC should, as VRANAS (2007) suggests, be taken to be defined as capability. So now, let’s define capability a bit more precisely according to VRANAS (2007, p. 170): Capability is the temporally indexed conjunction of ability and opportunity. Where ability is the set of skills, knowledge and bodily dispositions available for enacting some action; and opportunity is the condition external to the agential control that allows them to apply their ability. For instance, I cannot fly an airplane because I lack the knowledge and the skills necessary for doing so. Had I learned them, I would be able to fly an airplane. Yet, if there is no airplane, or the airplane is broken, I would not be capable of flying it— I would have the ability, but not the opportunity. So OIC means ‘ought’ implies the ability and the opportunity conjoined at the same time.

It is rather plain that capability is an objective matter. I may be capable to do something and not know it, or believe I am capable and not be. The sense in which it is objective is rather tricky, for one may not do some things they think they are incapable of doing, even if they are capable, thus never doing them which prevent a simple dispositional view of capability. So, the best way to capture this objectivity is counterfactually. The person S is capable to ϕ if had they tried in the opportune contexts, they would succeed (or have a plausible likelihood to succeed). Being as it may, the ‘Can’ in OIC is an objective feature of the world. It is also a descriptive and empirical one. Then, what about the ‘Ought’ part of OIC?¹³

2.2. The Objectivity of Obligations

Now, we need to delve into GRAHAM’S (2010) paper to discuss why we should consider obligations to be objective. Unsurprisingly, the discussion is rather complex. Before I set up the formal definitions, we can take a step back and assess the debate. What is at stake is, in a sense,

¹²But see BERGLUND (2005) for an at-the-time exhaustive overview and a still worthwhile introduction.

¹³I am citing only VRANAS (2007) because my aim is not in defending OIC, per se. His contributions to this discussion, however, surpasses that paper and cover almost all the literature in one way or the other. They are all great papers and deserve the read. See: VRANAS (2018a,2018b, 2024). Vranas is contrasted by the book by KING (2019) who argues against OIC.

where we should put the “moral camera” so to speak¹⁴. Should we follow the agent’s perspective, or should we take a broad take on the whole scene and consider what would be the really best option? So, to keep the analogy, imagine the following scenario: we are watching one of those movies that the punch of the picture is to present a complicated moral situation. The same story was shot by the director in two different cameras (as far as I know, this is standard in the industry). In the first camera, we follow with the main character. We get some views as if it was first-person. In this situation, we see a fallen, broken chair with mud on it; we see the victim fallen down near the chair; and we see the knife with blood by the victims’ side. Later on, we meet a person who was a suspect and they had the mud in their feet. They also had motive, and there was a whole talk about the person liking knives for whatever reason. Ok, the spectator is set to blame this person for the murder.

The other camera, however, keeps the broad perspective over the whole house. In this camera, we can see another person hiding in the living-room, and when chance allows, escaping. At that time, we also get a close in their feet touching the source of mud: a puddle outside the house that anybody coming in or passing by would have to step on. Oh, wow, the villain was not the suspect from before! The question now is: who will the main character blame?

Well, the terrible screenplay above suggests forcefully that it will be the suspect, not the villain. The ethical question comes now: who *should* the main character blame? (I am a bad screenwriter, so, if the above case seems insufficient for blame, just imagine that whatever the evidence needed for blaming the suspect is filmed in the first camera; and that a full explanation of those pieces of evidence is shown in the second camera, as well as this additional, rather fatal evidence that the villain was in the crime scene.) The point is: in the first camera, the main character is fully justified in thinking that the suspect is to be blamed. In the second camera, we get the truth: that another person did it. Now, if we are objectivists, we would want to say that the main character should blame the actual villain. If we are subjectivists, then we would hold that the main character is correct in blaming the suspect, even if this would be the wrong (factually incorrect) choice.

With this intro, we can see the definitions Graham offers for subjectivism and objectivism. Starting with the former:

A moral theory, T, is ability-constrained-evidence-subjective=def. according to T, a person has the moral obligations that she has at a time solely in virtue of both facts about her abilities and facts about her evidential situation at (or prior to) that time (GRAHAM, 2010, p. 90).

And the latter is simply: “A moral theory, T, is objective=def. it is not the case that T is [ability-constrained-]¹⁵ evidence-subjective” (GRAHAM, 2010, p.89).

The definitions in that paper are done twice. The first to capture what subjectivists are saying. The second was to amend that definition to deal with a problem connected to taking OIC and subjectivism together. The problem is easy to see: the first definition of subjectivism did not include the ability restraint. So, if one believes in OIC, and that what they should do is based on the epistemic states of an agent, then, given that we can be wrong about our own capabilities, it follows that one could believe that they should ϕ even if they can’t. Which contradicts OIC. To mend this problem, the definition quoted above was made. Graham points out some reasons for

¹⁴ GREENE (2013, p. 133-1355) also uses the metaphor of the moral camera, although in a different context and with a different emphasis.

¹⁵ Author’s addition between brackets.

concern regarding that position. The subjectivist wants to hold that only the agent's epistemic stance can account for what they should do, *except* for their own ability. This ability is an objective feature in a subjectivist account. The only reason for the concession is the (reasonable) desire to keep OIC and subjectivism together. This is the definition of an *ad hoc* theoretical construction. Graham then asks: why not more objective conditions? And if more, why not *only*? (this last part is somewhat implied by his considerations).

I agree with this point. Subjectivism and OIC do not go hand-in-hand. If the agent, in a given point, was capable of getting more information, and, in those conditions, they were obliged to do that, OIC seems to warrant that the person *should get more information*. So, the fact that the person is not better informed is not a good reason to not *require* better information. If OIC is held, it is done so superfluously. If this is the case, then subjectivism would have to add a further argument for why someone's not *actually* knowing would entail that they *couldn't* know. The logical inference is obviously not valid, so they need additional reasons. If, on the other hand, they assume that not knowing does not prevent one from being obliged to know, then subjectivism seems untenable; given that if people could know better, and they should know better, then the fact that they do not know better is irrelevant for their moral situation, contrary to subjectivism.

Of course, holding OIC is not supposed to be the *only* excusing factor in the subjectivist theory. But the addition of other excusing factors must, at once, explain why not being better informed is a sufficient reason to forgo obligations to be better informed; furthermore, they would also maintain OIC saying that not being able to do something is sufficient for not being obliged to do it, independently of information people may have about their capability. Combining both claims we would have the following: what I know *now* limits what I should know; and what I cannot know limits what I should know. These combined propositions are not equivalent— but are dangerously similar— to a wholesale rejection that we have moral obligations to know things we do not currently know, or believe. I have not enough time to press the issue to its limits, but I hope that these remarks make it clear that any subjectivist accounts will have to account for this, and it is a *prima facie* objection to the combination of subjectivism and OIC. At best OIC is superfluous, at worst it offers a principle that claims a condition for exculpation that competes with subjectivism.

With the *prima facie* case delineated above, I move to Graham's two arguments for objectivism. First, we have the argument from correction, attributed to W.D. Ross's book "the Right and the Good". The argument is rather simple. Back to the two cameras above, according to the first-person camera, the main character should blame the suspect. In the second camera, they should blame the villain. The difference is that in the first, the main character has false, but justified belief. If this belief was corrected: going in accordance with the second camera, is it really true to say that their moral duties was changed? Isn't it more plausible that the person was *corrected*—implying that they were *wrong before*—rather than *redirected*—implying that the it was correct— but had to change for some reason? The argument suggests that when people go from imperfect knowledge to perfect knowledge, what change is not their moral obligations, but the correctness of their beliefs: they were supposed to act according to the corrected beliefs all-along¹⁶. This suggests that the objectively correct thing to do also matters and objectivism is correct.

¹⁶This argument has a question-begging ring to it. But so is true for most intuition-based arguments.

The second, and more persuasive, argument is, in Graham's own words:

The second argument for objectivism goes as follows: the question I want answered when I ask myself what my moral obligations are is the same as that which I want answered when, in seeking your help, I ask you what they are; but, to adequately answer me you don't need to consider my evidence concerning my situation; therefore, my moral obligations don't depend on my evidence concerning my situation; so, objectivism is true (GRAHAM, 2010, p. 91).

Graham is very direct and succinct, warranting the direct quote. However, if I could suggest one amend: the adviser does not need *only* to consider your evidence. The evidence does matter for an adviser, but further, independent truth-claims *also* matter. This is already sufficient to undermine subjectivism (who claims that *only* ability and evidence-dependent fact matter for moral obligation), which makes this a very compelling argument.

Graham's article moves from that to discussing plausible arguments for subjectivism. He develops interesting counters to them, nevertheless, the nuances of the discussion will not matter so much for us. If someone can hold a subjectivist position that deals with the challenge from OIC above and give a counter for the two arguments, they will be off-the-hook from my arguments below. In general, I think the subjectivist will just deny OIC and the arguments below will also not apply anymore¹⁷. In those arguments, objectivism will be assumed as part of the presuppositions about when someone is to be blamed or not. And on the existence of criteria about when is it fair to blame someone. If the fairness of acts of blaming is a moral imperative, as I contend it is; and objectivism is true, then the fairness of blaming is also an objective matter. Either a person, when judging other, does so fairly, or not fairly and there is a fact of the matter about which it is.

3. The Biology of Choice and Opportunities

The final stage-setting Section 1 need is connected to the notion of capability. As capabilities are agent's properties and those are instantiated by biological systems, it seems rather important to think about how we come to learn what someone is capable of doing. There are two answers to this question. The first is the dispositional answer. This one is rather simple to answer:

A subject S is justified in believing that some agent Ag is dispositionally capable to do some action A in a set of contexts, if the agent has a track-record of doing A in those contexts; or if a relevantly similar agent Ag has shown this track record and S has no reason to think that Ag* and Ag differ in regarding to the particular context being analyzed.*

This criteria could be improved upon, but it is serviceable. At any rate, this is not the epistemological sense that I think brings difficulties. Rather, the sense I mean is what I will call: "situated capability". This notion of capability is that applied in a particular, predefined context. For instance: you can open doors. But that does not mean that you are capable of opening the heavy, metal door of the example above. Dispositions are not undermined by few counterexamples. You might be able to drive, and drive well, even if you cannot drive after taking a powerful anesthetic. Unfortunately, it does not seem that we have a neat epistemological recipe to learn about situated capabilities. This will be discussed below, but it is important to check some general details of why that might be so difficult.

Vranas defined ability in terms of knowledge, skills and bodily capabilities. Opportunities are those things that allow the agent to put their abilities to good use. Vranas is explicit that psycho-

¹⁷ In fact, these arguments may be taken as evidence against objectivism!

logical properties are not included in the ability part. But I think that it is impossible for them to *not* be included in the opportunity condition for capability. A person's biological make-up might be able to run ten kilometers, they might have done that before, but had they had a bad infection recently, they might not be able to run, say, 3 kilometers. A few weeks later, with good recovery, they might be back to the long-distance running. It is clear that ability was not compromised. All the breathing, rhythm and muscular structure was still there, as well as the know-how of running. What changed is that the body took a hit and had to recover. This is opportunity, not ability, and it matters.

With this in mind, I wish to review the argument put forward by R. SAPOLSKY (2023, ch.3) that we have a nested structure of influences operating under every decision we make. I will not do this to draw his Hard Determinist conclusions. Rather, the claim I wish to make is rather plain: certain factors about our bodies are out of our control, they oscillate based on many, many, variables and those may affect our moral (and otherwise) decisions. The upshot is this: even if libertarian free will is true, and most of our decisions are our own as a non-caused cause; by being possible that we might be affected in the ways that the literature Sapolsky considers, then maybe a particular choice of ours every once in a while was not free. This is compatible with libertarian free will because the fact that external forces might affect our decisions in decisive ways is a plain fact of life. We might be libertarians drugged with special substances that make us temporarily incapable of choosing, for instance. So, I just wish to point out that we may, at times, be in a situation analogous to the drugged libertarian, and on those rare events, we might not be able to choose.

3.1. Sapolsky's argument in a nutshell

I will make a short presentation of the first part of Sapolsky's argument. His earlier book, SAPOLSKY (2017), makes a prolonged case for the same observations, with the important caveat offered by Sapolsky himself (SAPOLSKY, 2023, p. 47 fn.) that his first book suffered from citing some papers that failed the replication crisis in psychology. For brevity, and for the epistemological advantage of only relying on replicated results, I will follow only his (2023, ch.3).

The chapter in question aims at thinking about where does intent come from. The answer, unsurprisingly given that the author is a neurobiologist, has a lot to do with the brain. But not only— it has also much to do with culture and society as well. The review here will be only about the brain, for it suffices to make clear the way in which the opportunity condition may be hard to determine.

3.1.1. Second to minutes before

There is a behavior, why has it occurred? We are talking about humans, and humans are biological creatures. Our brains have something to do with decisions and personality, as the famous case of Phineas Gage illustrates. Brains have substantial control over two crucial systems: neurological and endocrinological systems, or oversimplifying: neurons and hormones.

This section of the book argues for the unconscious, unexpected effects of sensorial stimuli in

moral (or evaluative) decisions¹⁸. Two main of such cases deserve mentioning in this summary, both are connected to disgust and moral rigidity. This could be inferred weakly from the nature of moral condemnations being frequently associated to disgust, but it is more than words. In a series of experiments, subjects are put in a condition to display moral choices in both uncomfortable situations and disgusting ones. The uncomfortable situation was a bowl of ice-cold water. The disgusting one is an imitation of vomit. The two groups were given the same set of moral situations. Some of them are purity-related moral infractions and others are not (e.g. John rubbed someone's toothbrush in a public toilet's floor vs John risked someone's car with a key). Subjects in both situations had to decide the strength of punishment for the infractions.

Disgusted subjects gave harsher punishments to purity-related infractions than they did to the others. Moreover, their punishment was harsher than the punishment given by the uncomfortable group of subjects (SAPOLSKY, 2023, p. 48).

Still in the connection between disgust and moral rigidity, subjects are asked to make moral judgements about sex and sexuality. One group of subjects was asked to do that in a room with a mild (but noticeable) disgusting smell. The other had a non-disgusting smell in the room. Result: the group in the smelly room leaned more towards conservatism than the other group (SAPOLSKY, 2023, p.47). Who would imagine that a courthouse being built near a smelly area might affect the average decision of judges there?¹⁹

3.1.2. Minutes to days before

If the previous section presents some ways our brains fail inasmuch as neuronal activations are concerned, in this section we will discuss the problems caused by hormones. The first case to consider is testosterone. Contrary to popular belief, testosterone does not lead to increase in aggression per se. It makes people more prone to defend their social status, especially displaying aggressivity towards those that are *under* them in the hierarchy. Additionally, testosterone affects people who are prone to violence to have a lower threshold to begin violent behavior. Non-violent people seem to be less effected. Finally, it makes us more prone to interpret socially ambiguous interactions (e.g. people's faces) as aggressive interactions.

What causes rises in testosterone in people and in its effects? Time of the day, recent activities (fights, discussions, sex, etc). Individual variations. These individual variations are due to genes; fetal, pre and post-natal environments. That means that we have very little control over the levels of testosterone we have when we are about to interpret a social situation in a morally significant way (SAPOLSKY, 2023, p. 52-53).

Similarly, the hormone oxytocin is connected to human bonding. It has multiple social functions and explains much of our social lives. Unfortunately, beyond some beneficial outcomes, it has also surprising effect in in-group favoritism: the psychological disposition to prefer those perceived as in-group than those perceived as out-group. (SAPOLSKY, 2023, p. 54-55). Once more, the effect of this hormone is determined by a myriad set of factors and we can hardly hope to know how it effected a subject in a particular situation in real life.

¹⁸I will not dwell in the suggested neurological explanations offered by Sapolsky. But each of the examples receives an explanation in the same page in which they are described.

¹⁹This is me extrapolating from the data. Similar results were tested on legal decisions, but there is more to be done to make a categorical claim that the result is important.

3.1.3. Weeks to years before

This section is about neuroplasticity and bacteria in one's digestive system. The neuroplasticity is attached to how the brain changes over time, building and breaking synapses and/or neurons. It can affect morality in substantial ways. Mental illness and stress can have overarching effects in moral cognition, increasing fear, aggressivity and the like. Likewise, being chronically stressed. Changes for the best can happen also, if the environment is positive (SAPOLSKY, 2023, p. 58).

About the bacteria, we have less to think about in terms of morality directly. But that bacteria in one's digestive system can affect behavioral patterns is a rather surprising fact. Sapolsky lists the following effects observable by switching bacterial makeup between individuals: appetite and food cravings; gene expression patterns in the neurons; disposition towards anxiety and the speed of development of some mental illnesses (SAPOLSKY, 2023, p. 59). If bacteria in one's stomach can alter their behavior, maybe they could do so in a morally relevant context too?

To be very clear, from now on I will be assuming libertarian free will. So, even if bacteria could have some effect, it would usually be out-competed by the agential powers people have. But as a powerful drug might temporarily undermine this agential power, so too a combination of multiple effects like the abovementioned might at times overpower the agential control. The proposition I will be assuming is:

Force-balance (FoB): a strong combination of factors might, on occasion, undermine the agential powers a given libertarian-free-willed agent might have.

If this is the case, the agent loses the opportunity to apply their abilities to the satisfaction of their objectives, inclusive moral ones.

4. Normative Demands, Obligations and Fair Blaming²⁰

The next thing I need to defend in my argument is that there is a class of actions that are underexplored, namely, the actions associated to normatively demanding of others. These "normative demands" are second-order actions (actions about actions) concerning how people engage in the practices of morality. This is not the same as, to choose a single term: "normative evaluations". I will store "normative evaluations" to the traditional ethical stance connected to what people (really) ought to do, and normative demands to the acts people engage in when they hold others (and themselves) to some standard, being it a good one or not. Henceforth, I will just say 'demands' and 'evaluations' to mean their normative kinds.

Demands can cause bad or good moral effects on people. For instance, a professor who hire a new PhD student to be their assistant but only demands of them relatively easy cognitive tasks like ("summarize a chapter of an undergrad introduction book") is harming the student by displaying a neglect over their true capacities. In the other direction, if the same professor demands too many high-level tasks ("one knock-down argument against the big players in the literature per week!") this will also cause harm, in this case of overdemanding. In the opposite direction, giving people responsibilities that track their capability is a way of recognizing them as full participants in a given community and to help them to develop. Because of this potential for harm

²⁰ Professor Macnamara argues against demands in connection to blame. Although I cannot do full justice to her arguments here, one thing I can say is that her view on demands is very different from mine, and most of her objections do not even begin to apply to my work (MACNAMARA, 2013).

and good, I contend that demanding is a kind of action that is passive of moral evaluation. Good demands being called “fair” and bad demands being “unfair”. The set of neutral plus fair demands would be called “adequate” demands, whereas “inadequate” is the same as unfair. Presumably, fair demands track obligations, for it would be truly weird if there was no merit in following the best process in ascertaining the adequacy of our demands. Nevertheless, the relation might not be bi-conditional, in fact, there is reason to think that it is not.

I contend that, if what I have argued in Section 2 is correct, then evaluations will correspond to the objectivist perspective taken into account earlier. Whereas, it seems that the fairness of demands are the most plausible candidates for what the subjectivist’s intuitions are following. In this interpretation, the Objectivist/Subjectivist debate is actually based on the systematic ambiguity of thinking about what is actually correct versus thinking about what we should fairly demand of people in a given context.

In a nutshell, the issue is the following: We can evaluate agent’s demands and conclude that they are bad (or good) in the number of ways that are representative of normal evaluation of actions. But, the differences between evaluations and demands is that of perspective. What is correct for one to demand of others depends on many factors about the demanding agent as well as the demanded subject. This is not supposed to be the case in an objectivist view of evaluation. For this objectivist angle, there are things people are obligated to do, that they are allowed to do, and that they are forbidden to do. These things do not depend on who is evaluating. These judgements are supposed to be (as close as possible to) true-to-the-eyes-of-a-god. On the other hand, subjectivism seems to track the norms about demands quite well, even if it does not do so with obligations. The person in the movie (see Section 2) might be obliged to accuse the murderer for the murder in the living room instead of the suspect, and yet, it might not be fair to demand that they do so because they did what we could reasonably demand of them to do. If I am correct, this shift in the moral camera that captures the Objectivist/Subjectivist debate is a shift from the morality of obligations to the morality of demands.

The existence of such a difference of perspective suggests that the sense in which we are allowed to demand things from other people is tied to their objective obligations only partially: there are also the standards to which we, as judges, should obey.

This is discussed in COATES & TOGNAZZINI (2013, p. 18-23), where the authors distinguish three ways in which blaming might be adequate. The first are the conditions of blameworthiness; the second are the conditions of *jurisdiction*; and the third are the conditions of procedure. Without prolonging too much on them, the first kind is about what criteria should we look for when deciding if the person is to be blamed. The second— and most relevant here— contends that even if a subject is to be blamed, not all people can blame that person in the same way. Some might be in such a position as to be hypocritical when blaming, and thus losing the fairness of their action. Others might be so far away of the case (in a broad sense of ‘far away’, not only the geographical one) that their blaming is unfair for not being adequately engaged with some relevant aspects. The third condition is somewhat close to the second, but emphasizes the way in which we pursue blaming. If we administer it with proportionality, or following a good investigation.

I argue here that it is part of the second condition that we should not demand of others more than we could expect of ourselves in the same situation, or more than we can expect of a rele-

vantly similar person in the same situation. How this requirements are to be spelled out is work for the future, yet it is reasonable to say that, within reason, we know with whom to fairly compare cases, and what we should expect of ourselves²¹.

Although this works in general for practical purposes, if we wish these intuitive rules of thumb to actually track really existing obligations, we need more epistemological work to be done. Intuitions can fail us in many ways. The trouble in the next section is entirely directed to this asymmetry: we intuitively know when and how to blame people but given that obligations are objective facts; and, if OIC is true, CC is proposition we must be able to affirm or deny in particular cases, then the epistemological requirement for a fair blame will have to be stronger than the intuitions. This, I will argue, is what we cannot deliver.

That said, it would be odd to say that being obligated is not-at-all related to what we are allowed to demand of the subject, thus it is crucial to think about the relations between what we are normatively allowed to demand and what people are obliged to do. Some seemingly plausible rules about it will be relevant for the arguments in Section 5. Let ‘agents’ to be those people who are demanding of others and ‘subjects’ to be those who are being demanded:

Obligations and Abstract Demands (OAD): If a given subject has failed to fulfill an obligation, there is, at least in principle, an agent (which might include themselves) who is in position to fairly demand that they had not done so.

Lack of Obligation is the Lack of Demands (LOLD): If a subject is not obliged to Φ , then there is no agent for whom it would be fair to (normatively) demand that the subject should Φ , and/or blaming the subject for not having Φ -ed.

From Obligations to Demands 1 (FOD-1): if a subject is obliged to ϕ , then it is fair for some agent to demand of them that they should ϕ , so long as the agent is justified in believing that the subject is capable of ϕ -ing; and the agent has justified belief that the subject is obliged to ϕ .

As those *prima facie* norms are normal moral norms, they should be objectively valid, if objectivism is true. The main reason to hold fair demands, I think, is exactly because, as the principles claim, there is a close relation between a fair demand and objective obligations. This is as it should be, given that, otherwise, morality itself would be inconsistent (and we hope this is not true, although I have very few argument to claim it isn’t).

4.1. Fair blame

The foreplay is almost done, so let’s set up the last few assumptions needed for the arguments in Section 5. First, I assume libertarianism. And I assume that there is an important sense in which people are responsible for their actions, in the sense of being proprietary, answerable causal beginners that acted for motives connected to the action, and who could have acted otherwise. Therefore, I am presupposing that there are legitimate ways to blame and punish people for what they have done. This being said, I am assuming that libertarianism is true, but not that we

²¹ I am certain that the reader is skeptical so here is some remarks. We might always be allowed to expect other subjects to (i) act in conformity to instrumental rationality. (ii) act in accordance with their role’s duties (e.g. to treat people if one is a doctor); and (iii) to perform in accordance with objectively valid conditions (e.g. in accordance with the true moral system, or in ways that are avoiding clear failures).

know that it is true, anymore than that we know it right now. That is, although people will be uncaused causes of their own behavior most of times, we, as outsiders evaluating people and making demands on them will have no more access to this fact than we currently do. There is no “free-will-do-metre”.

To assume, however, that blame is sometimes due is not to claim that it cannot be misdirected; being applied to people who do not deserve (so much of) it. Hence, granting that blame is morally acceptable or even necessary is not the whole story, we must discuss under which conditions that blame can be fairly attributed. If it cannot be fairly attributed then it must either be only a descriptive tool to capture a given inclination of resented people, or it is an incorrect, unfair response and thus immoral. To be clear, I have the following definition of fair blame:

Fair Blaming (FB): A judgment of blame is done fairly, if the agent doing it is: i) *well-informed*, meaning that the judgement respects a conscientious amount of morally relevant information about the judged situation; ii) *well-reasoned*, meaning that the arguments are sound, bear relevant premises and are not defeated by the full-body of morally relevant information taken into consideration, or that ought to be taken into consideration; iii) *adequately universal*, meaning that the individual, the arguments and principles and all relevant moral information lead to a judgment of blame that would apply equally to any person in the same situation. Finally, iv) *proportionality*: the judgment of blame must accept some degree of proportionality between the badness of the action, the amount of responsibility the subject had and the amount of blame we allocate to that subject.

First, *well-informedness* requires conscientious appeal to the available morally relevant properties. This means both taking in consideration all the evidence already available and applying adequate standards for further investigation. This is obviously vague and highly contextual. But, although the details matter greatly in general, for this argument the reader is invited to fulfill it with their favorite theory of epistemic inquiry²². Regardless, it is hard to imagine a world in which a poor investigation would yield a fair blame. Additionally, if no good account of what an adequate investigation into blameworthiness is possible, my main argument will run even better, so for the sake of strengthening the position of my adversaries, I shall concede that some such investigation is possible. *Well-reasonedness* does not need much defense. Obviously, if an agent tries to blame some subject based on poor reasoning, the judgement will not satisfy FOD-1 and, hence, will not be a demand that tracks a duty. *Adequate universality*, is the condition of non-ad-hoc-ness and this is equally obviously necessary. If I just make up standards for each person by fiat this arbitrariness is bound to bring forth injustice. And finally, *proportionality* is not very important for my argument, but something like it must be added if we wish to have an adequately complete account of fair blaming. In the end, I think these defining conditions do offer us a good candidate of a definition for Fair Blaming. If not, they are at least necessary conditions, and this much is what I hope to argue from.

So, if i)-iv) are at least necessary conditions, what is the problem I wish to advance? The problem is that i)-iii) are not plausibly met when it comes to token people. I will take some time to present this concept.

²² Mine would be that of GOLDBERG's (2018) and his Hybrid account. Although the theory faces some push-back from SIMION (2024) and PEELS (2019).

4.2. Token and Type people

Abstraction is a crucial part of hypothetical reasoning, but many judgements of blame are actually conducted against a living, flesh-and-bones person. The difference between hypothetical people like the best possible Brazilian, or Jack who must decide if he pulls the lever in a Trolley Problem-type, or the average German; and concrete people, is that in the first case, their doings— and the morally relevant properties of their situations— are all established by stipulation or abstraction, not so when it comes to the concrete ones out there. Let's call the abstract people: *type people*; and the concrete kind: *token people*.

The great difference between those two kinds of people is the situations in which they enact moral actions, type-people do so in thought experiments whereas token people do so in the real world. More precisely, although we know all about hypothetical scenarios, and we can wish out of existence some unfriendly complexities; in real life, all actual events are imparting in the moral action, as the results discussed in Section 3 above suggest. This is not to say that there are no irrelevant properties for moral consideration, sure there are. Usually, it is an irrelevant fact about a murder attempt whether or not it happened near a quartz rock. Therefore, I grant that there are cases of properties that are morally irrelevant and that those can be ignored when thinking about type or token people's moral character or blameworthiness.

The problem is quite different, in fact. It concerns the range of morally relevant properties being investigated and found out by both moral psychology and philosophy. This number is growing every year with new findings about how humans actually think morally and what they take into consideration when judging to act. This is not to say that the morality of a given decision is itself to be found in the lab behavior of monkeys or in brain activity. Rather, the issue is about what people are actually capable of doing: the degree to which the moral capability of people is affected by those ever-growing numbers of known morally relevant properties. The reason this matters is because OIC makes capability into a robust moral property: one whose presence or absence changes the fairness of a judgement by itself.

4.3. Token OIC

OIC applies to moral duties, and those apply to people. If there are two kinds of people, then it seems fair to consider whether we have two kinds of OIC. It is probable that there are because the fairness of an agent's demands over a subject that they should ϕ , depends on whether the subject had ability and opportunities and whether the agent knew that²³; and this issue can depend on whether the context is hypothetical or empirical. If it is hypothetical, then all the capacities and opportunities can be determined by fiat. If it was empirical, then we have to make a case that the particular person had, on that particular case, the ability and the opportunity. Given that 'ability' are the skills, knowledge and agential bodily functions, it is usually easier to determine if the person was able. However, the opportunity depends on the person having the right conditions at her disposal to enact their ability. But, given that the variability of those properties that impact our behavior can be very wide, and in some extremes they are like a powerful drug in their incapacitating powers, it is possible that the subject — even having the ability— might be biologically constrained, thus lacking the opportunity to impart that ability

²³ Given LOLD and OIC, if there is a lack of capability, then there can be no fair demands to the effect that Φ .

into getting a duty done. To repeat: this is compatible with libertarianism, for libertarians only would contend that *most of the time* we have free will, not that *all of the time* we have such. The comparison with the drug is important. If we have free will *all of the time* then a drugged person is fully responsible. But they are not. Hence, we do not. Equally, I hold, for the specific conditions in which FoB (Section 3) applies, which I am assuming are quite rare.

Given that the issue I wish to bring out is on the fairness of demands, we will focus on the agent (the person blaming), the subject's capability being plain or not (having the opportunity or not) is not very relevant. It would be if we had a 'free-will-do-metre', but we do not. Absent that, the problem remains: How could an agent demanding that a subject should do ϕ be justified in believing that the subject had (at that particular frame of time) the opportunity to ϕ ? Without that justification, the demand will not track obligations failing FOD-1; and will not satisfy the condition i) of FB. Hence, without that justification, the demand for ϕ -ing would not be fair, and is unlikely to be connected to a true obligation the subject has. This is the crux of the argument, so the reader is advised to go back and reacquaint with the principles and definitions mentioned.

Then, the crux of the matter is that if we, as agents, cannot determine if the subject is in one of those rare moments in which multiple biological variables came together to undermine the subject's opportunity to impart their abilities into fulfilling their duties, we would fail FB and FOD-1 above. The reason being that, if OIC is true, capability is a robust morally relevant property: one whose presence or absence makes a difference by itself into blame or not-blame.

We have two ways of ascertaining moral blame, one is strictly speaking abstract and hypothetical, the other is highly empirical. Both depend on the same principles and definitions of fairness, and thus both should be equally valuable in guiding our moral judgements. But, naturally, the empirical version—the one directed at blaming token people—is much harder to do, and a lot more consequential.

5. The Epistemological Arguments Against Token Blame

5.1. The argument from full investigation

In this section, I advance the most direct argument against the possibility of guaranteeing the conditions of FB for token people. I will call "the capability of a token person to fulfill their duties" their 'token capability'; and "subject's capability" when it goes for both type and token people. In the next section, I will advance a more nuanced version, yet, both start with the same premises, which will be simply numbered; and when the premises are unique to one of the two arguments it will be named with numbers and lowercase letters. This first argument has two versions. It moves like this:

- 1) We can only blame subjects fairly if we conscientiously investigate the morally relevant properties regarding their situation under consideration (intuitively plausible).
 - 2) The subjects' capability is a morally relevant property regarding the situation, of such a kind and relevance that all conscientious investigations must justifiably ascertain their presence or absence (from OIC and capability being robust).
- C1) We can only blame token people fairly if we justifiably take into consideration their token

capabilities (follows from 1,2).

3a) The number of known conditions for capability has increased as research in moral psychology, moral cognition, neuroscience, social psychology, philosophy, etc. has developed (SAPOLSKY (2017, 2023), briefly summarized in Section 3 above²⁴).

4) The list of moral variables is, now, so vast and difficult to account for; and the research of their impact is so contextual, that we cannot justifiably determine whether their co-occurrence in a given context would be sufficient to undermine subject's capability. (Section 3, SAPOLSKY, 2017; SAPOLSKY, 2023); Section 1: CC's requirement fulfillment is underdetermined)

5a) Additionally, we cannot determine if the situation's description is ever complete and include all the opportunity-relevant properties. (given that we are learning more and more, and that we lack a general law for predictions, we cannot affirm that our current knowledge is exhaustive).

6a) If we cannot know that a given description of opportunity-relevant properties is complete, nor when the overall co-occurrence of those properties will satisfy FoB, then we cannot be sure when judging the presence of token capabilities.

C2) If we are never sure whether token capabilities are present, we always fail necessary conditions for token-Fair Blaming.

C2 leads to the rejection of Fair Blaming because we probably fail FB's Criterium i), for there is no way to decide if all the relevant information was considered, hence if we do actually consider them it was out of luck. And we fail criterium iii) because by not knowing if we have all the relevant properties we cannot be sure which cases are morally analogous to which and, hence, reasoned universality is fatally compromised.

I find this argument particularly convincing against condition i). To sum up the argument, I claimed as follows. Fair blame requires conscientious analysis of the morally relevant properties. A crucial one is the capability people have to act according to duty. This property is crucial because of OIC; but it is also an empirical matter whether a given token person has the opportunity to actually fulfill their duties. It is not only an empirical matter as it is also an impossible practical task, given that it requires understanding factors that go back to evolutionary developments and cultural implications, among others. Therefore, we are never able to perform an adequate analysis of the morally relevant properties in token people, 'opportunity' will allude us. This is not to say that the subject had, or didn't have, the opportunity. It is only to say that the agent is not justified in their belief one way or the other; and because fairness requires this justification according to FOD-1 and FB, the demand is not fair—the act of blaming is not fair. But because this situation is pervasive, no acts of blame could be fair.

5.1.1. A second version of the same argument

Much the same argument can be run in a slightly different way. This time the argument is based on a different attempt to select morally relevant properties. It departs from a different approach to fair blaming, namely: take universality and proportionality as the sources of your judgement.

Method: We start with a past case whose blame is really plausible and infer from that case that the current one is also to be blamed. This strategy has a juridical ring to it, but it promises to

²⁴ See TIBERIUS (2014) for a book-length introductory review.

avoid the problem above if only by assuming that, at times, we have blamed correctly. I do not want to draw a complete skeptical scenario from the outset, so I will grant the hypothesis.

Hence, the strategy we are contemplating is the following: Stage 1- get a comparison case (C-Case) from which to judge your actual case (A-Case). Stage 2- highlight similarities and differences between C-Case and A-Case. Stage 3- punish the A-Case in accordance with the C-Case plus or minus the aggravating or exculpating factors given by Stage 2-. Plausibly, if these three stages are done appropriately, then we can grant a satisfaction to FB necessary conditions, even if they are slightly weaker in this case. So, under this background, the new version of the argument goes like this:

- 1) We can only blame subjects fairly if we conscientiously investigate the morally relevant properties regarding their situation under consideration.
- 2) The subjects' capability is a morally relevant property regarding the situation of such a kind and relevance that all conscientious investigations must justifiably ascertain its presence or absence.
- C1) We can only blame token people fairly if we justifiably take into consideration their token capabilities.
- 3b) We have a C-Case from which to compare and get the information about our A-Case's token capabilities. (by hypothesis)
- 4) The list of moral variables is, now, so vast and difficult to account for; and the research of their impact is so contextual, that we cannot determine whether their co-occurrence in a given context would not be sufficient to undermine subject's capability.
- 5b) The C-Case is only useful in blaming token people if we can conscientiously assess the differences between the C-Case's token capabilities and the A-Case's token capabilities. (plausible assumption about analogical reasoning)
- 6b) if 4 is true, we cannot empirically guarantee that if C-Case was fairly blamed, then A-Case was fairly blamed, for the effect of C-Case's morally relevant properties might have been context-dependent, or some exclusively A-Case set of properties might defeat the inference that held for C-Case.

Hence, 6b concludes the same as C2 before. This argument, actually, requires that we can separate well the properties of the cases between those that matter and those that do not. But, because in token people we cannot, the analogy will not offer justification, rather, it will presuppose that all morally relevant properties (for a conscientious investigation) are the same between both cases. But the very issue at hand is to suggest that we cannot guarantee that! This new strategy, however, opens up to a second line of criticism that attacks the assumption of analogousness.

5.2. The Argument from the Lack of Analogies

The second kind of argument takes pains to show that, if we are considering token people, we cannot just assume that two cases will be analogous. I will defend later that there are properties with what I will call "idiosyncratic moral effects." These are defined as properties whose moral significance varies substantially – they can in one case be able to allow for fair blaming and in

another to allow for full exculpation – depending on with which other properties they are co-instantiated; and it requires that consequence to be different both in quality and in intensity.

1) We can only blame subjects fairly if we conscientiously investigate the morally relevant properties regarding their situation under consideration.

2) The subject's capability is a morally relevant property regarding the situation of such a kind and relevance that all conscientious investigations must justifiably ascertain its presence or absence.

C1) We can only blame token people fairly if we justifiably take into consideration their token capabilities.

3c) To argue from analogy between different token cases (or from type people to token people) we must guarantee that the analogy holds and agent's capabilities are the same in both cases (at least inasmuch as the presence or absence is the same).

4) The list of moral variables is, now, so vast and difficult to account for; and the research of their impact is so contextual, that we cannot determine whether their co-occurrence in a given context would not be sufficient to undermine subject's capability.

Two options for analogies to hold.

Option 1:

— 5c)²⁵ An analogy holds for morality if the differences between cases, even if they are morally relevant, are not morally relevant for the current situation (foundational principle of analogical reasoning).

—— 6c) If properties can have idiosyncratic moral effects, there is no class of reference from which to justifiably infer whether a given moral property was morally relevant to another situation or not (follows from 5c and the definition of idiosyncratic moral effects)

——— Consequently, the analogy is not justified.

Option 2:

— 7c) An analogy holds for morality if the differences between cases, even if they are morally relevant, are not as impactful in the current situation as the similarities.

—— 8c) If properties can have idiosyncratic moral effects, we cannot rule out the hypothesis that the differing properties might have more impact in one case than in the other such as to trump the impact of the similarities (due to 4 and the definition of idiosyncratic moral effects).

——— Consequently, the analogy is not justified.

9c) if either 6c) or 8c) are true; and if properties can have idiosyncratic moral effects, then we have no justified source to affirm the analogy between the fairness of demands due to the supposed parity between two seemingly equivalent cases. (We are gratuitously assuming that the differences are not (sufficiently) relevant i.e. begging the question).

10c) Properties can have idiosyncratic moral effects.

C3) We cannot obtain fair judgements of blame by analogy with either other token people or with type people.

²⁵ '—' is being used, as is standard in Natural Deduction, to mark a hypothesis and the hierarchy of dependence. As the hypothetical are not included in a single body of argument, but bifurcate in two different ones.

C4 Hence, if both C2) and C3), then we cannot satisfy Fair Blame at all.

This argument rejected the possibility of plausible universality, undermining FB's condition iii). But, by doing that, it rejected an important strategy to highlight morally relevant properties— analogy— which makes the argument also undermine i) by making it explicit that analogies are incapable of offering us all the needed information about morally relevant properties. The above argument requires two supporting pieces of evidence. One to argue for premise 10c, and another to argue for the importance of properties having idiosyncratic moral effects.

To defend 10c), although this is not a discussion to be settled with a single argument, I find that the hypothesis is highly plausible. It just states that in different cases, the same property has different moral effects to the extent that the very same property might make one person to be blamed in a situation and to be exculpated in another due only to the contribution of these properties to the other morally relevant ones. Moreover, it requires that consequence to be different both in quality and in intensity. Given the variability of human moral behavior, the idea that a set of three properties like “short-temper, obsessive motivation and high conceptual intelligence” could not lead to different outcomes in different people is rather unlikely. Additionally, a plausible example might suggest that they are not even that hard to come by.

Imagine a subject who don't have their right arm. Not having a right arm may be irrelevant when considering the fairness of a demand over that subject for donating money for charity. Nevertheless, it might be very relevant when deciding whether the subject was a murderer in a case in which forensics guarantee that the gun was held in the suspect's right hand. If this example is plausible, the property “having a right arm” has idiosyncratic moral effects in the sense I am using the terms. It absolutely does not matter in one case, and absolutely exculpates in the other.

Alright, if properties with idiosyncratic moral effects can exist, why do they matter for premises 6c,8c, and 9c? The reason is quite simple. If a given property might, when added to others, bring about *surprising*, novel moral effects, then by choosing the two supposedly analogous cases based on their similarity in moral properties alone will not suffice, given that this idiosyncrasy emerges from the overall combination of properties—moral and otherwise. On the other hand, stipulating that the cases are similar enough will just beg the question. This, once more, only applies to token people. When it comes to type people, we can stipulate that the cases are analogous given that we can determine the absence of idiosyncrasies in the cases.

The last thing that I must argue is that some of those novel moral effects may affect capability. But this is rather obvious. Imagine two men: Johnny and Najib, who are both, at around the same time, in a very similar situation. Both men had received orders from their spouses to hold on to a plate of sweets and candies while they prepare everything for their children's party. There are not many candies and if even one is taken the beautiful decoration their spouses worked hard to get will be undone. Both are sweet-lovers and those sweets are their favorites. Once more, both men had had super stressful weeks, and they are prone to sugar eating when wishing to compensate for stress. However, Johnny's mother had a substantial trauma when he was still a fetus and because of this, he acquired a decreased self-control relatively to Najib, Johnny usually can deal with that quite well. However, it so happens that when three factors get together, Johnny's self-control runs out. The conditions are: stressful weeks, favorite sweets, and low blood sugar caused by

more-than-average time without eating²⁶. Consequently, Johnny, but not Najib, eats one sweet and get their spouse's angry response for it. Now, the property "having had a stressful week" prevented Johnny, but not Najib from fulfilling a duty. Had they been instantiated in slightly different conditions, both men would had been analogous relatively to this particular scenario. But in the next scenario, we would not be able to know. This is why the existence of properties with idiosyncratic moral effects is a problem for these kinds of token people analogies. The only issue remaining is whether this sort of example is compatible with the dialectical assumption of libertarian free will. For the particular case that I am making here, it doesn't matter, for the objective is to show that idiosyncratic moral effects can affect capability by undermining opportunity. But if there is an inconsistency with libertarian free will, then there might exist more implications for the argument above. I cannot pursue the issue here, for it would require a precise presentation of free will accounts.

To sum up, this argument, if successful, undermines our capacity to learn about morally relevant properties in situations of blaming by asking us to justify our analogies regarding whether or not a given token agent exhibits an ability.

6. Recapitulating and exploring exits

So, here is the eagle-eye point of view. We started discussing blame as an objective thing grounded on the way people are and in what they do. I further added that obligations are also objective things grounded on the way the world is. I discussed some reasons to think that, even if libertarian free will is true, there might be some combination of co-occurrent properties that, on occasion, could block the opportunity to act according to one's own free will. I argued that this is analogous to the case of a strong drug making people do things they would not ordinarily do. I then presented a distinction between token and type people. I moved to argue that there is one morally relevant property i.e. a subject's capability that must be known to be present or not, and, still about which we get to learn very little when it comes to token people. Because conscientious investigation was one condition for fair blaming, and because those arguments show that this is impossible to do, we concluded that fair blaming was impossible.

A conclusion nobody wants, I guess. At any rate, I don't. So, this section explores how to avoid it, that is, I will take the arguments in the previous section to be *reductio ad absurdum* of the set of propositions we brought in the beginning, with the exclusion of libertarian free will. That view plays only the didactical role of showing that the discussion here is very loosely connected to free will. The set of assumptions that lead to the problems are these: Realism about blame, objectivism about obligation, moral normativity over demands, epistemic fallibility in the sense of not being able to know all the propositions we could possibly know in the relevant time, OIC (capability), and the empirical discoverability of new morally relevant properties. Idiosyncratic moral effects are not necessary for the first argument, but are for the second. Finally, of course, there was the two conditions for FB (i) and iii), and FOD-1. We might try to deny each of them, and although this could be done in a myriad ways, I will discuss just the most obvious and broad consequences.

²⁶ For these kinds of effects in real life, Sapolsky's 2023 book offer a very useful compilation and explanation. The fetal effect is derived from the really observed consequences of the Dutch Hunger: See ROSEBOOM *et al* (2011) for a review of these effects. For space's sake I preferred to just add the small, grounded fantasy above than doing a dive into those papers here.

6.1. First option, undermining OIC

If OIC is false, then the individual's ability to fulfill their duties becomes less important, and the rest of the morally relevant properties may be more tractable epistemically speaking. The OIC-sceptic must, however, still contend with the intuitive idea that people's capability matters to their duties. Hence, this strategy requires a substitutive theory, not merely a rejection of OIC. For different reasons, KING (2019, ch.4) explores such a view. I think her account is unsuccessful to face this challenge, and it fails on some internal grounds, but it would be a way out²⁷.

6.2. Second option: Embracing

It is interesting— and somewhat ironic— to think about the implications of OIC to avoid the conundrum above. Given that the knowledge of capability is impossible, and, hence, we have a practical impossibility, if OIC is true, then the duty to blame fairly may be wavered. In which case, we are in the clear to blame unfairly, as we both cannot avoid to blame (for an argument to this effect SHER (2006 ch.1) and cannot hope to blame fairly. That is a consistent way out of my argument, if only a bit unorthodox. In other words, the application of OIC to the act of judgement makes it the case that we might not be obliged to blame fairly, because we are incapable of doing so. Hence, blaming may just not be a thing we can ever do morally. This denies the normativity of demands above. Some people like PEREBOOM (2001) and SAPOLSKY (2023) might wish to just agree with this move, after all, they already do not believe that there is such a thing as fair blame.

6.3. Third option: change the definition

Maybe there is no requirement of conscientiousness in the moral inquiry required for fair blaming. This may well be, but most of my arguments above only needed an argued position about the effect of one property of the agent— namely: 'opportunity'— consequently, to avoid the issue definitionally might be harder than one might think. This option combined with rejecting OIC may seem more promising. At any rate, it seems hard to reject that demanding— and blaming as a subset of that— require some adequate evidential standard, and this standard clearly seems to need to account for all the properties that might, on their own, undermine blame: the 'robust properties' as I called them.

This is the case of opportunity. If there was opportunity, the person might be fairly blamed. If there wasn't, they cannot. The fact that we do not know either way may be taken to be negligible because most times people can act according to their duties. But I don't think this is so easy. We know that there are excusing circumstances at multiple times— sometimes even fully exculpating ones— so the odds that FoB happens is not *so low* as to be negligible by fiat. When we do the empirical investigation, we frequently see those factors playing in, FoB conditions are not *ad-hoc*. That is different from considering a claim for the lack of opportunity given the intervention from 4D alien Spacecraft, or any some such scenario. In this new scenario, but not in FoB, we have *no evidence at all* of the existence of these aliens. This point needs to receive more attention if this argument is to truly fly, but it is *prima facie* plausible to maintain that all robust moral properties should be adequately investigated.

²⁷ A forceful attempt against OIC is offered by J. GOLDWATER (2020).

6.4. Fourth option: Acceptance

I have been dealing with the hypothesis that moral judgement justification must be somewhat strong, that is because most people take that morality is something we are very good at; maybe we can even do it *a priori* based on intuitions like Moore seem to have defended. Therefore, the epistemic conditions have high standards. But maybe we are bad at moral judgements, in this case, we might just be high fallibilists about morality and just go with it. In which case, the definition of Fair Blaming might be: the argument to the effect of blame was plausible. This has many meta-ethical implications. WEDGWOOD (2007) does argue for a fallibilist view on our moral epistemology, but given the above set-up, the fallibilism will be very strong. We will be mostly adding a huge luck element in our investigations to the effect that, if I investigated all I could in the best way available to me, it is likely that the fairness of blame will track obligations adequately.

6.5. Fifth option: blame is directly related to demands, not to obligations

Maybe the problem with this whole project is the metaphysical realist background. I tend to think that this is correct. Perhaps the idea that people cause bad things and then get to be blamed by them, given that not doing bad things was an obligation they really had is what does not mix well with the empirical nature of our actions. The suggestions here would be the following. Perhaps we should decouple blame from failing an obligation and link it more directly to the fairness of demanding that the person had not done that. Then, blame would require the justified belief that the subject of blame has causally participated in the coming-to-be of the bad outcome (in the relevant way that exclude spurious causal claims) and the fairness of the demand that the person had not done so. Now, because the fairness of the demand only require justified beliefs, and the absence of opportunity might not lead to unjustified beliefs about obligations, only untruthful ones, it would follow that the argument above is dismantled. The mere possibility that the subject might not be capable of delivering on their duty would not undermine the fairness of an agent's demands about it. By taking blame one step back from obligation, we can hold blameworthiness even if in the end the person was not truly too obligated, due to their incapacity. It will all depend on the fairness of the demands. This exit denies the mentioned principle in Section 4.1 called LOLD. The counter intuitive side of it is that we can be properly blamed by things that we, in the end, were not truly obligated to do (or for not failing to do things we were not obliged to do). Indeed, this is a big bullet to bite.

Conclusion

This paper argues for the unexpected tension between some standard positions. OIC, objectivity about blame and obligation, and the empirical nature of human capability. If the arguments are correct, something about it must be rejected. This tension, I think, was neglected for long because a crucial difference, with robust methodological implications, was not observed, or not observed enough: the difference between token and type people. Taking that difference seriously puts us into the path of facing the epistemic challenges any person trying to judge the blameworthiness of another must face. This paper is not a defense of the claim that there is no fair blaming, rather it is a defense of the option five (Subsection 6.5), or it would had been had I had the space for that. In any case, these arguments presented in section V are the only arguments I

know of that reject the possibility of fair blaming, and consequently, large parts of moral responsibility, without touching upon debates about free will, except in the lightest of ways. This, if not anything else, may show that they merit philosophical attention.

Bibliographic References

- BASSFORD, A. D. 2022. Ought implies can or could have. *The Review of Metaphysics*, [S.l.], v. 75, n. 4, p. 779-807.
- BERGLUND, A. 2005. *From conceivability to possibility: an essay in modal epistemology*. 197 p. Aarhus, Denmark. Doctoral dissertation. Filosofi och lingvistik. Aarhus University.
- COATES, D. J.; TOGNAZZINI, N. A. 2013. The contours of blame. In: COATES, D. J. (Ed.); TOGNAZZINI, N. A. (Ed.). *Blame: its nature and norms*. Oxford: Oxford University Press.
- FRANKFURT, H. G. 1969. Moral responsibility and the principle of alternative possibilities. *Journal of Philosophy*, New York, v. 66, n. 23, p. 829-839.
- GOLDBERG, S. 2018. *To the best of our knowledge: social expectations and epistemic normativity*. Oxford: Oxford University Press.
- GOLDWATER, J. 2020. Six Arguments Against 'Ought Implies Can'. *Southwest Philosophy Review*, [S.l.], v. 36, n. 1, p. 45-54.
- GREENE, J. 2013. *Moral tribes: emotion, reason, and the gap between us and them*. New York: Penguin.
- GRAHAM, P. A. 2010. In defense of objectivism about moral obligation. *Ethics*, Chicago, v. 121, n. 1, 88-115, Oct 2010.
- HIERONYMI, P. 2020. *Freedom, resentment, and the metaphysics of morals*. Princeton: Princeton University Press.
- KING, A. 2019. *What we ought and what we can*. London: Routledge.
- MACNAMARA, C. 2013. Taking demands out of blame. In: COATES, D. J. (Ed.); TOGNAZZINI, N. A. (Ed.). *Blame: its nature and norms*. Oxford: Oxford University Press.
- PARFIT, D. 1971. Personal identity. *Philosophical Review*, New York, v. 80, n. 1, p. 3-27.
- PEELS, R. 2019. The social dimension of responsible belief: response to Sanford Goldberg. *Journal of Philosophical Research*, Virginia, v. 44, p. 79-88. <https://doi.org/10.5840/jpr201944150>.
- PEREBOOM, D. 2001. *Living without free will*. Cambridge: Cambridge University Press.
- ROSEBOOM T. J.; PAINTER R. C.; VAN ABEELLEN A. F.; VEENENDAAL M. V.; DE ROOIJ S. R. 2011. Hungry in the womb: what are the consequences?: lessons from the Dutch famine. *Maturitas*, Ireland, v. 70, n. 2, p. 141-145, Oct 2011. <https://doi.org/10.1016/j.maturitas>.
- SAPOLSKY, R. M. 2023. *Determined: Life without free will*. New York: Random House.
- SAPOLSKY, R. M. 2017. *Behave: The bestselling exploration of why humans behave as they do*. New York, Random House.
- SCANLON, T. M. 2013. Interpreting blame. In: COATES, D. J. (Ed.); TOGNAZZINI, N. A. (Ed.). *Blame: its nature and norms*. Oxford: Oxford University Press.
- SCANLON, T. M. 2008. *Moral Dimensions: permissibility, meaning, blame*. Cambridge, Mass: Harvard University Press.
- SHER, G. 2005. *In praise of blame*. Oxford: Oxford University Press.

- SIMION, M. 2024. *Resistance to evidence*. Cambridge: Cambridge University Press.
- SMITH, A. 2013. Moral blame and moral protest. In: COATES, D. J. (Ed.); TOGNAZZINI, N. A. (Ed.). *Blame: its nature and norms*. Oxford: Oxford University Press.
- STRAWSON, P. F. 2008. *Freedom and resentment and other essays*. London: Routledge.
- TIBERIUS, V. 2014. *Moral psychology: a contemporary introduction*. London: Routledge.
- VRANAS, P. B. 2024. Beyond ought-implies-can: impersonal obligatoriness implies historical contingency. *Journal of Ethics and Social Philosophy*, v. 29, n. 1, p. 29-61.
- VRANAS, P. B. 2018. "Ought" implies "can" but does not imply "must": an asymmetry between becoming infeasible and becoming overridden. *Philosophical Review*, New York, v. 127, n. 4, p. 487-514.
- VRANAS, P. B. 2018. I ought, therefore I can obey. *Philosophers' Imprint*, Michigan, v. 18, n. 1, p. 1-36, Jan 2018.
- VRANAS, P. B. 2007. I ought, therefore I can. *Philosophical studies*, Netherlands, v. 136, p. 167-216, Jul 2007.
- WEDGWOOD, R. 2018. The unity of normativity. In: STAR, D. (Ed.). *Oxford handbook of reasons and normativity*. Oxford: Oxford University Press.
- WEDGWOOD, R. 2016. Objective and subjective 'ought'. In: CHARLOW, N. (Ed.); CHRISMAN, M. (Ed.). *Deontic modality*. New York: Oxford University Press.
- WEDGWOOD, R. 2007. *The nature of normativity*. Oxford: Oxford University Press.
- WILLIAMS, B. 1973. *The problems of the self*. Cambridge: Cambridge University Press.