Synthese (2024) 203:163
https://doi.org/10.1007/s11229-024-04578-z

ORIGINAL RESEARCH

# Epistemic ownership and the practical/epistemic parallelism

**Jesús Navarro**[1,2]

Received: 27 September 2023 / Accepted: 24 March 2024


**Abstract**
We may succed in the fulfilment of our desires but still fail to properly own our practical life, perhaps because we acted as addicts, driven by desires that are alien to our will, or as "wantons," satisfying the desires that we simply happen to have (Frankfurt, 1988). May we equally fail to own the outcomes of our epistemic life? If so, how may we attain epistemic ownership over it? This paper explores the structural parallellism between practical and epistemic rationality, building on Williamson's (2002) suggestion that we should commence with successful performances as the foundation for both domains, be it action or knowledge. By highlighting the limitations of higher-order regulative approaches in epistemology, exemplified by Sosa (2007, 2011, 2015, 2021), the paper introduces a form of teleological epistemic constitutivism inspired by Velleman (2000, 2009). The proposal is that epistemic ownership is not attained in the mere pursuit of truth or knowledge, but requires furthermore a struggle to understand what we know.

**Keywords** Ownership · Practical rationality · Epistemic rationality · Knowledge-first · Virtue epistemology · Wanton · Understanding · Constitutivism.

## 1 Introduction

There are many things that we do, but that we do not fully recognise as our own actions. We are impelled by our impulses, driven by our desires, which give rise to intentions that effectively bring about events in a manner that isn't particularly devi-

✉ Jesús Navarro
jnr@us.es

[1] Department of Metaphysics and Contemporary Philosophy, University of Seville, Seville, Spain

[2] African Centre for Epistemology and Philosophy of Science, University of Johannesburgh, Johannesburgh, South Africa

Published online: 08 May 2024

ant or faulty — but we never really identified ourselves with those desires from the outset. Instead of functioning as autonomous agents, we assume the role of passive spectators throughout the entire process. Addicts often struggle but fail to overcome that status. So-called "wantons" are worse: they do not even care about any of this, simply fulfilling their desires through their actions, whatever those desires happen to be. Nevertheless, what their actions lack, often referred to as practical ownership, appears to be an essential condition for fully autonomous action.

This is a well-established idea in action theory, but is there anything similar in epistemology? If so, just as in the practical domain, this would be an important question affecting the nature of epistemically autonomous agents.

In this paper, my intention is to illuminate the issue of epistemic ownership by applying, mutatis mutandis, the insights gained from discussions regarding the nature of practical ownership. My argument hinges on drawing an analogy between two debates that independently originated in action theory and epistemology. In essence, the idea is that we can tell to Ernest Sosa in epistemology what J. David Velleman told to Harry G. Frankfurt in action theory.

What is it that Velleman told to Frankfurt? In brief, that higher-order regulation of our actions alone is insufficient to overcome the state of wantonness; instead, we need to pursue a goal that is constitutive of practical rationality, namely, doing something intelligible. My analogous message to Sosa, as well as to those epistemologists emphasizing that autonomous epistemic agents are those who manage to regulate appropriately their own cognition in the pursuit of knowledge, is that this strategy may not allow us to transcend epistemic wantonness. Rather than engaging in higher order regulation, what we must do is to strive to achieve a goal that is constitutive of epistemic rationality, namely, understanding what we know. To elaborate on this analogy, it is crucial to start off with the right foot on each side of the practical/epistemic parallelism. I aim to do so by following certain cues from Timothy Williamson's knowledge-first proposal.

The structure of the paper is as follows. In § 2 I introduce a debate in action theory originally instigated by Donald Davidson's (1980) causalism and subsequently advanced by Harry Frankfurt's (1988) higher-order theory of agency, whose limitations prompted J. David Velleman's (2000, 2009) teleological constitutivism. § 3 lays the theoretical groundwork for establishing a proper parallelism between practical and theoretical rationality. I follow Timothy Williamson's (2002) advice, emphasizing that we should not consider action and belief as the primary counterparts in that parallelism, but rather action and knowledge. In § 4, I narrate a parallel story in epistemology, commencing with Alvin Goldman's (1979, 1986) process reliabilism, akin to Davidson's practical causalism, and Ernest Sosa's (2007, 2011, 2015, 2021) virtue epistemology, comparable to Frankfurt's higher-order regulative theory. § 5 reveals that a certain reading of Sosa's epistemology shares the same limitation highlighted by Velleman in Frankfurt's action theory. In both cases, the complaint is that to attain epistemic ownership, we must aim at a goal that is constitutive, rather than merely regulative, of autonomous rationality. In § 6, I present my version of epistemic constitutivism, which is teleological in nature, and I contrast it with several existing alternatives. In a somewhat programmatic manner, I propose that the constitutive goal of autonomous epistemic rationality is understanding, mirroring the constitutive

goal of autonomous practical rationality in Velleman's framework, which is intelligibility. § 7 addresses objections to this view and argues that while both goals target the same endpoint, sense making, they approach it with opposing directions of fit. Finally, in § 8, I draw my conclusions.

## 2 The debate on practical ownership: Davidson, Frankfurt, Velleman

Davidson (1980) famously defended a naturalistic account of rational action, according to which practical rationalizations are a species of causal explanation: the existence of the appropriate rational-cum-causal connection between mental states and bodily movements is what differentiates, in his view, mere bodily movements from purposeful behaviour based on reasons. A prominent challenge for causalism is to specify the appropriate causal relations between desires and actions, which cannot be merely fortuitous. In a famous example: a waiter that has the desire, and even the intention, to annoy her boss by dropping the dishes; those mental states make her nervous, so that her hands sweat, and that makes her drop the dishes. We would not count this as an intentional action of hers. According to Davidson, the connection between intentional states, the agent's body, and her surroundings, must take place in non-deviant ways, and the agent's desires ought to manifest their respective paradigmatic motivational force, instantiating an act of 'choosing'.

For some time, the focal point of the discussion revolved mainly around whether Davidson's framework managed to save the day for causalism dealing with this and other puzzles — an open question to this day[1]. However, the landscape of this discourse was significantly transformed by Frankfurt (1988), who showed that even if Davidson were to meet the aforementioned challenges, a merely causal explanation of successful intentional action would still possess a fundamental flaw: it would fail to accommodate the role of agents themselves in practical deliberation. In essence, it's possible for actions to be causally linked to the agent's mental attitudes in non-deviant ways, yet the agent may not recognize herself as the originator of those actions. Frankfurt's examples in this regard often featured drug addicts and other impulsive individuals who, despite their actions being impeccably linked to their choices, still felt alienated in their own endeavors. In such scenarios, a Davidsonian rational-cum-causal explanation may be in force, but the agents themselves would be conspicuously absent from the narrative, assuming the role of passive spectators in their own lives.

This concern leaves us with a pressing issue regarding the nature of practical autonomy: we must elucidate our practical lives in a manner that bestows upon agents a form of ownership over their own existence. To address this, Frankfurt introduced a higher-order theory of human cognition. The core concept behind this theory is that there exist basic agents who possess only first-order desires that can be appropriately combined to generate first-order intentions, mediated by Davidsonian 'choices,' —

[1] We may still find some convincing reinstatements of the Wittgensteinian view that reasons are not causes, such as Tanney (2005). See Aguilar & Buckareff (eds.) (2010).

yet these individuals would still lack concern for the quality of their conative states. Frankfurt insightfully referred to such entities as 'wantons':

> The essential characteristic of a wanton is that he does not care about his will. His desires move him to do certain things, without its being true of him either that he wants to be moved by those desires or that he prefers to be moved by other desires. The class of wantons includes all nonhuman animals that have desires and all very young children. Perhaps it also includes some adult human beings as well. In any case, adult humans may be more or less wanton; they may act wantonly, in response to first-order desires concerning which they have no volitions of the second order, more or less frequently (Frankfurt, 1988, 16–7).

In the realm of practical ownership, the status of the wanton is even more precarious than that of the addict. This is because the wanton doesn't even exhibit any concern for what the addict fails to achieve. Frankfurt's critique of Davidson's causalist theory lies in its potential limitation to account for wantons, thus failing to capture the critical distinction that sets them apart from autonomous agents.

To elucidate the latter, Frankfurt proposed that we manifest higher-order attitudes, encompassing desires and volitions, which serve to regulate our own first-order conative states. Through this process, we generate second-order desires that transform into second-order volitions, thereby possessing the capacity to govern our first-order attitudes. This higher-order perspective serves as the vantage point from which agents can establish a rightful place for themselves in their intentional actions, ultimately achieving practical ownership of their lives.

While Frankfurt's regulative strategy may appear plausible on the surface, it inevitably confronts a possible vicious regress — the endemic disease of higher-order theories. The agent who initially failed to recognize herself in her first-order desires may indeed generate second-order desires and volitions, but she may also falter in recognizing herself in these second-order states, potentially necessitating the creation of third-order states, and so forth. In the end, higher-order states may essentially function just as mental attitudes endowed with rational-cum-causal capabilities, seemingly lacking any intrinsic properties to halt the regress.

Frankfurt was certainly aware of this concern, which he sought to address by asserting that "When a person identifies himself decisively with one of his first-order desires, this commitment 'resounds' throughout the potentially endless array of higher orders" (1988, 21). In essence, his argument rested on the idea that these higher-order attitudes need not be actively instantiated but merely held dispositionally. The strength of the original commitment would theoretically extend throughout any order.

However, the issue with Frankfurt's solution lies in that it merely has the capacity to alleviate the symptoms without curing the disease. The problem raised by the vicious regress is not practical, but theoretical. Even though the dispositional solution may appear to resolve the impossibility of implementing the regress in practice, it doesn't fully address the root cause of the problem — the fact that the higher-order attitude seems to lack what the lower-order one was missing. The agent could still find herself as a passive spectator to the activation of her higher-order dispositions,

leaving it unclear why higher-order regulation inherently carries the seed of practical autonomy.

To address this deficiency, we require an alternative diagnosis of wantons, one according to which their lack of ownership would not result from the fact that they fail to regulate their first order desires from a higher order. Frankfurt's merely regulative solution should be substituted by one that inherently accommodates the presence of the agent herself within the deliberative process across all orders, both lower and higher.

A significant step in this direction was taken by J. David Velleman in his versions of practical constitutivism (2000, 2009). Velleman's perspective provides an alternative framework for understanding practical ownership, grounded in the notion that autonomous action possesses a constitutive goal, which, in his view, is sense-making[2]. Whenever an agent acts autonomously, she is not solely focused on satisfying a specific desire; she is also engaged in the pursuit of intelligibility, both for herself and others:

> You govern yourself, it seems to me, when you seek to grasp yourself as part of an intelligible world and consequently gravitate toward being intelligible. […] The appeal of this view, for me, is that it locates autonomy in a part of the personality from which you truly cannot dissociate yourself. This part of your personality constitutes your essential self, in the sense that it invariably presents a reflexive aspect to your thinking: it invariably appears to you as "me" from any perspective, however self-critical or detached. […] You can dissociate yourself from other springs of action within you, by reflecting on them from a critical or contemplative distance. But you cannot attain a similar distance from your understanding, because it is something that you must take along, so to speak, no matter how far you retreat in seeking a perspective on yourself (2000, 30).

According to Velleman, intelligibility is not merely a contingent goal that the autonomous agent may or may not choose to act upon to regulate her behavior. An action is deemed autonomous precisely because it was carried out in pursuit of this goal.

The idea that all autonomous actions are inherently oriented toward a constitutive goal, such as intelligibility, offers a principled approach to addressing Frankfurt's objection to Davidson. Merely non-deviant causal connections are insufficient for agents to assert ownership over their practical lives. In addition to this, autonomous agents must act under the guidance of a self-conception, striving to perform actions that align with their self-identity and make sense within the given social context (Velleman, 2009). While being regulated by higher orders of assessment can be significant, it holds value primarily for its instrumental role in the pursuit of intelligibility. An agent does not attain autonomy merely by forming higher-order volitions but by

[2] Velleman's characterization of the constitutive aim of action has evolved over time. In his 1996 paper, 'The Possibility of Practical Reason,' he initially aligned with Kant by defining this aim as "autonomy itself" (2000, 193). However, he later shifted from this view and identified the goal as "self-knowledge," in the sense that rational agents seek "to know what we are doing" (2000, 26–7). Gradually, his formulation has evolved to encompass notions of self-understanding, intelligibility, and "making sense" (2009, 26). It is this latter formulation that I prefer for reasons that will become evident in the final section.

purposefully aiming to act in a manner that makes sense, considering her own identity and the circumstances in which she finds herself.

Henceforth, Velleman's alternative solution to Frankfurt's puzzle is to identify a *constitutive goal* of autonomous action. The distinction between Velleman's teleological constitutivism and other deontological varieties, such as Korsgaard's (2008), lies in the nature of its objective as a goal rather than a norm[3]. And the fact that the goal constitutes autonomy distinguishes his view from regulative theories. The goal is constitutive of practical ownership and full autonomy not in the sense that it ought to be attained, but that it must be pursued[4]. One's autonomous actions are constituted by the struggle to make sense, not merely regulated by it[5]. To be "constitutive" implies that this goal defines the very activity itself. An agent is not truly engaged in the activity unless she actively pursues this goal, and her engagement in the activity is contingent upon her pursuit of this goal. This is akin to how one only plays chess in so far as one endeavors to checkmate the opponent, regardless of whether success is ultimately achieved[6]. Crucially, the constitutive goal is not directed at *improving* performance in accordance with external standards; it is the very internal target that defines the activity.[7]

From this perspective, the issue with the addict lies in her inability to attain a constitutive goal she nonetheless aims at. She successfully fulfills her desires in non-

[3] For a detailed exploration of the distinction between teleological and deontological constitutive demands, see Shah (2003), Shah and Velleman (2005), and McHugh (2011, 370-1). It's worth noting that Mitova (2016, 201), although she explicitly draws this distinction and advocates for a teleological form of constitutivism, situates it within a broader deontological agenda. She does so by delving into the question of why we should care (motivation or goal) for normative constraints (deontological requirements). While Mitova's interpretation of Velleman has significantly influenced my work, our perspectives diverge at this juncture and in some other instances, that I will indicate in their due place.

[4] As suggested by Grajner and Schmechtig (2016), a norm of the form 'If C then X is N" would require some condition C to be met for the action X to have the appropriate normative status N. Applied within the context of our discussion, if an action X is deemed intelligible, it would imply that one is in the relevant normative state (rationally justified, permitted, obliged, and so on) to execute it. This normative interpretation of the constitutivist requirement would imply that an agent can only act autonomously if she successfully satisfies the condition of being intelligible. In my perspective, such a requirement may be overly restrictive.

[5] I say "not merely" because the distinction between regulation and constitution is not one of opposition. Constitutive goals may also serve to regulate performances, as elaborated by Wedgwood (2002, 268).

[6] At various points, Velleman introduces the view as a definition of the constitutive goals of *rationality*, while at other junctures, he applies it to *autonomy*. My inclination is to interpret it as a statement concerning the latter, not the former. In accordance with the perspective I will defend, what the wanton fails to achieve is the status of an *autonomous* agent, even though she may attain the status of a rational agent. This distinction may bear significance in discussions regarding various accounts of rationality, but I do not consider it essential in addressing the issues of autonomy and ownership, as I intend to do here.

[7] Mitova's assertion is indeed compelling: the pursuit of sense-making represents "the distinguishing mark of agency" (2016, 207) because one cannot readily forsake this goal while retaining their status as an agent (211). I would further refine this by specifying that it pertains specifically to *autonomous* agents — one may not truly be an autonomous agent unless they aim at this constitutive goal. From my perspective, wantons do engage in actions, albeit without actively pursuing the constitutive goal of making sense. Consequently, they qualify as agents, but not as autonomous ones. This may align with what Mitova terms "full-blown action" (212). However, I perceive an ambiguity in Mitova's position regarding whether intelligibility is the constitutive goal of action, period, or whether it denotes something more stringent, as I interpret it — pertaining to autonomous or "full-blown" action.

deviant ways, yet she ultimately falters in the pursuit of intelligibility, as if persistently questioning the purpose of her own actions. But the problem with the wanton is more profound; she does not even aim for that constitutive goal, and the mere instantiation of higher orders of assessment, *pace* Frankfurt, does not alter this fundamental deficiency. In stark contrast with both the addict and the wanton, the autonomous agent actively pursues not only the fulfilment of her contingent desires, but also the constitutive goal of sense-making, which ensures her presence in the performance and her ownership of the results[8].

## 3 Drawing the practical/epistemic parallelism right: Williamson

The remainder of the paper delves into the field of epistemology, aiming to draw a paralellism as closely aligned as possible with the narrative presented earlier in the domain of action theory. However, it is crucial to embark on this endeavor with care, and this section is dedicated to setting the appropriate groundwork for that purpose.

Structural parallelisms between practical and epistemic rationality, and their respective conative and cognitive attitudes, are generalized nowadays in the theory of rationality — see Singh (2019). This parallelism often originates from the comparison of action on one side and belief on the other, as the initial relata. For example, in the debate on doxastic freedom it is frequently asserted that we may not believe at will as, allegedly, we are able to act at will — e.g., see Rinard (2019) or Vitz (2021)[9]. Similarly, discussions on the "basing relation" define it as a connection that exists between an action and its reason in the practical realm, or between a belief and its justification in the epistemic realm — e.g., see Blake-Turner (2021). Accordingly, the concern that I address here, namely ownership, has been framed in epistemology under the rubric of *doxastic* ownership, focusing on our relationship with beliefs — as will be detailed in § 6 below — whereas its equivalent within the practical side of this parallelism has revolved around our relationship with our actions — as described in § 2 above.

However, despite their long-standing relationship, belief and action may not be the most suitable counterparts, and we might find better matches for each of them. In this regard, I will adopt the alternative approach proposed by Williamson (2002). According to his perspective, if we initially pair action and belief in the parallelism between practical and epistemic rationality, we will inevitably encounter a significant mismatch: action inherently implies success, whereas belief does not. As discussed

[8] The distinction between the wanton and the autonomous agent does not depend on the former's lack of appreciation for the *value* of the constitutive goal of action. Instead, it lies in her failure to actively *pursue* that goal, regardless of whether she values it or not. For instance, an agent may play chess guided by its constitutive aim (i.e., checkmating the opponent's king) but do so half-heartedly and without appreciating the value of winning. Similarly, an agent may engage in autonomous action, pursuing the goal of understanding, even if she does not wholeheartedly *value* that goal. And, conversely, an agent would persist as a wanton if, despite appreciating the value of the constitutive goal of autonomous agency, she neglected to actively *pursue* it. I am grateful to a referee for *Synthese* who urged me to explicitly articulate this point, as it represents a crucial detail.

[9] For an alternative arrangement of the terms in this debate that aligns more closely with the perspective I advocate for here, see McHugh (2014).

in the previous section, action not only involves success but also entails the successful execution, in causally appropriate ways, of some intentional attitude, typically a desire. When success is not achieved in a performance, it ceases to be classified as action and is labeled as an "attempt" at best. In contrast, beliefs are propositional attitudes that need not be successful, i.e., true, to retain their status as beliefs. Moreover, they do not necessarily have to be causally linked in any manner, whether deviant or not, to their representational content — *pace* semantic externalists. Therefore, initiating the practical/epistemic parallelism by aligning belief with action would be like buttoning up one's shirt starting with the wrong button.

Here is Williamson's alternative proposal:

> Knowledge and action are the central relations between mind and world. In action, world is adapted to mind. In knowledge, mind is adapted to world. When world is maladapted to mind, there is a residue of desire. When mind is maladapted to world, there is a residue of belief. Desire aspires to action; belief aspires to knowledge. The point of desire is action; the point of belief is knowledge (Williamson, 2002, 1)[10].

I will not endorse here the rest of the tenets of Williamson's knowledge-first program, but this is one I will: *what stands to action on the epistemic side is not belief, but knowledge, as the fulfilment of the intentional state of belief when properly related to the fact it represents*.

Recall that Davidson's causalist account was primarily concerned with intentional action, which, in his perspective, constituted the rational-cum-causally appropriate fulfillment of certain desires. Consequently, we should now seek an analogous causalist account of knowledge within the realm of epistemology, and a prominent candidate for this role is Alvin Goldman's reliabilism. Goldman initially posited that knowledge equates to true belief caused in the right way (1967), a position that later evolved into what we now recognize as "process reliabilism" (1979, 1986), wherein knowledge is defined as true belief caused by a reliable cognitive process. In this context, reliability denotes the property of being conducive to truth in a sufficient proportion of relevant cases.

Just as Davidson's causalism aimed to elucidate why rational explanations transcend mere rationalizations, serving as genuine explanations for why agents act as they do, Goldman's reliabilism sought to address a similar issue pertaining to epistemic internalism. Specifically, it grappled with the challenge that, although a purely internalist account of justification might offer a coherent narrative about how an agent justifies the belief they form, it would not establish why that belief qualifies as knowledge unless it is causally linked to the fact it purports to represent.

[10] Williamson has made a recent modification to his proposal, replacing desire with intention in the structural analogy (2017), but he does so in line with the demands of his knowledge-first project — as discussed in Miracchi and Carter (2022). This is a framework I do not subscribe to in this context though. This substitution leads him to adopt an intellectualist stance that I find unconvincing, specifically the attempt of understanding desire as a variant of belief concerning the goodness of a particular state of affairs (2017, 165).

Despite its considerable influence, Goldman's position, much like Davidson's, faced several theoretical challenges. A prominent one was explaining why causal connections that are merely deviant fail to produce knowledge. Goldman's causalist account was specifically designed to address standard Gettier cases, which are the epistemic equivalents of actions resulting from deviant causal chains. In these cases, according to Goldman, the appropriate causal link between the fact and the belief is either absent or entirely inappropriate. However, just as Davidson was at pains to deal with causal deviancy in a principled way, it remains unclear why a strictly reliabilist account of causal deviancy is not overly ad hoc. This challenge becomes especially apparent in situations where agents that are generally deemed reliable find themselves in environments where the likelihood of error is notably high, as exemplified by the famous fake-barn scenario (Goldman, 1976, 772).

Now, as Frankfurt did with Davidson's causalism, we may simply grant reliabilists that their account of knowledge will eventually have resources to deal in a principled way with causal deviancy. Even so, the concern that I want to raise is that a further challenge would still await, one that would mirror the structural problem Frankfurt highlighted concerning autonomous action: namely, the difficulty of the causalist account of knowledge to explain how the agent may achieve ownership of her cognitive performances as an autonomous rational being. While it may not be required for an agent to achieve this level of ownership concerning each and every one of her cognitive accomplishments, it is reasonable to anticipate that, in certain instances, autonomous agents could — or even should — attain such ownership[11]. And the worry is that the reliabilist framework alone will fall short in accounting for this, as it merely characterizes agents who may possess knowledge without necessarily meeting the requisites of epistemic ownership. Consequently, in such an account, intellectual autonomy will remain elusive[12].

But how exactly may agents be epistemically successful, in the sense of attaining beliefs that are not only true, but appropriately related to the facts in rational-cum-causal ways, and still fall short of taking ownership of those beliefs? The next section will flesh this possibility out.

[11] I leave it open for discussion whether, for instance, functional beliefs produced by the automatic operation of our perceptual system would be subject to the requirement of being *owned* by the agent, or if such ownership is even feasible in those cases. The recognition that epistemic ownership is a possibility, and a significant aspect of our epistemic lifes that demands an explanation, is enough for my point.

[12] A similar concern has been articulated by Broncano-Rodríguez and Vega-Encabo (2011) in terms of the agent's "epistemic engagement". McHugh (2013) also delves into a closely related topic, but he frames it in terms of "doxastic responsibility". However, I have objections to both parts of McHugh's lable. On the one hand, I don't believe that responsibility is the central issue here. This is because an agent may meet the conditions for epistemic ownership but still hold their beliefs in an irresponsible manner, as per Scanlon's substantive sense of being unable to respond to the appropriate reasons (1998, 22, 248). From my perspective, ownership represents a necessary — though possibly not sufficient — condition for autonomous agency, but an agent can be autonomous while falling short in terms of substantive responsibility demands. On the other hand, due to the reasons outlined in § 3, I find it problematic to articulate this concern in doxastic terms — see also below, note 16.

## 4 Epistemic addicts and epistemic wantons

I will follow a similar strategy to Frankfurt's, initially identifying agents who fall short of achieving ownership despite their efforts (epistemic addicts), and addressing then the more challenging case of agents who are indifferent to the possibility of such failure (epistemic wantons)[13].

The counterpart of Frankfurt's addict in the realm of epistemology is a specific type of epistemic akratic agent: one who knows but fails to recognize herself as the owner of her knowledge because she believes she shouldn't hold the beliefs she knows to be true[14]. Such agent would consistently form true beliefs (as she wouldn't know otherwise, according to causalism) while simultaneously thinking she lacks compelling reasons to maintain these beliefs. This agent would be considered epistemically flawed, much like Frankfurt's addicts were considered practically flawed. Frankfurt's addict successfully achieved their intentions, satisfying the criteria for intentional action, but failed to align their desires with what they truly wanted. Similarly, this specific form of epistemic akratic attains what, according to the causalist account, qualifies as knowledge (i.e., belief that is true as a result of a reliable cognitive process), yet she fails to believe in accordance with what she thinks she should[15].

But things could get even worse. Our second epistemic pathology, that of the epistemic wanton, wouldn't be a case of failure like the akratic, but rather that of someone who doesn't even make an attempt to achieve what the epistemic addict failed to attain. Recall that Frankfurt complained tha the practical wanton acted on her desires and successfully performed intentional actions but didn't concern herself with the quality of those desires. Similarly, our epistemic wanton would perhaps manage to form beliefs that are reliably true, but would lack a personal stance about the quality of her belief-formation processes.

The worry at this point is that a basic causalist account may only be able to account for one type of epistemic agent: epistemic wantons. These individuals would be present in the realm of epistemic deliberation but would lack autonomy over their belief-forming processes, just as the practical wanton was absent from practical deliberation as a practically autonomous agent.

[13] With his idea of a wanton Frankfurt originally expressed a concern for both desires and beliefs, although his ellaboration of it was mostly focussed on the practical aspect. He explored an idea that is akin to the one of an epistemic wanton with his influential concept of "bullshit", as the kind of discourse where speakers manifest a despicable neglect for the epistemic quality of their own speech (1988, Ch. 10). However, bullshit is a matter of defective communicative intentions, and not of a problematic normative evaluation of our own cognitive lives in epistemic terms. What the bullshitter appears to lack is not epistemic ownership of their beliefs but rather a form of honesty in her communication with others.

[14] This relies on the denial of the KK principle, in line with Williamson's anti-luminosity arguments (2002, Ch. 4).

[15] I rely here on the idea that epistemic akrasia can be rational — see Horowitz (2014, 2.1). Importantly, if the agent believed she shouldn't form a particular belief due to the presence of significant epistemic defeaters, then she would not possess knowledge, and the scenario would not exemplify an epistemic addict. The situation in question should instead involve an agent who has high rational confidence in both a belief (p) and some misleading evidence that she shouldn't believe it — see Lasonen-Aarnio (2014).

But even this still is somewhat intangible, like a planet we can infer exists only through its gravitational influence on other celestial bodies. Here is an imaginary case that may put some more flesh in the character:

> BILL. In his role as a journalist, Bill is tasked with composing an article that provides an overview of the social and political conditions in a distant country he is unable to visit. His research process entails a comprehensive review of numerous information sources, including official reports, newspapers, expert opinions, layman perspectives, personal interviews, internet blogs, tweets, and more. Bill is fully aware that some of the information he encounters may be erroneous, and he acknowledges the presence of contradictions within the collected data. Despite this awareness, his method involves assimilating each piece of information into his body of evidence without ever scrutinizing the reliability of its source. When confronted with questions about the trustworthiness of his sources, he responds by asserting that lack of reliability does not necessarily imply falsehood, and there may be "some truth" even in unreliable sources. Consequently, all information is given equal weight in his research, consistently dismissing concerns related to source reliability as negligible, irrelevant, or burdensome. In his view, the larger the volume of evidence, the closer he will be to truth.
> It turns out that a substantial portion of the information Bill compiles is perfectly accurate, and not by luck, but because his information sources happen to be reliable — something Bill has no idea about. He meticulously gathers this information, identifies and resolves potential contradictions, and ultimately delivers an article that effectively and truthfully portrays the social and political landscape of that country.

In a sense, Bill finds himself in a more favorable position than the epistemic addict. Unlike the addict, he is not compelled impulsively to believe what he knows. Rather, Bill diligently seeks to ascertain the truth regarding the general conditions of the distant country, and he does so in a manner that aligns with the evidence he has amassed. However, Bill's peculiarity seems to be that he exhibits no concern whatsoever for the reliability of his information sources. While he places great emphasis on truth-seeking, as evidenced by his extensive collection of information, he appears entirely unconcerned about reliability. Fortunately, the world has been accommodating by furnishing Bill with an environment that, for the most part, is conducive to reliable information, and thanks to that he acquires a substantial body of knowledge about that remote country. In the light of this, one may wonder whether there is something inherently problematic, defective, or lacking in Bill's epistemic conduct.

There is indeed: Bill is an epistemic wanton. He is not really the owner of his epistemic life. He is just compulsively accumulating evidence without any concern for its quality. In a less auspicious informational environment, Bill would not have achieved knowledge, yet he never took the initiative to assess whether he was in such an environment. Frankfurt's practical wanton successfully fulfills her desires through non-deviant means but simply doesn't care about having the desires she ought to have. Similarly, Bill achieves knowledge by forming true beliefs through non-deviant

processes, but he disregards the epistemic quality of his beliefs, which represents a significant deficit in his exercise of intellectual autonomy. Bill's insensitivity to the epistemic quality of his evidence means he fails to truly own the success of his cognitive attitudes, regardless of how reliably they may have performed.

Is Bill merely an outlandish product of philosophical imagination? In his extreme form, he certainly is. However, Frankfurt's insight holds that wantonness exists on a spectrum. We, as adult humans, can exhibit varying degrees of epistemic wantonness, acting in ways that are more or less similar to Bill's conduct. This is especially relevant when considering our behavior on the Internet and social networks, where many consume information without filtering for reliability, assuming that there might be "some truth" in it. The sheer volume of available information can create a false sense that filtering is unnecessary, and individuals hope that accurate information will eventually emerge from the sea of falsehoods simply by keeping their eyes wide open. With a bit of luck, this gullible approach may yield accurate information, or even knowledge. But more often than not, it doesn't.

## 5 A regulative approach: Sosa

Now that we have envisaged the idea of an epistemic wanton, how do we manage to overcome that status? To begin with, one might expect a solution in epistemology structurally analogous to Frankfurt's. Agents would act as autonomous inquirers who take proper ownership of what they know by striving to attain a higher-order perspective. This response draws inspiration from the most prominent variety of virtue reliabilism: Ernest Sosa's telic virtue epistemology (2007, 2011, 2015, 2021)[16]. Sosa famously distinguishes two forms of knowledge: animal and reflective. Animal knowledge, or brute animal cognition, involves apt belief — belief that manifests the agent's cognitive faculties and virtues. This is something that, even if young children and non-human animals can achieve, is nonetheless absent in Gettier cases. However, more sophisticated agents can attain reflective knowledge, which arises from meta-apt belief. In this case, the agent possesses a correct perspective on the situation they are in. When this higher-order perspective successfully regulates the formation of first-order beliefs, the agent may attain true beliefs that are not only apt and meta-apt but also apt *because* they are meta-apt. In such instances, the agent would achieve

[16] I do not simply attribute this response to Sosa because his virtue reliabilism is not primarily intended to address the issue of epistemic ownership, but to clarify the nature of knowledge. I chose Sosa's view as my point of focus because I believe it is the best available account on that subject, as outlined in my specific defense in Navarro (2015, 2016). However, my contention is that Sosa's approach may provide only a regulatory strategy, akin to Frankfurt's one, which, when applied to the problem of epistemic ownership will fall short of fully explaining it. Accordingly, instead of as an objection, what follows should be seen as a suggestion for development of Sosa's framework. On the other hand, Conor McHugh, does address the issue of ownership of doxastic attitudes through a regulatory approach. According to McHugh's perspective, doxastic ownership and responsibility are acquired by exercising "epistemic guidance control" over our first-order doxastic attitudes (2013, 142-3). I could have also chosen to target McHugh's regulatory theory, but I opted not to do so because he frames the problem in doxastic terms, utilizing the practical/epistemic parallelism in a manner that I have rejected in § 2.

what Sosa (2011, 11) dubs "knowledge full well," namely animal knowledge enlightened and guided by a reflective stance concerning its reliability.

The similarity between Sosa's account and Frankfurt's is quite striking. Both provide descriptions of basic agents, such as young children and non-human animals, who successfully achieve their first-order attitudes. Just as Goldman's reliabilism, according to Sosa, falls short in explaining how agents can attain full knowledge, Davidson's causalism, according to Frankfurt, fails to fully explain autonomous intentional action. Moreover, not only are their diagnoses similar, but their proposed remedies are analogous as well. Sosa's solution to the limitations of process reliabilism is structurally akin to Frankfurt's solution to the deficiencies of Davidson's causalism. In both cases, they turn to the higher-order capacities of the agents. Just as Sosa suggests that a higher-order perspective concerning the agent's reliability (reflective knowledge) should regulate their first-order cognition (animal knowledge), in the case of Frankfurt, he holds that higher-order desires should generate higher-order volitions capable of regulating an agent's first-order desires[17].

It is reasonable to expect then that, if Sosa were confronted with the problem of the epistemic wanton, he would likely have a solution similar to Frankfurt's. Bill's shortcoming, from Sosa's perspective, would be that he failed to form higher-order attitudes regarding the reliability of his sources. He should have aligned his beliefs with a higher-order reliability assessment, but he failed to do so, being guided solely by his first-order attempt to get things right.

Now, if we've followed the analogy this far, why not taking one step further? A response inspired by Sosa would seem to be vulnerable to the same objection that threatened Frankfurt's theory. Recall, the objection questions whether merely ascending one level higher in cognitive orders truly makes a qualitative difference, raising the concern that the strategy may initiate a vicious regress. As Stephen Grimm plainly noted, "the question remains as to how brute reactions on the first-order level become performances of a significantly different kind with the addition of a second level brutely responding to stimuli from below" (2016, 193).

Sosa may attempt to address the threat of regress, much as Frankfurt did, by proposing a dispositional claim: in practice, there's no need to instantiate an infinite number of epistemic orders. It would be sufficient for the agent to be *disposed* to go higher if the situation demands it, whithin the sensible limits of her competence (see Sosa, 2015, 86 n25). This way, the agent's strong commitment to reliable truth would "resonate," to borrow Frankfurt's analogy, in every order of assessment she engages in.

But let's recall the response I provided earlier to Frankfurt: the regress itself isn't the illness, but just a symptom. The regress starts because nothing inherently enlightening occurs just by moving up the orders of assessment. The core problem lies in the fact that when the first order lacks illumination, a higher-order strategy suggests a way to shed light on it — by building a larger room around it to illuminate it. But this would only work if the new room comes with some light that the former one did

[17] Especially pertinent to this parallelism is Chap. 1 of Sosa (2015), in which he contrasts his account of knowledge, based on the idea of competence *manifestation*, with causal explanations, either of intentional action, as proposed by Davidson, or of perception, as explored by Grice.

not possess. The question then arises: why is the higher order of epistemic assessment more illuminated than the first one?

To further develop this objection, let's revisit the case of our epistemic wanton, Bill. Imagine that we are his supervisors and, out of the worry that he doesn't have any concern about the reliability of his information sources, we advise him to investigate each possible epistemic source thoroughly. Bill might learn, for example, that whenever he receives information from a particular witness in a social network, he should inquire into her trustworthiness. Similarly, when he reads new information in a local newspaper from the distant country in question, he should check its credibility. This could involve searching for additional information on a rating website or asking other informants for their perspective. The idea is to show Bill that he shouldn't indiscriminately accept any information that comes his way but should always be mindful of what his meta-sources say about the reliability of his primary sources.

However, would we be satisfied if Bill simply followed these instructions? I believe not, because Bill could interpret these instructions literally, meaning he would search for information about the quality of his sources *while not caring at all about the reliability of the meta-sources he encounters*. If he showed no concern whatsoever about the reliability of the sources that vouched for the trustworthiness of a particular testifier or newspaper, there would still be something amiss about his conduct. He would have learned to ascend one level higher in the hierarchy of assessments, but at this higher level, he would still act as a mere wanton, indiscriminately including any information he obtains into the bucket of his meta-evidence. What transformative change would elevate Bill from the category of a mere epistemic wanton if he remained just as nonchalant about the quality of his second-order evidence as he was about the quality of his first-order evidence?[18]

At this point, we can introduce the third step of our analogy, which provides a potential solution to Sosa's puzzle similar to the one offered by Velleman to Frankfurt's puzzle. Epistemic ownership, as a prerequisite for epistemic autonomy, may not be achieved solely by forming regulative attitudes that ascend higher and higher in the orders of assessment. Instead, it might be attained by actively pursuing a goal that is constitutive of epistemic rationality, one that, when genuinely pursued, makes a qualitative difference regardless of the cognitive order the agent is considering. But what may that goal be?

## 6 A constitutivist alternative

Insofar as the analogy between these debates in action theory and epistemology has been appropriately drawn in the previous sections, we are now faced with a clear objective: the definition of a constitutive goal for autonomous epistemic rationality, akin to the constitutive goal of autonomous practical rationality. Pursuing this goal

[18] It's worth noting that the concern about an agent lacking ownership of what they know, as I am presenting it, is not dependent on whether the agent acquired that knowledge from a testimonial source, as opposed to gaining it first hand. In my view, arguments about the demands of epistemic autonomy are too often affected by a bias against knowledge acquired from testimonial sources — for instance, in Lynch (2016) or Sosa (2021, 3–16).

should not be regarded as an optional or discretionary matter, wherein the agent may choose to regulate her cognition or not, but as a goal inherently pursued by any rational agent who seeks, as such, to figure out how the world is[19]. The question that arises is: what form should this constitutive goal take?

One initial candidate for this constitutive goal is truth, aligning with Bernard Williams' renowned adage, "belief aims at truth" (1973, 136). This tenet suggests that whenever an agent engages in deliberation regarding her beliefs, her conduct should be guided by the pursuit of truth. The resulting mental state qualifies as a belief only to the extent that this goal is pursued. In the realm of epistemology, many constitutivist theories have been proposed in line with this perspective. Some frame it teleologically, postulating truth as the ultimate *aim* or objective of belief[20], while others express it deontologically[21], treating truth as the *norm* governing belief. However, it should be noted, due to the analogical alignment established in § 3, that this candidate must be dismissed. The reason lies in the fact that this class of proposals is concerned with a different phenomenon from the one under examination here: belief. The challenge posed by epistemic wantons does not revolve around determining what constitutes beliefs, as epistemic wantons do indeed form beliefs—even if not *their* beliefs, in a sense to be elucidated. Bill's cognitive attitudes do not constitute examples of make-belief or self-deception; they are genuine beliefs, which pursue its constitutive goal or norm — if Williams is right, truth. Williams' maxim holds considerable sway in this respect: beliefs cannot be formed in ways that disregard the quest for truth. Nonetheless, even if this holds true, we still require an additional constitutive goal that accounts for the presence of agents themselves, in their capacity as autonomous agents, in the process of belief formation in such a way that they can take ownership of those beliefs. The sought-after constitutive goal cannot be truth, just as practical success, as the fulfillment of desires, even if aligned with the pursuit of the good, fails to suffice for an adequate account of practical ownership.

Another possible contender for the constitutive goal of epistemic rational deliberation is knowledge, aligning with Timothy Williamson's claim that "belief aims at knowledge (not just truth)" (2002, 47)[22]. This motto, stronger than Williams', also holds a certain appeal, but knowledge should be discarded for precisely the same reason as belief, if the previous rationale is correct: our epistemic wanton not only aspired to knowledge but actually *achieved* it, both in the sense of animal knowledge, and, after the variation, as reflective, or even full knowledge. Nonetheless, our appre-

[19] To the extent that the required element is a goal, essentially serving as a form of motivation, one could argue that the demand exhibits a closer affinity to virtue responsibilism rather than virtue reliabilism. Accordingly, it is plausible that Sosa would posit that such a statement aligns itself with intellectual ethics, a domain distinct from the core considerations of what he terms "gnoseology" (Sosa, 2021, 17–48). However, the manner in which the autonomous agent must pursue the constitutive goal, as previously noted in footnote 8, does not revolve around an appreciation of its value, a facet most evidently falling within the purview of intellectual ethics, but around the agent's attempt to achieve that goal, which is a matter of *telic* assessment. I will revisit this point in § 7.

[20] See Velleman (2000, 244), Sha & Velleman (2005, 499), or Chrisman (2016).

[21] See Wedgwood (2002, 272), or Shah (2003). As I earlier said, Mitova (2016) offers a somehow mixed account.

[22] See for instance McHugh (2011) for a similar claim.

hension persisted: Bill could still resort to higher-order deliberation wantonly, much as he originally considered his first-order evidence wantonly. In other words, the wanton may be driven by the quest for knowledge, irrespective of how demanding we conceive it to be, and still remain a wanton in so far as he is not also motivated by the goal that puts her in the picture, as the owner of this knowledge. Or, at the very least, this is how Velleman's objection to Frankfurt would be articulated when applied to Sosa. Merely aiming at knowledge will not enable the epistemic wanton to transcend the state of wantonness.

If neither truth nor knowledge serves as the constitutive goal of autonomous epistemic rationality, then what may it be? What is it that the wanton fails to aim for, thereby preventing him from truly owning what he knows?

My main objective in this paper has been to show that there must be *some* constitutive goal of epistemic rationality, without definitively determining what that goal may be. Nevertheless, I am compelled to put forth what I believe to be a plausible candidate for such a goal: understanding[23]. The epistemic wanton will persist as such unless she strives for understanding of what she knows, making sense of her knowledge. When we acquire knowledge, we collect individual pieces of a puzzle; however, understanding entails arranging these pieces together, attaining an enhanced perspective where all these known propositions coalesce meaningfully. Such understanding seems to be a stronger contender for the sought-after constitutive goal of autonomous epistemic rationality.

I base this idea on an approach to understanding that views knowledge as insufficient in itself for true comprehension, as outlined by Zagzebski:

> Understanding is not a state directed to a single proposition at all. This is not to deny that there is a sense in which one can be said to understand a proposition *p*. But the understanding of *p* is not directed primarily at *p* itself. One understands *p* as a part of and because of one's understanding of a system or network of truths (1998, 49).

Various attempts to explain the distinction between mere knowledge and understanding have been proposed, but I don't need to endorse any particular one here[24]. All my proposal requires is the possibility that someone can know a proposition (regardless of how stringent the requirements for knowledge are in terms of safety or reflective attitudes) and still fail to understand it, or at least fail to fully understand it in the context of its subject matter[25]. This would clarify why the wanton may strive to know a proposition but remain unconcerned about something she should necessarily be aiming for as an autonomous agent in that same rational process — namely, understanding.

[23] I am grateful to Santiago Echeverri for pressing me to pursue this line of thought.

[24] In favour of a strong distinction are Zagzebski (2001), Grimm (2001), or Pritchard (2014). Others are more inclined to account for understanding in terms of knowledge, while preserving the distinction, with different degrees of reduction — see Grimm (2006), Greco (2014), Lynch (2016), or Kelp (2021).

[25] I grant that propositional knowledge requires at least linguistic understanding of the sort needed for genuinely believing the relevant proposition. However, this attainment may occur even in the absence of a deeper understanding of why the fact in question is the case and how it fits into an overall picture of events.

Understanding goes beyond logical coherence. Two propositions, $p$ and $q$, may exhibit perfect compatibility in terms of logical coherence, yet it might still not make sense for them to be true simultaneously. Recognizing this demands something beyond acknowledging the truth of $p$ and $q$, something more than searching for evidence related to $p$ and $q$, or even seeking evidence for ($p$ & $q$): it requires establishing these facts within profound explanatory relationships with each other, and possibly with other pertinent propositions, thereby making sense of them within the broader context of one's perspective on the subject.

Sosa's invitation to regulate first-order cognition based on a higher order stance about one's reliability can be viewed as an instance of this. Along these same lines, he claims:

> Prominent among values that constitute the higher, reflective level is that of understanding. But this does not preclude a correlative level of knowledge allied to such understanding. It is in part because one understands how one knows that one's knowing reaches the higher level. A belief constitutive of such reflective knowledge is a higher epistemic accomplishment if it coheres properly with the believer's understanding of why it is true (and, for that matter, safe), of how the way in which it is sustained is reliably truth-conducive.
> That a belief cohere thus within the believer's perspective is, moreover, not irrelevant to that belief's being deeply attributable to the believer's epistemic agency (2001, 195).

In this passage, Sosa tackls the problem that has occupied me in this paper: epistemic ownership. In his view, the reflective stance provides the agent with a perspective about how it is that she knows, in terms of how reliable her beliefs are, which would be a particular form of understanding, one that would have some relation — it is "not irrelevant" — with the belief's attributability. All of this coheres with my proposal. After all, $p$ makes much more sense if it is derived from highly reliable sources than if it is not[26]. However, my key point is that one does not aim at understanding by just aiming at reliability. Aiming to regulate one's cognition in light of the reliability of one's sources must be seen as part of the effort to understand, which constitutes the agent's autonomous involvement in the process[27]. Concern for reliability is just one facet of the broader endeavor to understand, and it cannot replace the whole of it[28].

[26] When asked why it was so difficult for him to acknowledge his mistakes, the television character Frasier humorously said: "I have a degree from Harvard. Whenever I'm wrong, the world makes a little less sense".

[27] I assume that, in certain instances, this form of source-related understanding may suffice to partially address and overcome epistemic wantonness (something akin to holding that some proposition makes sense, given that it is asserted by a speaker that is epistemically trustworthy).

[28] This point sheds light on why some scholars argue that Sosa did not correctly identify the goal of the reflective perspective as knowledge. According to Grimm (2001, 186), for instance, its true goal may be understanding. However, I see no reason why one same cognitive performance cannot aim at *both* knowledge *and* understanding, just as an action may aim at both a primary goal—the satisfaction of those motivations that prompt it — and a constitutive goal — making sense. Propperly understood, these two goals are not in competition.

The difference is subtle but important. The recognition of understanding as a constitutive goal of epistemic rationality provides with a systematic solution to the problem of epistemic ownership that a merely regulative theory would lack. Recall that beneath the symptom of the vicious regress, the true endemic disease of higher order theories was that ascending to a higher order of assessment does not provide in and by itself what the agent lacked at the lower order. However, when the demand is not merely regulative but constitutive, the added value does not stem solely from the regulation by higher orders of assessment or from the increased reliability resulting from it. The enhanced epistemic value in terms of ownership arises from the fact that ascending the orders of assessment demonstrates the agent's dedication to a goal that defines the very essence of her intellectual pursuits—namely, making sense of the world[29]. Aiming for understanding beyond knowing is what inherently involves the agent, as it is not an optional regulative choice but a motivation that defines her role as an autonomous epistemic rational agent.

Let us now reconsider Bill's situation within the context of this proposal. Recall that we began with Bill in a state of wantonness, and our initial advice was to regulate his belief-formation processes by paying attention to reliability. That wouldn't compel him though to engage in genuine autonomous epistemic deliberation. According to a constitutivist approach, we should rather emphasize to Bill that he should aim at understanding whatever information he encounters, regardless of the order it belongs to. On occasion, we might encourage Bill to ascend to higher orders of assessment, perhaps if he had indications of the unreliability of his sources. After all, it does not make sense to believe proposition *p* based on evidence from a known unreliable or untrustworthy source. What sets the dispositionalist and regulative view apart from the constitutivist perspective is that the agent's inclination to revise lower-order beliefs isn't solely driven by the goal of safe belief. Instead, it is rooted in a genuine commitment to the constitutive objective of epistemic deliberation: comprehending every piece of knowledge in the context of the entirety of what one knows. Aiming at understanding isn't merely a commitment that originates at a lower level and potentially extends to higher orders of assessment. Rather, it serves as the guiding light that illuminates the entirety of our knowledge across all levels of assessment[30].

[29] I am inclined to think that this point equips us with the necessary tools to elucidate why intellectual autonomy holds epistemic value, as posed in Sosa (2021, Ch. 1) or Vega-Encabo (2021). However, this discussion must be reserved for another occasion.

[30] It is crucial to emphasize that my proposal does not advocate for the *substitution* of the specific and contingent goals of the agent with the constitutive goal. In the practical realm, the pursuit of intelligibility does not replace but rather complements the goal of satisfying one's desires through intentional action. Likewise, in the epistemic context, aiming for understanding does not substitute but rather complements the goal of 'satisfying' one's beliefs through the acquisition of knowledge. Otherwise, as Sosa himself acknowledges, "Beliefs could develop for years through the rare and imaginatively coherent thinking of an obsessive paranoid." (2021, 214). It is conceivable then for an agent to aspire to understanding while neglecting the pursuit of knowledge. Such an agent would undoubtedly fall short of the standards set for an ideal epistemic agent, but the nature of this shortfall differs from that of the epistemic wanton.

## 7 Some objections

One possible objection to my view is that full understanding is often an unattainable goal. Even when we believe we have it, that intuition may be fallible, leaving us uncertain about whether we've truly comprehended the subject. Are we not setting an overly elusive and unreachable objective for the wanton?

In response to this objection, it is important to note that having a constitutive goal for an activity does not necessarily mean that the goal must always be fully achieved to engage in that activity properly. What is required for this is that one strives toward that goal to the best of one's abilities. This sets my proposed view apart from deontological constitutivist perspectives. Simply teaching Bill a rule wouldn't be enough. What he ought to learn is that he also has to *aim* at something that constitutes his intervention in the whole process as autonomous agent.

To further illustrate this point, consider the analogy of teaching chess to someone. We can explain all the rules of the game, from how each piece moves on the board to various strategic and tactical guidelines. However, we also need to convey that the ultimate goal in chess is to checkmate the opponent, even though technically, not attempting to do so would not violate any rules. This goal defines the essence of the game, making it a goal-oriented activity, not just a rule-based one. Even if for our pupil it may be nearly impossible to achieve checkmate, she would not truly be playing chess if she were not aiming for it. Similarly, in processes of inquiry and epistemic deliberation, aiming for understanding is crucial, even if full understanding is challenging or unattainable. The act of aiming at understanding is what allows one to truly own one's thoughts and cognitive processes.

Here is another objection: the constitutive goal I propose for epistemic deliberation is not inherently epistemic but leans more towards the ethics of belief, or intellectual ethics, that is not concerned with the kind of assessment that is internal to epistemic endeavours (Sosa, 2021, 17)[31]. Similarly, Stephen R. Grimm, even if he has been critical with respect to Sosa's higher-order reflective stance because, just as I've been showing, it does not manage to demonstrate by itself why our beliefs are "fully our own", holds that what may be earned with such ownership is "better thought of as a moral or ethical gain, rather than an epistemic one" (2016, 194)[32]. What is at stake at this point goes beyond the scope of this paper, but I would say that understanding is still a genuinely epistemological goal under the pluralist assumption that understanding is an essential target of our cognitive engagement with the world. As epistemic agents, we do not only aim at knowledge and the verification of each proposition's truth. We also strive to comprehend how all elements interconnect, exploring the coherence of information — an objective that may also fall under telic

[31] See above, notes 8 and 19. Relatedly, Chrisman (2016) identifies regulative goals with the manifestation of skills (in contrast with constitutive aims, which would be related to the decision to participate in the activity in question). Given that Sosa identifies epistemic virtues with skills, this would explain why, in Sosa's views, constitutive *aims* would be left out of the epistemological picture.

[32] Zagzebski (1998, 259) famously defends this. A similar concern for the genuinely epistemic gain of intellectual autonomy has been recently pressed by Vega-Encabo (2021), who eventually claims that such gain is more related to an aspiration to perfect our own agency than to the obtention of any specific epistemic good.

normativity. The realm of intellectual ethics concerns itself with why and how the agent *values* this constitutive goal, an inquiry that invites an assessment within the broader context of the agent's overall motivation, encompassing prudential, moral, and even political substantial aspects. However, the emphasis of my argument is not tied to the agent's personal particular motivation. On the contrary, my teleological epistemic constitutivism underscores the agent's endeavor to attain a specific goal intrinsic to her epistemic pursuits, irrespective of the reasons underlying the goal's significance to her — which is the subject of intellectual ethics.

Relatedly, one may raise the objection that, the way I defined them, the constitutive goals of practical autonomy (intelligibility) and epistemic autonomy (understanding) are essentially identical—both centered on the idea of making sense. This would imply a potential collapse of practical reasoning into theoretical reasoning, or vice versa, raising doubts about the stability of the original parallelism that underpinned my argument[33]. Notably, scholars like Velleman have attempted to bridge this divide by interpreting practical reasoning in terms of theoretical reasoning. In this context, as Mitova suggests:

> Velleman understands practical deliberation as a species of theoretical deliberation. The conclusion of an episode of practical deliberation is an intention to φ which, according to him, just is a belief that I will φ. This automatically turns the norms governing practical sense-making into epistemic norms (Mitova, 2016, 213).

But I do not seek, nor do I need to pursue that reduction, in any of its two possible directions[34]. While intelligibility and understanding have an internal relation, as they both involve the pursuit of sense, they remain distinctly separate, similar to how the practical and the epistemic differ concerning truth — they instantiate different directions of fit. As autonomous practical agents, we engage in actions that create sense. When we succeed, the facts that are our deeds actively produce meaning. Conversely, as autonomous epistemic agents, our goal is to grasp the sense already inherent in facts. We do not generate this meaning; rather, we seek to discover it. In both cases, the motivation is rooted in the pursuit of sense, but as practical agents, we aim to produce it through intelligible actions, while as epistemic agents, we aim to grasp it through understanding. In this way, the goal of intelligibility can be constitutive of autonomous practical rationality without collapsing into the epistemic goal of understanding. We succeed as autonomous practical agents by doing something in the world that makes sense; we succeed as autonomous epistemic agents when we manage to find the sense that something in the world makes.

One last thought about the possibility of failure. Agents may satisfy their primary goals while still falling short in the quest for their constitutive goals, both in action

[33] Miracchi and Carter (2022), for instance, propose "to abandon from the very start the idea that knowledge and action (and their corresponding attempts) are 'mirrors' of each other — mirrors reversing direction of fit".

[34] I thus part ways here with Velleman (1989), as I intend to maintain neutrality on the potential intellectualist reduction of the intention to φ to the belief that one will φ.

and in knowledge. We may successfully satisfy our desires but find no meaning in our actions, or we may know the way things are without fully comprehending the sense they make. This is possible because the respective primary goals (satisfying desires or knowing truths) differ from the constitutive goals (doing something intelligible, or understanding). My proposal allows for the possibility of such failures, which is a strength because it reflects common situations. The constitutive goals of rational autonomy may indeed be elusive, and this gives rise to a form of skepticism that merits discussion. My aim has not been to dismiss this particular variety of skepticism but rather to provide a framework for its comprehension and analysis.

## 8 Conclusions

My working hypothesis has been that we can gain insights into epistemology by drawing on the literature on practical ownership. This parallelism is based on Williamson's idea that action corresponds to knowledge as desire corresponds to belief. From this starting point, I have identified certain analogies between various concepts in action theory and epistemology.

First, I equated Davidson's causalism in action theory with Goldman's process reliabilism in epistemology, in that both emphasize reliable processes for action and belief formation. Next, I likened cases of deviant causal chains that fall short of intentional action to Gettier cases that fall short of knowledge. Then I drew parallels between Frankfurt's addicts in action theory and a specific variety of epistemic akratic agents in epistemology, who know but are reluctant to accept the reasons why they know. As a different pathology, I identified the equivalent of practical wantons: epistemic wantons, as agents that may satisfy their beliefs in reliable ways, and thus know, but lack ownership of this knowledge, remaining as mere spectators of their own rational deliberation.

To address this lack of ownership, I distinguished between regulative and constitutive strategies. Frankfurt and Sosa represent regulative strategies, as they respectively propose higher-order theories of practical and epistemic rationality. However, I argued that these strategies face a common challenge: they only delay the issue of ownership rather than fully resolving it.

In contrast to those regulative strategies, I have advocated for a teleological constitutivist solution to address epistemic ownership, drawing an analogy to Velleman's stance in the practical realm. Agents attain ownership of what they do or know by aiming at sense-making. In practical constitutivism, this goal is doing something intelligible, as outlined by Velleman. In the form of epistemic constitutivism that I have presented, the goal is understanding what one knows. Merely aiming for the satisfaction of desires or beliefs — even if intentional action or knowledge is attained — is insufficient to overcome practical or epistemic wantonness. Instead, agents must pursue a constitutive goal of autonomous rational deliberation, which, in the epistemic case, is understanding what one knows.

**Acknowledgements** I began drafting this paper while serving as a visiting scholar at NYU in 2019, which afforded me the opportunity to discuss different versions of the draft with J. David Velleman, Ernest Sosa,

John Greco, and Santiago Echeverri. The paper languished for a period thereafter, only attaining its final form during a stay at the University of Glasgow in 2023, where it greatly benefited from my extensive discussions with members of the COGITO Epistemology Research Center, particularly with Tim Kearl and Mona Simion. I would also like to express my gratitude to Teresa Bejarano, Antonio Guillén, Modesto Gómez-Alonso, Jesús Vega, and my Ph.D. students at Sevilla (Daniel Barbarrusa, Nacho Gómez-Ledo, Dani Pino, and Lola M. Vizuete) for their invaluable ideas and criticism over those years. Additionally, I extend my thanks to the audiences at the Autonoma University of Madrid, Universitat de València, and University of Glasgow for their very insightful uptake. In terms of funding, I would like to acknowledge the Spanish Ministerio de Educación, Cultura y Deportes for supporting my stay at NYU through the Programa Estatal de Promoción del Talento y su Empleabilidad en I+D+I (Subprograma Estatal de Movilidad); to the Spanish Ministerio de Universidades for funding my stay at Glasgow, which was made possible by the European Union– NextGenerationEU initiative; and two distinct research projects: Nanorin (PID2021-123938NB-I00), and Metaprodes (PID2021-124152NB-I00), both funded by MCIN/ AEI / https://doi.org/10.13039/501100011033 / FEDER, UE.

**Funding** Funding for open access publishing: Universidad de Sevilla/CBUA

## Declarations

**Conflict of interest** The author has no competing interests to declare that are relevant to the content of this article.



## References

Aguilar, J. H., & Buckareff, A. A. (Eds.). (2010). *Causing human actions: New perspectives on the causal theory of action*. MIT Press.

Blake-Turner, C. (2021). Acting and believing on the basis of reasons. *Philosophy Compass*, *e12797*.

Broncano-Rodríguez, F., & Vega-Encabo, J. (2011). *Engaged Epistemic Agents Crítica*, 43(128), 55–79.

Chrisman, M. (2016). The Aim of Belief and the goal of Truth: Reflections on Rosenberg. In M. Grajner, & P. Schmechtig (Eds.), *Epistemic reasons, norms and goals* (pp. 357–381). De Gruyter.

Davidson, D. (1980). *Essays on actions and events*. Clarendon.

Frankfurt, H. G. (1988). *The importance of what we care about: Philosophical essays*. Cambridge University Press.

Goldman, A. I. (1967). A causal theory of knowing. *The Journal of Philosophy*, *64*(12), 355–372.

Goldman, A. I. (1976). Discrimination and perceptual knowledge. *The Journal of Philosophy*, *73*(20), 771.

Goldman, A. I. (1979). What is justified belief. In G. S. Pappas (Ed.), *Justification and knowledge*. Riedel.

Goldman, A. I. (1986). *Epistemology and Cognition*. Harvard University Press.

Grajner, M., & Schmechtig, P. (2016). *Epistemic reasons, norms and goals*. De Gruyter.

Greco, J. (2014). Episteme: Knowledge and understanding. In K. Timpe, & C. A. Boyd (Eds.), *Virtues and their vices* (pp. 285–302). Oxford University Press.

Grimm, S. R. (2001). Ernest Sosa, Knowledge, and understanding. *Philosophical Studies*, *106*(3), 171–191.

Grimm, S. R. (2006). Is understanding a species of knowledge? *The British Journal for the Philosophy of Science*, *57*(3), 515–535.

Grimm, S. R. (2016). The Value of Reflection. In M. Á. Fernández Vargas (Ed.), *Performance Epistemology: Foundations and Applications*. Oxford University Press, 183–195.
Horowitz, S. (2014). *Epistemic Akrasia Noûs*, 48(4), 718–744.
Kelp, C. (2021). *Inquiry, Knowledge, and understanding*. Oxford University Press.
Korsgaard, C. M. (2008). *The Constitution of Agency: Essays on practical reason and Moral psychology*. Oxford University Press.
Lasonen-Aarnio, M. (2014). Higher Order evidence and the limits of defeat. *Philosophy and Phenomenological Research*, *88*(2), 314–345.
Lynch, M. P. (2016). *The internet of us: Knowing more and understanding less in the age of Big Data*. Liveright Publishing Corporation.
McHugh, C. (2011). What do we Aim at when we believe? *Dialectica*, *65*(3), 369–392.
McHugh, C. (2013). Epistemic responsibility and Doxastic Agency. *Philosophical Issues*, *23*(1), 132–157.
McHugh, C. (2014). Exercising Doxastic Freedom. *Philosophy and Phenomenological Research*, *88*(1), 1–37.
Miracchi, L., & Carter, J. A. (2022). Refitting the mirrors: On structural analogies in epistemology and action theory. *Synthese*, *200*(1), 9.
Mitova, V. (2016). What do I care about epistemic norms? In M. Grajner, & P. Schmechtig (Eds.), *Epistemic reasons, norms and goals* (pp. 199–223). De Gruyter.
Navarro, J. (2015). No achievement beyond intention: A new defence of robust virtue epistemology. *Synthese*, *192*(10), 3339–3369.
Navarro, J. (2016). Acting in Order to know, knowing in Order to Act: Sosa on Epistemic and practical deliberation. *Disputatio International Journal of Philosophy*, *VIII*(42), 233–252.
Pritchard, D. (2014). Knowledge and understanding. In A. Fairweather (Ed.), *Virtue Scientia: Bridges between Virtue Epistemology and Philosophy of Science* (pp. 315–328). Springer.
Rinard, S. (2019). Equal treatment for belief. *Philosophical Studies*, *176*(7), 1923–1950.
Scanlon, T. (1998). *What we owe to each other*. Harvard University Press.
Shah, N. (2003). How Truth governs belief. *Philosophical Review*, *112*(4), 447–482.
Shah, N., & Velleman, J. D. (2005). Doxastic Deliberation. *Philosophical Review*, *114*(4), 497–534.
Singh, K. (2019). Acting and believing under the guise of normative reasons. *Philosophy and Phenomenological Research*, *99*(2), 409–430.
Sosa, E. (2001). Human Knowledge, Animal and Reflective. *Philosophical Studies, 106*, 193–196.
Sosa, E. (2007). *A Virtue Epistemology: Apt belief and reflective knowledge 1*. Clarendon.
Sosa, E. (2011). *Knowing full well*. Princeton University Press.
Sosa, E. (2015). *Judgment and Agency*. Oxford University Press.
Sosa, E. (2021). *Epistemic explanations: A theory of telic normativity, and what it explains*. Oxford University Press.
Tanney, J. (2005). Reason-explanation and the contents of the mind. *Ratio*, *18*(3), 338–351.
Vega-Encabo, J. (2021). Understanding and the Value of Intellectual Autonomy. In J. Matheson, & K. Lougheed (Eds.), *Epistemic autonomy* (pp. 195–214). Routledge.
Velleman, J. D. (1989). *Practical reflection*. Princeton University Press.
Velleman, J. D. (2000). *The possibility of practical reason*. Clarendon.
Velleman, J. D. (2009). *How we get along*. Cambridge University Press.
Vitz, R. (2021). Doxastic Voluntarism. In *Internet Encyclopaedia of Philosophy*. 12/28/2021, from https://iep.utm.edu/doxa-vol/.
Wedgwood, R. (2002). The Aim of Belief. *Noûs*, *36*(s16), 267–297.
Williams, B. (1973). *Problems of the self. Philosophical papers 1956–1972*. Cambridge University Press.
Williamson, T. (2002). *Knowledge and its limits*. Oxford University Press.
Williamson, T. (2017). Acting on Knowledge. In J. Adam Carter editor, Emma C. Gordon editor, & Benjamin W. Jarvis editor (Eds.), *Knowledge first: Approaches in epistemology and mind*. University Press.
Zagzebski, L. T. (1998). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge Univ. Press.
Zagzebski, L. T. (2001). Recovering understanding. *Knowledge, Truth, and Duty*. Oxford University Press.