THOUGHT
A JOURNAL OF PHILOSOPHY

Online First: September 24, 2024
DOI: 10.5840/tht202492035

Recruitment Revisited: Cognitive Extension and the Promise of Predictive Processing

Luke Kersten
University of Alberta

Abstract: The extended mind thesis maintains that cognitive processes and systems can, on occasion, stretch to include parts of the brain, body, and world. One outstanding puzzle facing this view is the "recruitment puzzle." The recruitment puzzle asks how cognisers are able to reliably recruit internal and external resources such that they form extended systems. Andy Clark has recently suggested that predictive processing helps to address this puzzle. I argue that, while promising, Clark's proposal remains incomplete. I suggest that Clark's proposal can be productively extended by disambiguating two important senses of recruitment: *ready-to-hand* and *adaptive* recruitment. After outlining the recruitment puzzle and Clark's proposal, I suggest that careful attention to these two senses of recruitment helps to reveal further constructive ways of developing the extended mind thesis.

Key words: extended mind, predictive processing, recruitment puzzle, prediction error minimisation; epistemic action

INTRODUCTION

The extended mind thesis (or hypothesis of extended cognition) claims that cognitive processes and systems can, on occasion, stretch to include parts of the brain, body, and world. It holds that the physical limits of the individual do not mark an important boundary when it comes to understanding cognition (Clark and Chalmers 1998; Wilson 2004; Clark 2008). While initial discussions of the thesis centred on thought experiments and metaphysical possibility, more recent work has engaged with the experimental and methodological practices of cognitive science (Favela et al. 2021). Yet despite its notable history and established status, a number of outstanding issues remain for the thesis.

One such issue is the 'recruitment puzzle' (Clark 2007, 2008, 2017, 2024). The recruitment puzzle concerns how and why cognisers are able to recruit internal and external resources into larger extended systems. It focuses on an apparent explanatory gap between our propensity to incorporate external props and tools into our problem-solving routines and the process or mechanism by which these activities occur. It is a 'puzzle' in the sense that there is no clear answer as to how or why such a recruitment process should unfold as it does in the case of extended cognitive systems.

Clark (2024) has recently sought to address the recruitment puzzle by drawing on the resources of 'predictive processing' (or PP for short). The suggestion is that the predictive brain's fundamental drive to reduce *expected future prediction error* explains both why and how external and internal resources come to be incorporated into extended systems. In what follows, I argue

Luke Kersten, Department of Philosophy, University of Alberta, 2-45 Assiniboia Hall, Edmonton AB T6G 1C9; https://orcid.org/0000-0001-7054-8942; kersten@ualberta.ca.

that discussion of the recruitment puzzle remains incomplete. I suggest that Clark's proposal, while promising, neglects an important distinction between two senses of recruitment: *ready-to-hand* and *adaptive* recruitment. I suggest that careful attention to these two senses of recruitment helps to reveal further constructive ways of developing the extended mind thesis, focusing attention on the different factors involved in extended investigations, such as conscious deliberation and long-term adaptive integration.

1. THE RECRUITMENT PUZZLE

The recruitment puzzle has its origins in what appears to be an initial gap within formulations of the extended mind thesis.[1] Clark, for instance, frames the issue as follows:

> [T]his overall vision (of cognition distributed among brain, body, and world) bequeaths a brand new set of puzzles. It invokes an ill-understood process of "recruitment" that soft-assembles a problem-solving whole from a candidate pool that may include neural storage and processing routines, perceptual and motoric routines, external storage and operations, and a variety of self-stimulating cycles involving self-produced material scaffolding. (Clark 2008: 138)

For Clark, it is unclear how the basic extended mind story should or can explain an agent's ability to actively exploit whatever mix of internal and external resources best support problem-solving with minimum effort. The recruitment puzzle emerges because it is unclear from what process or mechanism extended systems arise.

For instance, in the now classic example, Otto (an individual with a mild memory impairment) uses his trusty notebook (an external resource), alongside various brain-bound abilities (e.g., perception), to find his way to his favourite museum, the Museum of Modern Art (MoMA). The original idea was that the notebook (and its contents) played a sufficiently fine-grained and integrated functional role within Otto's mental life such that it should be included within his cognitive machinery (Clark and Chalmers 1998). Notice, though, that while the notebook forms a key part of Otto's extended system, it is unclear, at least initially, how the external resource is selected and coordinated with internal ones in the formation of an extended cognitive system. While Clark and Chalmers (1998), along with subsequent treatments, offer various conditions for cognitive extension, little is said about the recruitment machinery responsible for creating extended processing arrays.

One reason the recruitment puzzle should prove interesting to cognitive scientists and philosophers alike is that it bears directly on the scope of the extended mind thesis. If, for instance, the thesis is interpreted as a narrow metaphysical claim about the nature of mind, then there is little need to account for the formation of extended systems; there is indeed no puzzle to solve. But, if, as many maintain, the extended mind is a more expansive vision, one which also covers developmental and ontogenic questions, then the recruitment puzzle points to the need for a fuller, more complete story about the genesis and formation of extended systems and processes (see, e.g., Wilson 2004; Clark 2008; Wilson and Clark 2009). In pressing proponents to identify the underlying mechanisms and processes by which extended systems are formed and sustained, the recruitment puzzle speaks to the viability of the extended mind thesis as a full-scale vision of cognition. Resolution of the recruitment puzzle would speak directly to the thesis' plausibility, clarifying a further important piece of extended mind story.

This plausibility should also be of more general interest to philosophers given the thesis' application in a variety of recent philosophical debates, such as those surrounding consciousness (Ward 2012), affectivity (Colombetti and Roberts 2015), music cognition (Kersten 2017), epistemology (Carter et al. 2018), mechanistic explanations (Krickel 2020), and non-human animals

(Facchin and Leonetti 2024). Moreover, having been applied to such diverse fields as education (Pritchard et al. 2021), mental health (Roberts et al. 2019), human-computer interactions (Smart 2021), mathematics (Vold and Schlimm 2020), and neurocomputational theories (Kirchhoff and Kiverstein 2021), the thesis has also spread beyond strict philosophical circles. Such novel developments highlight the continued relevance of the thesis.

2. THE PROMISE OF PREDICTIVE PROCESSING

As mentioned, Clark (2024) has recently suggested that PP finally puts us in a position to address the recruitment puzzle.

Generally speaking, the core idea of PP is that all cognitive processes emerge from the brain's tireless efforts to minimise 'prediction error,' which is the mismatch between externally generated sensory signals and internally generated predictions about those signals issuing from a generative model (Clark 2013, 2016; Hohwy 2013). A generative model is an internal model which attempts to track the casual-probabilistic structure of the environment. It is hierarchically organised. At each level, evidence from a lower level is anticipated by predictions at a higher level, e.g., predictions at one level L + 1 provide a prediction of the evidence coming from level L. This recurrent process of 'prediction error minimisation' (i.e., the process of eliminating the discrepancy between what is predicted about the world and the evidence delivered from the senses) is what enables the brain to recover the structure of the environment from otherwise noisy and ambiguous signals (Williams 2019). The brain uses two strategies to achieve this. One is by updating the generative model, which brings the prior predictions in line with the world (perceptual inference); the other is by selectively sampling new sensory inputs via action, which brings the world in line with the brain's prior predictions (active inference).

Clark's interesting suggestion is that the coordination and recruitment of internal and external resources can be explained in light of the brain's more fundamental drive to reduce prediction error over time.[2] The proposal has two key elements. The first is what Clark (2024) calls "expected future prediction error." This is the idea that in virtue of being a prediction machine (i.e., a system which generates and tests predictions about the world) the brain selects those actions which best minimise counterfactual prediction error relative to its goals. That is, to minimise expected future prediction error, the brain looks ahead and compares what would happen if it adopted one action policy rather than another, e.g., testing the depth of a river using a stick versus one's foot. It achieves this by assigning, and continually updating, a weighting of prediction error signals. Error signals which have a high precision weighting are prioritised during a task, while signals which have a low precision weighting are given a lower priority (or even partly discounted). Different action policies are selected so as to reliably reduce prediction error in a given context via patterns of behaviour. I will say more on this later.[3] The second is the concept of "epistemic action." Epistemic actions are those behaviours which aim to improve or transform information for an agent in order to achieve some goal—for example, searching online for a film or switching a light bulb on and off to see if it needs replacing. In performing epistemic actions, agents gather information to reduce the gap between their current state and their goal state. Of note here is that while epistemic and practical actions are often intertwined, epistemic actions do not necessarily contribute to practical ends.

Putting these two elements together, the suggestion is that if an agent can gather information about the world (via epistemic actions) and it can select actions which minimise expected future prediction error (via precision weighting), then it will tend to discover effective routes to achieving its goals. Extended systems emerge when agents minimise expected future prediction error via *external resources* supported by epistemic actions. Or, as Clark writes, "[s]pecific extended processing arrays are recruited and deployed in temporally correct fashion without the need for deliberative agentive oversight, so as to best minimize expected future prediction error" (2024).

For example, returning to the MoMA case, in relying on an external resource to provide information about spatial location, Otto is using his notebook as a means of minimising future prediction error (failing to arrive at the museum). He performs an epistemic action on the notebook (looking up the address) to bring about his preferred goal state (arriving at the museum) via a course of action (walking). As the error minimising activities of the agent unfold, the epistemic possibilities provided by certain external resources, such as Otto's notebook, become woven into wider cognitive economies.

To be clear, though, the grounds for extension are still the 'trust' and 'glue' conditions articulated by Clark and Chalmers (1998). Clark (2024) is clear that it is only when external resources are sufficiently ingrained and well-practiced via sensorimotor loops (trust) as well as are robustly available (glue) that they qualify as part of an extended system. The point, instead, is that precision estimation (i.e., the context-specific attentional weighting of prediction error) is what enables various flows of information (via epistemic actions) to be included within an extended system on the basis of the reliability/confidence they engender. Clark's proposal is articulating the sub-personal, brain-based mechanism and processes responsible for external resource recruitment within extended systems.

A potential wrinkle in the story is how agents select amongst different actions. If an agent is faced with multiple or competing goals (e.g., seeing a film or completing a chapter of their new book), there is a question as to how they choose which better minimises long term prediction error. The answer for Clark lies again in precision-weighting. Precision-weighting, recall, involves a set of parameters that determine responsiveness to different prediction errors; 'precision' in this context quantifies the inverse variability or uncertainty associated with a probability distribution (Parr et al. 2019). Because brains employ a generative model (a mirroring of the causal-probabilistic structure of the external environment), agents possess prior beliefs (associated expectations) about the confidence of different action policies—new experiences encode different preferred outcomes (goals) into an agent's generative model, which ensures that multiple outcomes can be secured over time. By adjusting precision-weightings, an agent can flexibly adapt to the changing needs and circumstances of the environment. This enables them to prioritise one goal (predicted future state) over another, e.g., seeing a film instead of finishing a new chapter.

3. RECRUITMENT REVISITED

There is much to be said for Clark's proposal. It offers a novel solution to an important outstanding puzzle facing the extended mind. However, a good deal of the analysis hinges on the notion of 'recruitment.' One interesting point here is that Clark thinks recruitment should not be thought of as an effortful or deliberate process. He writes, for instance:

> What must be avoided, if we are to preserve space for this view [extended cognition], is the idea of recruitment as itself effortful and deliberative. Such deliberate effortful gathering would work against the idea that the extended regimes arise—like their internal biological counterparts—fluently and rapidly as task and context shift and alter. (2024)

While extended systems may serve agentive goals, they arise and dissolve fluidly and automatically. Clark again calls on the MoMA case to illustrate. He notes, for instance, that while Otto's top-level desire to go to MoMA seeds a rich processing cascade (e.g., checking his notebook), one capable of fluidly selecting whatever combination of inner and outer operations best enables progress on the desired outcome, the cascade itself unfolds without need for further conscious reflection or deliberation. As Clark (2024) makes the point: "The emergence of these larger arrays is not itself typically a matter of agentive deliberation. An extended mind serves agentive goals (Otto wants to get to MoMA after all). But extended coalitions arise and dissolve rather fluidly and automatically, as do webs of activity arising purely within the biological brain." For extended systems to emerge, top-level agentive goals must overarchingly act as catalysts for larger organisational wholes.[4]

One reason Clark and others are likely keen to restrict the notion of recruitment in this way is to preserve a role for "phenomenal transparency," the idea that once a bio-external resource has become sufficiently incorporated into a cognitive system it becomes hidden or phenomenally inconspicuous. For Clark and others, phenomenal transparency is often seen as a reliable guide to cognitive extension.[5]

What is interesting, though, is that deliberate, effortful recruitment does seem to play a role in many paradigmatic cases of extended cognition, not only in the sense of seeding a rich processing cascade but also in the sense of creating and sustaining extended processing arrays over time. For example, to solve complex mathematical problems, such as multiplying, people are often taught to break down problems into smaller steps, such as using partial products; each step consisting of easier to compute operations (Wilson 2004). In doing so, individuals are said to extend their cognitive workload onto the world; the numerals on the paper acting as a temporary memory buffer, for example. But notice that these processes require conscious and deliberate effort on the part of the agent, at least initially. The budding mathematician needs detailed instruction and training in how to transform the complex problem into a simpler one via world-involving activities.

A similar point holds in the case of musical performance. As Kersten and Wilson (2016) point out, the developmental trajectory of musical performance often runs through representational systems, such as standard Western notation systems. These systems supply crucial information to performers, allowing them to off-load the representational and cognitive workload. But, again, the learning process is an effortful and deliberate one, particularly in the early stages of training. Music performers require explicit, on-line coordination between their actions and the representational system, e.g., via sheet music. This does not imply, of course, that performers cannot eventually fluidly and automatically exploit such representational systems; they often do. But when the initial developmental trajectory of musical performance runs through the use of notation systems, conscious and deliberate processes are involved not only in setting the agentive goals but also in the ongoing selection and manipulation of external resources.

These are only two quick examples, but what they seem to suggest is that agents, at least in some typical cases, recruit external resources into larger processing wholes over time based on deliberate, conscious activities. But notice that this point does not sit particularly well with Clark's proposal. The suggestion was that there was an important difference between automatic, fluid recruitment, and the more deliberative, effortful activities directed by an agent with respect to extended processing regimes. And yet, it seems that many bona fide cases of cognitive extension revolve around deliberate, effortful recruitment activities on the part of the agent to some extent.

So, what is going on? Part of the answer, I submit, is that there are, in fact, two distinct senses of 'recruitment' that can be teased apart. The first is what might be called a *ready-to-hand* sense of recruitment. On this reading, agents recruit external resources to perform specific cognitive acts or tasks at a given point in time, such as Otto's using his notebook to navigate to MoMA. This is the sense of recruitment Clark invokes when speaking of fast and frugal cognising. It is the task of "selecting and orchestrating the right set of bio-internal and bio-external resources at the right time" (Clark 2024).

The second sense, however, is what might be called an *adaptive* sense of recruitment. On this reading, agents recruit external resources not for their immediate, real-time use but, rather, for downstream incorporation into larger extended processing wholes. Agents learn to use external resources in order to co-opt them for later use. The resource gets modified (consciously, deliberately) so that it fits the agent's cognitive needs. For example, if Otto's notebook begins to show signs of decay, he might pick up a new one to reduce the chance of getting lost next time he travels. While the new notebook is not yet part of an extended system, as it fails to meet the additional requirements on extension (i.e., the trust and glue conditions), it nonetheless stands poised to be included given sufficient time and integration. Otto has to fill in, consciously and deliberately, the relevant information he wishes to remember. This is the sense of recruitment Clark invokes when thinking about long term prediction error reduction.[6] He writes, for instance: "As this [recruitment] process unfolds, the epistemic possibilities provided by certain external items and resources, such as the smartphone, may gradually become woven into our daily lives in the additional ways required to count as true cognitive extensions" (2024).[7]

These two senses of recruitment are related in purpose (minimising expected future prediction error) but distinct in origin and design. For while ready-to-hand recruitment functions to resolve uncertainty in the here-and-now via perception-action loops, adaptive recruitment functions to reduce uncertainty in the long run by budgeting internal and external resources for later use. In the former's case, agents "consume" an external resource in the here-and-now in the service of some particular task or goal. In the latter's case, however, agents "train" themselves to consume a resource appropriately at a time further down the road. Even though both senses of recruitment are driven by a fundamental need to minimise prediction error, they are sustained by different processes, i.e., fast, automatic processing and deliberative, agentive oversight.[8]

To be clear, though, in drawing attention to the two senses of recruitment, I am not casting doubt on Clark's wider story. Rather, the suggestion is that the story is incomplete as presented. In neglecting to disambiguate what turns out to be a multifaceted concept (recruitment), Clark's proposal runs the risk of obscuring important differences in the way the recruitment of external resources unfolds. For this reason, in what follows, I want to complement and extend the analysis by showing how attention to these two senses of recruitment can provide further constructive ways of developing the extended mind thesis.

4. EXTENDED RECRUITMENT

There are few reasons to think pulling apart the ready-to-hand and adaptive senses of recruitment is a good idea in the context of extended cognition.

First, the distinction helps to direct attention to the different factors involved in extended investigations. One reason this is important is that the two senses of recruitment appear to place different epistemic demands on investigation. On the one hand, ready-to-hand recruitment requires understanding the on-the-spot assembly of internal and external resources that functions to solve a particular task or problem. This may involve looking at the various ways in which agents make trade-offs in decision-making about internal and external resources, such as those about the certainty of outcomes (Trautman and van de Kuilen 2015), expected value of rewards (Mehlhorn et al. 2015), and cost of information gathering (Juni et al. 2016). On the other hand, adaptive recruitment requires understanding the developmental trajectories of cognitive systems across time—as the previous examples highlight, external resources, such as crib sheets or notebooks, are poised not only for momentary use but also downstream incorporation. Different task horizons (how many future choices the decision-maker believes are before them), entail different levels of uncertainty, requiring agents to make different decisions about which resources to recruit on longer time-scales (Rich and Gureckis 2018). This may involve studying the ways in which agents actively alter their epistemic landscapes, for instance. A good example is spicing traditions, in which complex causal relationships are encoded into adaptive food practices. Spicing traditions allow individuals to model hidden causes of the environment which would otherwise remain obscure, such as how to kill foodborne pathogens in meat. By actively altering their epistemic environment over time, agents minimise long-term prediction error, and so achieve behavioural success (i.e., not getting sick).

In each case, understanding how and why an extended system is formed requires sensitivity to different factors. To effectively spell out the form and nature of an extension in its rich complexity, one needs to pay attention to not only the functional profile of on-the-spot neural-body-environment assemblies, but also how such assemblies come to be shaped over time by interactions with other resources and assemblies. In pulling apart the two senses of recruitment, we are better positioned to see how extended investigations can benefit from a simultaneous focus on local mechanistic questions and broader developmental ones.[9]

Second, the distinction helps to bring into focus the important role of conscious, deliberate thought in extended discussions. As we saw, conscious, deliberate recruitment can play a key role in the formation of extended systems. However, it is not entirely clear at present how deliberate, conscious reasoning fits within the PP framework. At first blush, conscious deliberation appears to involve the flow of logically connected ideas (or merely associated ones, or some combination of the two). It does not appear to involve the generation of prediction. There does not appear to be any mismatch with which an error signal might be generated, or prediction for conscious thought to be about. The structure of conscious deliberation does not appear to neatly fit within the PP framework.[10] Work on metacognition and problem solving may prove particularly relevant here, as self-explanation and error awareness, for example, are often found to be important metacognitive strategies in problem solving (Tajika et al. 2007; Erbas and Okur 2012). Thinking further about how such work fits within the wider subpersonal and personal story PP wants to tell may help to smooth out any apparent tension in the relationship. Attending to the different senses of recruitment may enable us to better understand how the fast, automatic processes involved in real-time precision estimation interact with the slower, more deliberate thought processes involved in epistemic activities spread over longer time scales. Each contribute to the creation of extended systems, but each does so in a different way. To be clear, I am not saying that PP cannot account for conscious deliberation. Rather, I am saying that in focusing on the ways in which extended cases are generated by conscious deliberation, the two senses of recruitment draw our attention to the need to spell out what the relationship between fast, automatic processes and slower, deliberate thought processes amounts to, and that holds promise for any future extended/PP story.

Finally, teasing apart the two senses helps to account for an important aspect of why cognitive extension is poised to occur at all. As Clark and others suggest, fluidity and automaticity are likely markers of bona fide extended systems. Yet, as we saw, in many cases, such as the budding mathematician or musician, there is something important to how external resources are recruited over longer time-scales; many resources are poised to acquire the requisite fluidity and "high bandwidth" with enough time. Ready-to-hand and adaptive recruitment preserve this key feature of extended cases. Part of the reason why, I want to suggest, is because they indirectly highlight the different processes underlying extended system formation, e.g., niche construction, scaffolding, etc. Ready-to-hand recruitment captures how well-understood and integrated resources are deployed in on-the-spot assemblies, while adaptive recruitment reflects the ways in which external props and resources gradually become woven into our cognitive economies such that they come to count as cognitive extensions. The ready-to-hand/adaptive distinction captures an important part of our thinking about extended cases.

CONCLUSION

So, to briefly sum up, I argued that Clark's (2024) solution to the recruitment puzzle, while instructive in several ways, comes up short. This is because it fails to distinguish between two important senses of recruitment: ready-to-hand and adaptive recruitment. I suggested in turn that further attention to these two senses of recruitment provides constructive ways of further developing the extended mind thesis, focusing attention on both the different factors involved in extended investigations and the role of conscious deliberation and long-term adaptive integration.

NOTES

1. Discussion of PP and extended cognition has sometimes been framed through the lens of 'Markov Blankets' (Hohwy 2016, 2018; Kirchhoff and Kiverstein 2019, 2021). I follow here Clark in avoiding such a framing. For a broader overview of the relationship between extended cognition and PP, see Kersten 2022.
2. For an earlier version of the account, see Clark 2017.
3. For a detailed discussion, see Smith and Sprevak 2023.
4. Elsewhere, Clark (2011: 450) also suggestively writes: "different neural sub-systems would have their own sub-personally mediated ways to 'call' or access the external resources, thus building the resource's reliable presence so deep into the information processing flow chart that it becomes visibly arbitrary and unhelpful to draw a single metabolically determined line dividing the truly 'cognitive.'"
5. For recent criticisms of the idea, see Facchin 2024 or Smart et al. 2022.
6. This sense of recruitment can also be found in the wider extended cognition literature. See Menary 2018 for discussion.
7. There is potentially a third sense of recruitment, in which the agent itself is modified in order to use a resource, what might be called "ontogenetic" recruitment. However, as this further sense does appear to directly bear on extended cases, I leave it to one side.
8. There is likely no sharp line here, as various engagements will recruit external resources along a continuum of more intentional and effortful less.
9. One line of response might be to suggest that adaptive recruitment is simply a way of "setting the scene" for downstream, on-the-spot cognitive extension (see, e.g., Clark 2011). In a sense, I think this is right. Adaptive recruitment, indeed, often serves this role. However, such a response does not address the specific point being made here: namely, that neither the subpersonal routines described by PP, nor the shared cultural or social practices described by Hutchins (2010), is sufficiently fine-grained enough so as to make sense of all bona fide extended cases. To explain the temporally extended dovetailing between the neural and extra-neural resources, we also have to consider the role played by consciously directed agentive level goals and reasoning in recruitment.
10. See Rupert 2022 for further discussion.

REFERENCES

Adams, F., and K. Aizawa. 2008. *The Bounds of Cognition*. Blackwell.

Bruineberg, J., K. Dolega, J. Dewhurst, and M. Baltieri. 2022. "The Emperor's New Markov Blankets," *Behavioral and Brain Sciences* 45: 1–63.

Carter, J. A., A. Clark, J. Kallestrup, S. O. Palermos, and D. Pritchard, eds. 2018. *Extended Epistemology*. Oxford University Press.

Clark, A. 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.

Clark, A. 2011. "Finding the Mind," *Philosophical Studies* 152: 447–461. https://doi.org/10.1007/s11098-010-9598-9

Clark, A. 2013. "Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science," *Behavioral and Brain Sciences* 36(3): 181–204.

Clark, A. 2016. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press.

Clark, A. 2017. "Busting Out: Predictive Brains, Embodied Minds, and the Puzzle of the Evidentiary Veil," *Noûs* 51(4): 727–753.

Clark, A. 2024. "Extending the Predictive Mind," *Australasian Journal of Philosophy* 102(1): 119–130. https://doi.org/10.1080/00048402.2022.2122523

Clark, A., and D. Chalmers. 1998. "The Extended Mind," *Analysis* 58(1): 7–19.

Colombetti, G., and T. Roberts. 2015. "Extending the Extended Mind: The Case for Extended Affectivity," *Philosophical Studies* 172: 1243–1263.

Constant, A., A. Clark, M. D. Kirchhoff, and K. J. Friston. 2022. "Extended Active Inference: Constructing Predictive Cognition Beyond Skulls," *Mind and Language* 37(3): 373–394. https://doi.org/10.1111/mila.12330

Erbas, A., and A. Okur. 2012. "Researching Students' Strategies, Episodes, And Metacognitions In Mathematical Problem Solving," *Quality and Quantity* 46(1): 89–102.

Facchin, M. 2023. "Extended Predictive Minds: Do Markov Blankets Matter?," *Review of Philosophy and Psychology* 14: 909–938. https://doi.org/10.1007/s13164-021-00607-9

Facchin, M. 2024. "Phenomenal Transparency, Cognitive Extension, and Predictive Processing," *Phenomenology and the Cognitive Sciences* 23: 305–327. https://doi.org/10.1007/s11097-022-09831-9

Facchin, M., and G. Leonetti. 2024. "Extended Animal Cognition," *Synthese* 203(5): 1–22.

Favela, L., M. Amon, L. Lobo, and A. Chemero. 2021. "Empirical Evidence for Extended Cognitive Systems," *Cognitive Science* 45(11): 1–45. https://doi.org/10.1111/cogs.13060

Gładziejewski, P. 2016. "Predictive Coding and Representationalism," *Synthese* 193(2): 559–582.

Hohwy, J. 2013. *The Predictive Mind*. Oxford University Press.

Hohwy, J. 2016. "The Self-evidencing Brain," *Nous* 50(2): 259–285.

Hohwy, J. 2018. "The Predictive Processing Hypothesis," in *The Oxford Handbook of 4E Cognition*, ed. A. Newen, S. Gallagher, and L. Bruin, 161–197. Oxford University Press.

Juni, M. Z., T. M. Gureckis, and L. T. Maloney. 2016. "Information Sampling Behaviour with Explicit Sampling Costs," *Decision* 3(3): 147–168.

Kersten, L. 2017. "Extended Music Cognition," *Philosophical Psychology* 30(8): 1078–1103. https://doi.org/10.1080/09515089.2017.1350946

Kersten, L. 2022. "A New Mark of the Cognitive? Predictive Processing and Extended Cognition," *Synthese* 200 (281): 1–25. https://doi.org/10.1007/s11229-022-03674-2

Kersten, L., and R. A. Wilson. 2016. "The Sound of Music: Externalist Style," *American Philosophical Quarterly* 52(4): 139–154.

Kirchhoff, M. D., and J. Kiverstein. 2019. *Extended Consciousness and Predictive Processing: A Third Wave View*. Routledge.

Kirchhoff, M. D., and J. Kiverstein. 2021. "How to Determine the Boundaries of the Mind: A Markov Blanket Proposal," *Synthese* 198: 4791–4810. https://doi.org/10.1007/s11229-019-02370-y

Kirsh, D., and P. Maglio. 1994. "On Distinguishing Epistemic from Pragmatic Actions," *Cognitive Science* 18(4): 513–549.

Krickel, B. 2020. "Extended Cognition, the New Mechanists' Mutual Manipulability Criterion, and the Challenge of Trivial Extendedness," *Mind and Language* 35(4): 539–561.

Mehlhorn, K., B. R. Newell, P. M. Todd, M. D. Lee, K. Morgan, V. A. Braithwaite, D. Hausmann, K. Fiedler, and C. Gonzalez. 2015. "Unpacking the Exploration-Exploitation Tradeoff: A Synthesis of Human and Animal Literatures," *Decision* 2(3): 191–215.

Menary, R. 2018. "Cognitive Integration: How Culture Transforms Us and Extends Our Cognitive Capabilities," in *The Oxford Handbook of 4E Cognition*, ed. S. Gallagher, A. Newen, and L. De Bruin, 187–215. Oxford University Press.

Parr, T., M. Dimitrije, K. Stefan, and K. Friston. 2019. "Neuronal Message Passing Using Mean-Field, Bethe, and Marginal Approximations," *Scientific Reports* 9 (1889): 1–18.

Pritchard, D., A. R. English, and J. Ravenscroft. 2021. "Extended Cognition, Assistive Technology and Education," *Synthese* 199(3–4): 8355–8377.

Rich, A. S., and T. M. Gureckis. 2018. "Exploratory Choice Reflects the Future Value of Information," *Decision* 5(3): 177–192.

Roberts, T., J. Kruger, and S. Glackin. 2019. "Psychiatry Beyond the Brain: Externalism, Mental Health, and Autistic Spectrum Disorder," *Philosophy, Psychiatry, and Psychology* 26(3): E-51–E-68. https://doi.org/10.1353/ppp.2019.0030

Rupert, R. 2022. "Cognitive Systems, Predictive Processing, and the Self," *Review of Philosophy and Psychology* 13: 947–972.

Smart, P. R. 2021. "Shedding Light on the Extended Mind: HoloLens, Holograms, and Internet-Extended Knowledge," *Frontiers in Psychology* 12: 675184.

Smart, P. R., G. Andrada, and R. Clowes. 2022. "Phenomenal Transparency and the Extended Mind," *Synthese* 200 (355). https://doi.org/10.1007/s11229-022-03824-6

Tajika, H., N. Nakatsu, H. Nozaki, E. Neumann, and S. Maruno. 2007. “Effects of Self-explanation as a Metacognitive Strategy for Solving Mathematical Word Problems,” *Japanese Psychological Research* 49(3): 222–233.

Trautmann, S. T., and G. van de Kuilen. 2015. “Ambiguity Attitudes,” in *The Wiley-Blackwell Handbook of Judgment and Decision Making*, ed. G. Keren and G. Wu, 89–116. John Wiley and Sons, Inc.

Vold, K., and D. Schlimm. 2020. “Extended Mathematical Cognition: External Representations with Non-Derived Content,” *Synthese* 197(9): 3757–3777.

William, D. 2023. “Hierarchical Minds and the Perception/Cognition Distinction,” *Inquiry* 66(2): 275–297. https://doi.org/10.1080/0020174X.2019.1610045

Wilson, R. A. 2004. *Boundaries of the Mind: The Individual in the Fragile Sciences*. Cambridge University Press.

Wilson, R. A., and A. Clark. 2009. “How to Situate Cognition: Letting Nature Take Its Course,” in *The Cambridge Handbook of Situated Cognition*, ed. P. Robbins and M. Aydede, 55–77. Cambridge University Press.