

UC Irvine

UC Irvine Previously Published Works

Title

Understanding the Financial Markets as a Sociotechnical Extended Cognitive System

Permalink

<https://escholarship.org/uc/item/9539f4ww>

Author

Pritchard, Duncan

Publication Date

2025-11-17

DOI

10.4324/9781003567622-5

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

For *Digital Development in Artificial Systems: Technology, Ethics and Governance*,
(eds.) M. Farina, X. Yu & J. Chen (London: Routledge).

UNDERSTANDING THE FINANCIAL MARKETS AS A SOCIOTECHNICAL EXTENDED COGNITIVE SYSTEM

Duncan Pritchard

University of California, Irvine

dhpritch@uci.edu

ORCID: [0000-0002-5997-0752](https://orcid.org/0000-0002-5997-0752)

ABSTRACT. Our interest is in understanding the financial markets as a sociotechnical cognitive system, akin to the kind of sociotechnical cognitive systems that we find in other highly collaborative and technologically dependent cognitive systems, such as in certain kinds of scientific inquiry. It is proposed that we should unpack this notion of sociotechnical cognition through the lens of the *extended cognition* research program. A three-tiered account of sociotechnical cognition is set out—along with the corresponding variety of sociotechnical knowledge that results—and applied to the special case of financial expertise employed within a market setting. The first tier is *sociotechnical facilitated cognition*. This results in *sociotechnical facilitated knowledge* that is individual and non-extended. The second tier is *sociotechnical extended cognition*. This generates *sociotechnical extended knowledge* that is individual and extended. The third tier is *sociotechnical distributed cognition*. This generates *sociotechnical distributed knowledge* that is irreducibly attributable to a group agent rather than to individuals within that group. It is argued that financial expertise employed in market settings is at least typically at most a form of sociotechnical extended cognition (thus leading to only sociotechnical extended knowledge).

KEYWORDS: Epistemology; Distributed Knowledge; Extended Cognition; Extended Epistemology; Philosophy of Economics.

1. INTRODUCTORY REMARKS

It is natural to think of the financial markets as a type of cognitive system that is essentially both social and technological in nature. Skillfully operating in these markets requires an extensive dependence on both technology (e.g., financial technology, or ‘fintech’ as it is more widely known) and the expertise of others (e.g., experts on a given stock, a particular financial market, the geopolitical situation in a region that is important to a trade, and so on). This entails that the kinds of theoretical frameworks that apply to sociotechnical cognitive systems ought to be, in principle at least, applicable to our understanding of financial markets and thereby cast light on how they function.

It is in that spirit that there has been a growing discussion in some quarters regarding how best to understand the sociotechnical cognitive system in play in the financial markets. For example, one might ask to what extent they are akin to advanced scientific research teams, given that these are also regarded as sociotechnical cognitive systems (we will come back to this analogy). Indeed, some have even postulated a quite radical hypothesis in this regard, which is that we should take seriously the idea of there being a ‘market mind’, where this is to be understood as not simply the aggregation of, say, individual traders and their actions but rather as some kind of emergent property of the financial cognitive system.¹

While I find this topic—and in particular the radical proposal of the market mind hypothesis—to be very interesting, I also think that there are a lot of important philosophical distinctions that often get glossed over in this discussion of the cognitive system involved in the financial markets. I therefore propose to get back to first principles. This will involve mapping out the conceptual terrain and then seeing how it applies to the case of financial markets *qua* cognitive system.

As we will see, I will be especially concerned with the epistemological implications of different ways of understanding this sociotechnical system. For if we take the idea of the financial markets as a sociotechnical cognitive system seriously, then that implies that the financial expertise involved in this sociotechnical cognitive system produces a kind of sociotechnical financial knowledge. (By way of comparison, think of how the sociotechnical system of advanced scientific research terms involves scientific expertise that leads to sociotechnical scientific knowledge). How is this sociotechnical financial knowledge to be understood?

Drawing on previous work, I will be proposing a three-tiered conception of sociotechnical cognition and the knowledge that corresponds to it.² As we will see, this approach will make appeal to the *extended cognition* research program in the philosophy of cognitive science.³ In particular, it will appeal to the epistemological ramifications of that research program, what I have previously described as *extended knowledge*.⁴

At the first level, there is *sociotechnical facilitated cognition*. This generates *individual non-extended knowledge*—i.e., knowledge that is possessed by the agent alone and which is attributable to her regular (non-extended) cognitive processes. At the second level, there is *sociotechnical extended cognition*. This generates *individual extended knowledge*. At this level, this knowledge is also individual knowledge, but it differs from the knowledge that is acquired via *sociotechnical facilitated cognition* in that it is attributable to extended cognitive processes. Roughly (we will explain this distinction in more detail below), this is when sociotechnological features of the subject's environment become a proper part of the subject's cognitive processes. At the third level there is *sociotechnical distributed cognition*. Unlike the other two levels, this results not in individual knowledge but rather *sociotechnical distributed knowledge*. This is when the knowledge that is being generated is properly attributable to a group of agents rather than to any individual agent within that group.

I will be suggesting that market knowledge—i.e., the kind of knowledge market experts acquire—is often best understood as, at most, sociotechnical extended knowledge. If that's right, then this has an important bearing on how we should construe an idea like the 'market mind'. In particular, we are not to think of this as being a sociotechnical distributed cognitive system that generates distributed knowledge (which is what I think is the most natural interpretation of how a market mind should be understood).

A final point before we get started. I will be understandings market knowledge in a very specific sense, on the grounds that I think this is the proper target for any discussion about the market as a cognitive system. My focus will be on the expertise exhibited by agents within this cognitive system, such as financial traders. In particular, I will not be interested in the role that, say, an inexperienced amateur investor might play in the markets, since although they are clearly participants in this cognitive system, they are not exhibiting the relevant expertise that the seasoned professionals exhibit. Relatedly, I am also not interested in another sense of financial expertise that one might have in mind in this context, which is concerned not with the participants of the financial cognitive system but with those who

study and theorize about it (e.g., economists). While that is genuine financial expertise, it is one step removed from the cognitive system itself, and hence falls outside the bounds of our discussion.

2. THE 'MARKET MIND'

I noted in the last section that there is a theoretical discussion about the cognitive system involved in the financial markets that characterizes it as a 'market mind'. So let's begin by reflecting on what might be meant by thinking of the financial cognitive system in this way. I think at least two readings come to the fore (there are no doubt more).

I suspect the reading that most theorists who hear this terminology would think of is that the market is an emergent, conscious, group mind. The idea would be that although there are obviously individual subjects engaging in this cognitive system, there is nonetheless an emergent property that results from these interactions that is irreducible to the minded activity of the participating subjects.

That's certainly one possibility, though it is metaphysically quite extravagant. One would need very good reasons to think that a market mind in this sense was on display in the financial markets. This leads us to a second reading of the market mind, one that is less dramatic. Perhaps the market mind refers to a kind of group agency that is reducible (albeit perhaps in a complex fashion) to the agency of the participating subjects. That is not as metaphysically extravagant as the first reading, as it doesn't imply that there is an irreducible conscious group mind. But it does capture the idea that there might be something quite interesting that comes into being when financial experts are embedded within the sociotechnical cognitive system that makes up the financial markets.

The hard question raised by these two readings is what would be required to determine whether we should opt for either reading. Since the second reading doesn't commit to a strong metaphysical thesis, then it can lay claim to being the default in this regard. But that doesn't answer the question of when we should opt for the first reading over the second; it merely says that if in doubt we should prefer the second reading over the first. What would it take for a sociotechnical cognitive system like that in play in the financial

markets to manifest an emergent agency that is irreducible to the individual agencies involved in that cognitive system?

With this question in mind, I want to shift the focus away from the market mind metaphor for now. My suggestion is that we should simply treat the financial cognitive system as a sociotechnical cognitive system and then try to determine what kind of sociotechnical cognitive system it is. This approach has the advantage of not settling in advance how we should understand this particular sociotechnical cognitive system. Rather, our strategy is to first offer a taxonomy of different kinds of sociotechnical cognitive systems and then, second, ask the question of where the specific sociotechnical cognitive system on display in the financial markets appears within this taxonomy. In this way, we can appeal to the wider literature on sociotechnical cognitive systems, rather than immediately being drawn into the specific details of the cognitive system involved in the financial markets.

3. SOCIOTECHNICAL COGNITIVE SYSTEMS

The expertise involved in generating financial market knowledge looks a lot like other cases of expertise that is embedded within sociotechnical cognitive systems. Consider, for example, the scientific sociotechnical cognitive networks involved in High Energy Physics, which have been widely studied.⁵ This involves a complex interchange between a varied array of experts with very specific skills, complex technology, temporal and geographical distribution of the main actors, widespread epistemic dependence on both technology and the expertise of others (with the rational trust that this implies), and so on.

Financial market expertise exhibits many of these features. It certainly involves specialization of expertise (a trader might be an expert only within one particular market, for example). It also involves epistemic dependence on both the expertise of others and technology. The decision maker in a contemporary financial brokerage will be relying not just on the expertise of her team, with their very specific skill sets, but also on the technology that they employ in order to reach their recommendations. There is also often a temporal and geographic distribution of the main actors (including, even, the technology). If one is managing an international investment portfolio, then one might be soliciting expertise from

a range of people across the globe, and in the process relying on technology that is also globally distributed.

Theorists of the financial markets make explicit the extent to which it is a sociotechnical cognitive system. For example, Donald MacKenzie, one of the leading sociologists of science and technology who has specialized in examining the financial markets, has emphasized how such markets constitute complex cognitive systems that involve high levels of epistemic dependence on both technology and expertise. For example, MacKenzie (2021) offers a comprehensive account of a particular kind of highly technological financial activity, known as high-frequency trading, where this epistemic dependence is quite stark. Essentially, for this trading to generate profits, gaining trading information even a fraction of a second faster than the rest of the market is crucial. But for that to occur, one needs to rely on the latest technology and the very best expertise.

Moreover, social scientists who study the financial markets emphasize how the technological and collaborative nature of the financial cognitive system leads to feedback loops, just as it does in other sociotechnical cognitive systems. For example, consider how the particular technological features of a scientific inquiry could lead to the scientists being either physically co-present as they inquire (e.g., if the inquiry can only be done in a specific laboratory) or could lead to them being geographically dispersed (e.g., if it is important to the measurements being taken that the scientists are located at specific places worldwide). Something similar occurs with financial trading, as MacKenzie here notes:

“[...] trading is not reducible to human action alone: it is a material process that is shaped strongly by non-human actants such as technologies, whether the latter are as simple as the octagonal, steppe amphitheatres of Chicago’s traditional trading pits or as complex as the electronic infrastructures of today’s markets.” (MacKenzie 2019)⁶

Take high-frequency trading again. As MacKenzie points out, the need to gain a fraction of a second on one’s competitors leads these traders to set up shop as close to the source of the financial information as possible, to maximize their ability to gain this information quickly.

This is not to say that there aren’t differences between the sociotechnical cognitive systems involved in highly theoretical scientific collaborations and in the kind of financial trading that we are interested in. Indeed, even within scientific collaborations there are differences of epistemic culture. Karin Knorr-Cetina (1999), one of the foremost social science theorists of scientific collaboration, has highlighted how, for example, a biological

laboratory will often have a vertical and localized epistemic culture, in the sense that everyone in the research team is tied to that particular laboratory and the person who is head of the laboratory is in charge of the direction of the scientific inquiry. In contrast, other kinds of scientific collaboration might be more horizontal and distributed (as seems to be the case with much of the high-profile physics research conducted at the famous CERN laboratory in Geneva, otherwise known as the European Laboratory for Particle Physics).

Where would we expect to find the financial market cognitive system in this regard? Well, it is certainly usually distributed rather than local, for the reasons given above. But there are also grounds for thinking that such a cognitive system is likely to be more vertical than horizontal due to the corporate nature of the structures involved. There is a lot of money to be made (or lost), and a lot hangs on financial results. Accordingly, one would expect there to be a culture in the financial markets that exhibits tighter chains of authority than one finds in scientific collaborations.

Note, however, that the differences here are ultimately of degree rather than kind. A lot of money can hang on a series tests run at CERN, for example, and the extra pressure might lead to a more vertical epistemic culture there as a result. Moreover, as we've already noted, not every kind of collaborative and highly technical scientific inquiry has the same epistemic culture anyway. What is important for our purposes is rather that the financial markets exhibit essentially the same kind of features as other well-studied sociotechnical cognitive systems, even if they may differ from these other sociotechnical cognitive systems when it comes to the details. Accordingly, we are on strong ground in treating them as a sociotechnical cognitive system and then determining the specific kind of sociotechnical cognitive system they exhibit.⁷

4. EXTENDED COGNITION AND EXTENDED KNOWLEDGE

We can understand sociotechnical cognitive systems through the lens of *extended cognition*. Extended cognition occurs when features of the subject's cognitive environment, such as technology, become seamlessly integrated into that subject's on-board cognitive processes to such a degree that they become a proper part of that cognitive process. In the right conditions, this is more than just technological assistance with cognitive tasks; it is *cognitive*

augmentation.⁸ For our purposes, it will be useful to stipulate that by extended cognition we mean that the technology in question is either information-processing or information-retaining and that it also plays a pivotal role in the formation of beliefs (as opposed to subdoxastic states).⁹

Think, for example, of how (if you're as old as I am anyway) you used to have 20-30 phone numbers that you had memorized (friends and family mainly), but that now you off-load that cognitive task to your iPhone (and also to your iCloud, so that the information is on all your devices should you need it). Since your iPhone is always with you (and the iCloud always available), there is no need to rely on one's biological memory in this way, when one can instead rely on the non-biological storage that is available. The crux of the matter is that if one's use of the iPhone to perform this cognitive task is sufficiently integrated into one's cognitive practices—where that means that one's use of the non-biological 'memory' provided by the iPhone is roughly analogous to how one uses one's biological memory, in that it is generally seamless, automatic, and so forth—then it counts as an extended cognitive process. This means that the technology has become a proper part of one's memorial cognitive processes, such that this cognitive process no longer includes only one's on-board, biological, memory, but also includes the contribution of this technology. Put another way, one's memorial cognition now includes *extended* cognitive processes that have the technology as a proper part.¹⁰

Once we allow that there can be extended cognitive processes involving technology, however, then it is but a short step to allowing that there can be extended cognitive processes involving other agents. For if features of the environment such as technology can function as proper parts of one's cognitive processes, then why couldn't social features of one's environment—i.e., other people—also function as proper parts of one's cognitive processes? We can thus distinguish between technologically extended cognition and socially extended cognition.

Examples of socially extended cognition include transactional memorial systems. Paradigmatically, this concerns couples of long-standing where over time they form a cognitive system whereby one individual takes on responsibility for some aspects of memorial cognition and the other individual takes on responsibility for other aspects of memorial cognition. So perhaps individual 1 is the person who remembers dates (e.g., of important past and future events, of family and friends' birthdays, and so on), whereas

individual 2 just devolves this memorial activity to person 1. And perhaps individual 2 has become the designated person to remember important personal information that is in common to both individuals (e.g., they remember banking information or information about pensions), while individual 1 has over time simply devolved this memorial task to individual 2.

The result is that we have an extended cognitive system that is essentially social in nature. Rather than technology being integrated into one's cognitive processes, there is instead the systematic contributions of another agent. Accordingly, just as in cases of technological cognitive extension where we attribute the cognitive success to the extended cognitive process that includes the technology as a proper part, so in cases of social cognitive extension, like transactive memorial systems, we attribute the cognitive success to the extended cognitive process that includes the other agent. For example, for individual 2 to remember a notable date, like a birthday, we would consider the interventions of individual 1 to be akin to the technological augmentation that we considered above, whereby they dependably manifest the information retention that individual 2 employs as a route to memorial knowledge.¹¹

Now that we have the distinction between technological extended cognition and social extended cognition in place, we can also make sense of sociotechnical extended cognition, where both technology and other agents are proper parts of the target cognitive process. Let's return to our example of scientific collaboration on a highly technical problem. In such collaborations there is not only an epistemic dependency on technology but also on the expertise of other members of the team. Moreover, over time this dependence on technology and expertise can become so seamless and integrated into one's cognitive practices that the use of the technology and expertise becomes a proper part of one's cognitive processes. In short, they become extended cognitive processes, albeit where the extension is both technological and social. The result is what we are here calling sociotechnical extended cognition.

Can extended cognitive processes lead to knowledge in essentially the same way that non-extended cognitive processes can lead to knowledge? Does the contribution of factors that are external to one's (biological) cognitive agency (in this case both technological and social) to one's cognitive success entail that this cognitive success is thereby not of the right epistemic pedigree to count as knowledge? I think there are a number of reasons to resist

this line of thinking, and so regard extended cognitive processes as not being in their nature epistemically problematic.

To begin with, as I've argued elsewhere, it is not essential that one's cognitive success should be primarily attributable to one's cognitive agency in order for cognitive success to amount to knowledge. Instead, all that is required is that one's cognitive success is at least significantly (i.e., non-negligibly) attributable to one's cognitive agency. This is a modest form of what is known as *virtue epistemology*, a version that allows for the possibility that factors external to one's agency can also play a significant explanatory role in one's cognitive success without thereby precluding it from being knowledge.¹²

This is why, for example, one can come to know something by for the most part trusting the word of a recognized expert. Clearly, one's cognitive success in this scenario is not primarily due to one's own cognitive agency, as the expert's cognitive agency is carrying a significant explanatory load here too. And yet, since one's cognitive agency is also bearing a significant explanatory load—one consulted the recognized expert, remember—then one's cognitive success can count as knowledge.

Note that in offering this line of argument we are effectively conceding that in cases of extended cognition the technology and expertise really is external to one's own cognitive agency. But, on the face of it, that is already to concede too much. Think of the scientific collaborators, where I think we would be very inclined to regard their sociotechnical extended cognitive successes as instances of knowledge. It is not as if in such cases the scientists are blindly employing the technology and expertise that they depend on. Rather, they have over time integrated this technology and expertise into their cognitive processes. Moreover, that this integration has taken place will reflect the fact that they regard these factors as good routes to the truth.

A better way to think about this issue is thus that this transition to sociotechnical extended cognitive processes are manifestations of cognitive agency. Indeed, isn't this the very spirit of extended cognition? Put another way, rather than the 'real' cognitive agent being the 'on-board' biological cognitive processes, we should instead take seriously that extended cognitive agency is a kind of cognitive agency on a par with biological cognitive agency.

If that's right, then we don't even need to concede that extended cognition entails that one's cognitive success will tend to be less attributable to one's cognitive agency than in

parallel cases involving non-extended cognition. Instead, we should expand the scope of virtue epistemology—call this *extended virtue epistemology*—so that it allows that one’s cognitive agency can be understood so that it can include extended cognitive processes.¹³ Yes, we can distinguish between extended knowledge and regular knowledge, but that distinction simply reflects the etiology of the cognitive success in question and doesn’t signal that the former is any less legitimate than the latter.

With the category of sociotechnical extended knowledge on the table, we can now make a distinction that will be very useful for our purposes. Recall that in order for a cognitive process to become extended it needs to meet certain criteria—not any use of technology or, say, expertise is thereby an extended cognitive process. We have noted, for example, that the external factor needs to be properly integrated within the cognitive processes over time, such that it starts to exhibit the same kind of features as the corresponding on-board cognitive processes (it is employed seamlessly, automatically, and so on).

With this point in mind, we can draw a distinction between merely sociotechnical *facilitated* knowledge and sociotechnical *extended* knowledge. The latter will meet the criteria for sociotechnical extended knowledge while the former will not. In particular, the former will involve the subject using technology and the testimony of other agents as simply instruments that their cognitive agency employs to get to the truth. In this sense, the contribution of the technology/other agents will be external to the subject’s own cognitive agency and merely resources for that agency. In contrast, when it comes to genuine cases of sociotechnical extended knowledge, this will involve the subject utilizing the technology and other people’s cognitive agency in ways that are fully integrated into their cognitive processes such that it counts as a proper part of those processes, and hence manifests their extended cognitive agency.

We are also now in a position to introduce a third, even more demanding, tier of sociotechnical cognition. This is *sociotechnical distributed cognition*. This occurs when the sociotechnical cognitive system produces epistemic outputs, like knowledge, that are not individual at all, but rather distributed across the cognitive system as a whole. Take our collaborative scientific team, for example. Could it be that the level of integrated collaboration on display here is such that there are epistemic outputs that are not attributable to any specific member of the team, but rather attributable to the team as a whole? For

example, that there is scientific knowledge that is arising from their collaborations that is not attributable to any individual within that team but which are attributable to the cognitive system itself? If so, then this would be a case of sociotechnical distributed cognition leading to sociotechnical distributed knowledge.

Whether there are any genuine cases of sociotechnical distributed cognition, as just described anyway, is controversial.¹⁴ It's certainly true that we often attribute knowledge to a group when there are tight-knit collaborations in play, particularly ones involving high levels of expertise. But usually we would also attribute that knowledge to individuals within the group too, so that wouldn't qualify as sociotechnical distributed cognition as we just defined it. Moreover, if we discovered that no-one within a collaborative scientific team had knowledge of the target claim, then that would prompt us to doubt whether the team as a whole could possibly know that claim.

It can be tempting to run sociotechnical extended knowledge and sociotechnical distributed knowledge together, as if they were one and the same. Both involve a kind of social cognition, after all. But the difference between the two is significant. Sociotechnical extended knowledge is still individual knowledge, after all, unlike sociotechnical distributed knowledge, which is by its nature group knowledge that isn't manifested in the individuals in that group. Once we recognize this distinction, it becomes hard to understand how scientific research teams could be interpreted in terms of sociotechnical distributed knowledge as opposed to sociotechnical extended knowledge.

As previously noted, we are usually only happy to attribute knowledge to a scientific team when we also believe that individual members of that team have that knowledge. That goes against the idea of sociotechnical distributed knowledge. There is also another reason for thinking that scientific collaborations are at most sociotechnical extended cognitive systems. This is that it is part of good scientific practice that one keeps track of who is responsible for what in terms of the epistemic outputs. Interestingly, we see this conspicuously in action when it comes to scientific malpractice, as when a particular researcher in a research team fakes results. The other researchers are usually able to chronicle their own contributions to the research endeavor in quite detailed ways in order to show that they are blameless in this scientific malpractice, despite the fact that their name is now associated with research results that are now problematic. The fact that this is possible shows how the attribution relations in collaborative scientific endeavors are primarily at the

individual rather than the group level, at least in the sense that the latter can be reduced to the former if required. That indicates that collaborative scientific teams are not sociotechnical distributed cognitive systems but, at most, sociotechnical extended cognitive systems. (Relatedly, what they produce is at most sociotechnical extended knowledge and not sociotechnical distributed knowledge).¹⁵

5. UNDERSTANDING MARKET EXPERTISE

We are now in a position to understand the canonical kind of market expertise that interests us. We have already listed the parallels between highly theoretical scientific collaborations and the cognitive system in play in market expertise. As we noted, both involve a high level of epistemic dependence on both technology and other peoples' expertise. That ensures that these are sociotechnical cognitive systems.

Moreover, this epistemic dependence on technology and expertise is often deeply integrated within both cognitive systems, so that it employed seamlessly. When this occurs, it is plausible that these are not just sociotechnical cognitive systems, but are also, more specifically, *extended* sociotechnical cognitive systems. This means that the knowledge that they produce is extended knowledge.

We have argued, however, that extended knowledge is, at least typically anyway, individual knowledge—i.e., knowledge that is properly attributable to individual agents within that cognitive system. In particular, we have drawn a distinction between sociotechnical extended cognitive systems and sociotechnical distributed cognitive systems. It is only in the latter that knowledge (or some other positive epistemic output) is produced that is properly attributable to the cognitive whole and not to any of the individual agents who are part of that cognitive whole.

Furthermore, we have argued that when it comes to scientific collaborations, it is not credible to suppose that these are sociotechnical distributed cognitive systems. The reason is that it is built into the proper functioning of such scientific collaborations that they track individual attribution relations. It is only by doing so that one can ensure the reliability of such collaborative inquiries. When these inquiries are properly conducted, the individual attribution relations may not be made explicit. But we can see from when inquiries are not

properly conducted, and epistemic blame is appropriate, that these individual attribution relations have been tracked all along.

What goes here for scientific collaboration applies with even more force to the kind of sociotechnical cognitive system involved in the financial markets. The reason for this is that there is even more need to track individual attribution relations when it comes to this cognitive system. We can distinguish at least two crucial factors in this regard.

The first returns us to a point that we made earlier: the financial stakes are higher. If there is error, then large sums of money may be lost. Moreover, even where money is made rather than lost, the attribution relations become especially important if one is aiming to make a case for one's own share of these profits (whether directly via commissions, or indirectly via pay raises or bonuses).¹⁶ Relatedly, while the professionals operating in the financial markets are also, like scientists in a research team, collaborating with other experts in their field, they will likely be more inclined to regard themselves, and their cognitive contribution to the cognitive outcome, in individualistic terms, at least when compared with a scientific collaboration.

The second is that the legal stakes are also higher. If one messes up a scientific collaboration due to incompetence or, worse, deception (faking evidence, say), then one is unlikely to face any legal consequences, only consequences for one's career. In contrast, incompetence or, worse, deception when operating in the financial markets could lead to both civil and criminal prosecution. It is thus imperative that records are kept that track the individual contributions of each party (and the associated technology) to each financial transaction.

The upshot is that there is even more reason to be skeptical of thinking of the cognitive system involved in the financial markets in distributed sociotechnical terms, as opposed to merely as an extended sociotechnical system.

6. CONCLUDING REMARKS: THE MARKET MIND, REDUX

So whither now the 'market mind'? Let's review our conclusions. We have distinguished between three levels of sociotechnical cognition and the social knowledge that results. The

first level is sociotechnical facilitated cognition. This results in sociotechnical knowledge that is both individual in nature and non-extended. The second level is sociotechnical extended cognition. This results in sociotechnical extended knowledge. We have seen that there are good grounds to think that the sociotechnical cognitive system involved in manifestations of financial expertise tends to conform to this second level.

We have also noted a third level of sociotechnical cognition. This is sociotechnical distributed cognition. This is where the cognition is at the level of the cognitive whole and not manifested by any of the individual agents who are part of this cognitive system. This means that while the cognitive system exhibits positive epistemic outputs, like knowledge, these are not manifested by the individual agents who are part of this cognitive system.

We have argued that not only is sociotechnical distributed cognition a bad way of understanding scientific collaboration, but it is also an even worse way of understanding the kind of epistemic dependence on technology and expertise involved in the financial markets. The distinctive nature of financial trading means that the considerations that make us skeptical of thinking of scientific collaborations along distributed cognition lines are even more stark when it comes to the kinds of collaborations involved in the financial cognitive system.

So where does that leave us with the idea of the ‘market mind’? I don’t think that it means that we should reject it entirely, but it does entail that we should be wary about taking it too literally. If the idea of the financial markets in terms of sociotechnical distributed cognition is not credible, then there cannot be an emergent mind at work within this cognitive system, as the former would surely be a pre-requisite for the latter. Nonetheless, that doesn’t mean that there isn’t some truth to the idea of a market mind. After all, we have seen that it is plausible to regard the cognitive system at issue in the financial markets as a form of sociotechnical extended cognition. This type of cognitive system generates sociotechnical extended knowledge, and that’s at least one way—if not the most immediate way—in which we might think that the financial markets involve a market mind. At the very least, it captures the thought that the financial markets are more than just cognitive agents acting individually within contexts that are highly reliant on technology and expertise, but are rather cognitive systems where the technology and expertise become proper parts of those cognitive agents’ cognitive processes.¹⁷

REFERENCES

- Adams, F., & Aizawa, K. (2001). 'The Bounds of Cognition', *Philosophical Psychology* 14, 43-64.
- (2008). *The Bounds of Cognition*, Blackwell, Oxford.
- Alfano, M., & Skorburg, G. (2016). 'The Embedded and Extended Character Hypotheses', *Routledge Handbook of Philosophy of the Social Mind*, (ed.) J. Kiverstein, 465-78, London: Routledge.
- Bourdieu, P. (2005). *The Social Structures of the Economy*, London: Polity Press.
- Bird, A. (2017). 'When is There a Group that Knows? Distributed Cognition, Scientific Knowledge, and the Social Epistemic Subject', *Essays in Collective Epistemology*, (ed.) J. Lackey, ch. 2, Oxford: Oxford University Press.
- (2019). 'Group Belief and Knowledge', *Routledge Handbook of Social Epistemology*, (eds.) M. Fricker, P. Graham, D. Henderson & N. Pedersen, 274-83, London: Routledge.
- Carter, J. A. (2013). 'Extended Cognition and Epistemic Luck', *Synthese* 190, 4201-14.
- Carter, J. A., Kallestrup, J., Palermos, S. O., & Pritchard, D. H. (2014). 'Varieties of Externalism', *Philosophical Issues* 24, 63-109.
- Carter, J. A., Clark, A., Kallestrup, J., Palermos, S. O., & Pritchard, D. H. (Eds.), (2018a). *Extended Epistemology*, Oxford: Oxford University Press.
- (Eds.), (2018b). *Socially Extended Epistemology*, Oxford: Oxford University Press.
- Carter, J. A., & Pritchard, D. H. (2019). 'The Epistemology of Cognitive Enhancement', *Journal of Medicine & Philosophy* 44, 220-42.
- Carter, J. A., & Rupert, R. D. (2020). 'Epistemic Value in the Subpersonal Vale', *Synthese* 198, 9243-72.
- Chemero, T. (2009). *Radical Embodied Cognitive Science*, Cambridge, MA: MIT Press.
- Clark, A., (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*, Oxford: Oxford University Press.
- (2010). 'Memento's Revenge: The Extended Mind, Extended', *The Extended Mind*, (ed.) R. Menary, 43-66, Cambridge, MA: MIT Press.
- (2013). 'Gesture as Thought', *The Hand, an Organ of the Mind: What the Manual Tells the Mental*, (ed.) Z. Radman, 255-68, Cambridge, MA: MIT Press.
- (2015a). 'What 'Extended Me' Knows', *Synthese* 192, 3757-75.
- (2015b). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*, Oxford: Oxford University Press.
- (2023). *The Experience Machine: How Our Minds Predict and Shape Reality*, New York: Pantheon.
- Clark, A., & Chalmers, D. (1998). 'The Extended Mind', *Analysis* 58, 7-19.
- Gallagher, S. (2005). *How the Body Shapes the Mind*, Oxford: Oxford University Press.
- Gilbert, M. (2004). 'Collective Epistemology', *Episteme* 1, 95-107.
- Goldman, A. (1986). *Epistemology and Cognition*, Cambridge, MA: Harvard University Press.
- Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*, Cambridge: Cambridge University Press.
- Huebner, B. (2016). 'Transactive Memory Reconstructed: Rethinking Wegner's Research Program', *Southern Journal of Philosophy* 54, 48-69.
- Jarvis, B. (2015). 'Epistemology and Radically Extended Cognition', *Episteme* 12, 459-78.
- Knorr-Cetina, K. (1999). *Epistemic Cultures: How the Sciences Make Knowledge*, Cambridge, MA: Harvard University Press.
- Knorr Cetina, K., & Brugger, U. (2002). 'Traders' Engagement with Markets: A Postsocial Relationship', *Theory Culture Society* 19, 161-85.
- Knorr Cetina, K., & Preda, A. (eds) (2004). *The Sociology of Financial Markets*, Oxford: Oxford

- University Press.
- Lackey, J. (2021). *The Epistemology of Groups*, Oxford: Oxford University Press.
- Latour, B. (1999). *Pandora's Hope: An Essay on the Reality of Science Studies*, Cambridge, MA: Harvard University Press.
- Lynch, M. P. (2014). 'Neuromedia, Extended Knowledge, and Understanding', *Philosophical Issues* 24, 299-313.
- MacKenzie, D. (2019). 'Market Devices and Structural Dependency: The Origins and Development of 'Dark Pools'', *Finance and Society* 5, 1-19.
- (2021). *Trading at the Speed of Light: How Ultrafast Algorithms Are Transforming Financial Markets*, Princeton, NJ: Princeton UP.
- Menary, R. (2010). *Cognitive Integration: Mind and Cognition Unbounded*, London: Palgrave.
- Michaelian, K., & Sutton, J. (2018). 'Collective Memory', *Routledge Handbook of Collective Intentionality*, (eds.) M. Jankovic & K. Ludwig, 140-51, London: Routledge.
- Michaelian, K., & Arango-Muñoz, S. (2018). 'Collaborative Memory Knowledge: A Distributed Reliabilist Perspective', *Collaborative Remembering: Theories, Research, and Applications*, (eds.) M. Meade, A. Barnier, P. Van Bergen, C. Harris & J. Sutton, 231-47, Oxford: Oxford University Press.
- Noë, A. (2004). *Action in Perception*, Cambridge, MA: MIT Press.
- Palermos, S. O. (2011). 'Belief-Forming Processes, Extended', *Review of Philosophy and Psychology* 2, 741-65.
- (2014a). 'Knowledge and Cognitive Integration', *Synthese* 191, 1931-51.
- (2014b). 'Loops, Constitution, and Cognitive Extension', *Cognitive Systems Research* 27, 25-41.
- (2020). 'Epistemic Collaborations: Distributed Cognition and Virtue Reliabilism', *Erkenntnis* [Online First, DOI: 10.1007/s10670-020-00258-9].
- Pauw, W. T. J. L., de Nooijer, J. A., & Paas, F. (2014). 'Toward a More Embedded/Extended Perspective on the Cognitive Function of Gestures', *Frontiers in Psychology* 5, 359.
- Plantinga, A. (1993). *Warrant: The Current Debate*, Oxford: Oxford University Press.
- Pritchard, D. H. (2010). 'Cognitive Ability and the Extended Cognition Thesis', *Synthese* 175, 133-51.
- (2012). 'Anti-Luck Virtue Epistemology', *Journal of Philosophy* 109, 247-79.
- (2018a). 'Extended Knowledge', *Extended Epistemology*, (eds.) J. A. Carter, A. Clark, J. Kallestrup, S. O. Palermos & D. H. Pritchard, 90-104, Oxford: Oxford University Press.
- (2018b). 'Extended Virtue Epistemology', *Inquiry* 61, 632-47.
- (2018c). 'Neuromedia and the Epistemology of Education', *Metaphilosophy* 49, 328-49.
- (2022). 'Socially Extended Scientific Knowledge', *Frontiers in Psychology* [DOI: 10.3389/fpsyg.2022.894738].
- (2025). 'Extended Social Epistemology', *Oxford Handbook of Social Epistemology*, (eds.) J. Lackey & A. McGlynn, Oxford: Oxford University Press.
- Pritchard, D. H., Millar, A., & Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations*, Oxford: Oxford University Press.
- Rowlands, M. (2009). *The Body in Mind*, New York: Cambridge University Press.
- Rupert, R. D. (2004). 'Challenges to the Hypothesis of Extended Cognition', *Journal of Philosophy* 101, 389-428.
- (2009). *Cognitive Systems and the Extended Mind*, Oxford: Oxford University Press.
- Schotanus, P. (2023). *The Market Mind Hypothesis: Understanding Markets and Minds Through Cognitive Economics*, Berlin, Germany: De Gruyter.
- Shapiro, L. (2011). *Embodied Cognition*, New York: Routledge.

- Sosa, E. (1991). *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge: Cambridge University Press.
- (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge (vol. 1)*, Oxford: Oxford University Press.
- (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge (vol. 2)*, Oxford: Oxford University Press.
- (2015). *Judgment and Agency*, Oxford: Oxford University Press.
- Sprevak, M. (2009). ‘Extended Cognition and Functionalism’, *Journal of Philosophy* 106, 503-27.
- (2019). ‘Extended Cognition’, *Routledge Encyclopedia of Philosophy*, (ed.) T. Crane, London: Routledge [DOI: 10.4324/9780415249126-V049-1].
- Sutton, J. (2008). ‘Between Individual and Collective Memory: Coordination, Interaction, Distribution’, *Social Research* 75, 23-48.
- Sutton, J., Barnier, A., Harris, C., & Wilson, R. (2008). ‘A Conceptual and Empirical Framework for the Social Distribution of Cognition: The Case of Memory’, *Cognitive Systems Research* 1-2, 33-51.
- Tuomela, R. (2004). ‘Group Knowledge Analyzed’, *Episteme* 1, 109-27.
- Tollefsen, D. (2019). ‘The Epistemology of Groups’, *Routledge Handbook of Social Epistemology*, (eds.) M. Fricker, P. Graham, D. Henderson & N. Pedersen, 263-73, London: Routledge.
- Wegner, D., Giuliano, T., & Hertel, P. (1985). ‘Cognitive Interdependence in Close Relationships’, *Compatible and Incompatible Relationships*, (ed.) W. J. Ickes, 253-76, New York: Springer-Verlag.
- Wegner, D. (1986). ‘Transactive Memory: A Contemporary Analysis of the Group Mind’, *Theories of Group Behavior*, (eds.) B. Mullen & G. R. Goethals, 185-208, New York: Springer-Verlag.
- Wheeler, M. (2005). *Reconstructing the Cognitive World*, Cambridge, MA: MIT Press.
- (2013). ‘Is Cognition Embedded or Extended? The Case of Gestures’, *The Hand, an Organ of the Mind: What the Manual Tells the Mental*, (ed.) Z. Radman, 269-301, Cambridge, MA: MIT Press.
- Wikforss, A. (2014). ‘Extended Belief and Extended Knowledge’, *Philosophical Issues* 24, 469-81.
- Wray, K. B. (2018). ‘How Far can Extended Knowledge be Extended? The Asymmetry Between Research Teams and Artifacts’, *Socially Extended Epistemology*, (eds.) J. A. Carter, A. Clark, J. Kallestrup, S. O. Palermos & D. H. Pritchard, 11-23, Oxford: Oxford University Press.
- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge: Cambridge University Press.

NOTES

¹ For example, the University of Edinburgh’s ‘Edinburgh Futures Institute’ hosts a project that is explicitly devoted to exploring the market mind hypothesis:

□ <https://www.marketmind.org>

There have also been international conferences devoted to this hypothesis. (Indeed, an earlier version of this paper was presented at one such conference in Edinburgh in May 2022). For a recent book-length defence of this hypothesis, see Schotanus (2023).

² See, especially, Pritchard (2022), where I apply this three-tiered approach to understanding certain kinds of highly theoretical scientific collaborations. See also Pritchard (2025).

³ The *locus classicus* with regard to the extended cognition program is Clark & Chalmers (1998). Note, however, that the overarching aim of Clark & Chalmers (1998) is not just to argue that cognitive processes can be extended in the relevant sense, but also that it follows that the mind can also, therefore, be similarly extended (which is why they refer to their proposal as the *extended mind* thesis). For our purposes, I think we can set to one side whether the mind is extended in the way that Clark and Chalmers propose, and focus our attention on whether cognitive processes can be extended. At the very least, the latter is a plausible precondition for the former, so it makes sense to target our attention on this claim. For a useful discussion of this contrast, and its theoretical import, see Sprevak (2019). For some other key works on extended cognition/mind, see Wheeler (2005), Clark (2008; 2010), and Menary (2010).

⁴ For further discussion of the epistemological ramifications of extended cognition, and thus of the idea of extended epistemology, see Pritchard (2010; 2018a; 2018b), Palermos (2011; 2014a), Carter (2013), Wikforss (2014), Clark (2015), and Jarvis (2015). See also Carter, Kallestrup, Palermos & Pritchard (2014) and the papers collected in Carter, Kallestrup, Palermos & Pritchard (2014) and in Carter, Clark, Kallestrup, Palermos & Pritchard (2018a; 2018b).

⁵ The *locus classicus* with regard to the contemporary discussion of this topic is, of course, Knorr-Cetina (1999).

⁶ In this quotation, MacKenzie is specifically discussing private ‘dark pool’ electronic share trading systems. This is a market where one can trade securities without having to route orders through a public stock exchange. Moreover, MacKenzie’s focus in this quotation is on actor-network economic sociology, though he is trying to bring together Latour-style work in this vein with Bourdieu’s ‘field theory’. See Latour (1999) and Bourdieu (2005).

⁷ Interestingly, Knorr Cetina has herself discussed the relationship between highly theoretical and collaborative scientific cognitive systems to the sociotechnical cognitive system involved in the financial markets. See, for example, Knorr Cetina & Brugger (2002). See also Knorr Cetina & Preda (2004).

⁸ For further discussion of the epistemology of cognitive enhancement in the context of the extended cognition program, see Carter & Pritchard (2019).

⁹ Both constraints are important. The first is important because there is a variety of extended cognition where the extended cognitive processes do not employ information-processing features of the environment. This is known as *embodied cognition*—see Noë (2004), Gallagher (2005), Chemero (2009), Rowlands (2009), and Shapiro (2011). Think, for example, of how using a walking stick might become integrated into one’s way of life. The walking stick has now become part of one’s way of being, but it is not doing any information-processing for you. Note that this is not to claim that embodied cognition is irrelevant to one’s cognitive processes. For example, there is lots of evidence that certain kinds of cognition are improved if one can freely employ one’s hands. For discussion, see Clark (2013), Wheeler (2013), and Pauw *et al* (2014). The second constraint is important because one can also envisage an extended cognition program that applies at the sub-personal level, and so extends beyond beliefs to sub-doxastic states. I take it that this is what Clark is describing in his important work on predictive coding—see Clark (2015b; 2023). See also Clark (2015a), where he explicitly makes the case for an epistemology of extended cognition at the sub-personal level. See also Carter & Rupert (2020).

¹⁰ Naturally, the extended cognition thesis has its detractors. For some of the key recent literature that critiques extended cognition, see Adams & Aizawa (2001; 2008), Rupert (2004; 2009), and Sprevak (2009). There is also a lively debate within the extended cognition literature regarding just what is required to ensure that the target technology has been sufficiently integrated into the cognitive process to count as a proper part of that cognitive process. See Palermos (2014a; 2014b) and Alfano & Skorburg (2016). Notice that the technology being integrated into the cognitive system need not be external to the brain and nervous system of the subject. So-called *neuromedia*, for example, which is inserted into the brain, would count as an extended cognitive process (even though there is a sense in which it is ‘internal’ to the brain of the subject). For further discussion of neuromedia, see Lynch (2014) and Pritchard (2018c).

¹¹ I discuss such transactional memorial systems as socially extended cognitive systems in Pritchard (2025). For some important discussions of transactional memory systems in the context of group and extended cognition, see Wegner *et al*

(1985), Wegner (1986), Sutton (2008), Sutton *et al* (2008), Huebner (2016), Michaelian & Sutton (2018), Michaelian & Arango-Muñoz (2018), and Palermos (2020).

¹² I've argued for this modest version of virtue epistemology in a number of places, but see, especially, Pritchard, Haddock & Millar (2010, ch. 3) and Pritchard (2012). For some of the key defences of a stronger 'robust' version of virtue epistemology whereby knowledge involves cognitive success that is primarily attributable to one's cognitive agency, see Sosa (1991; 2007; 2009; 2015), Zagzebski (1995), and Greco (2010).

¹³ For more details about what an extended virtue epistemology might look like, see Pritchard (2018*b*). Even though virtue epistemology can accommodate extended cognitive processes, it is undoubtedly the case that many contemporary epistemologists implicitly suppose that all cognition (of a kind that's relevant to knowledge at any rate) is unextended. Sometimes this commitment is made explicit—see, for example, Goldman (1986, 51).

¹⁴ For some helpful discussions of the core issues regarding group knowledge, see Gilbert (2004), Tuomela (2004), Bird (2017; 2019), Tollefsen (2019), and Lackey (2021).

¹⁵ I expand on this point about how best to understand scientific collaborative teams in Pritchard (2022). Ironically, the point I make in this regard is very similar to that made by Wray (2018) in his *critique* of thinking of scientific research teams as involving extended cognition. As I point out, Wray's mistake is to fail to recognize that extended knowledge is still individual knowledge (which means that his critique is only actually effective against distributed scientific cognition/knowledge).

¹⁶ Of course, there are some rare cases—like the kinds of experiments conducted at CERN that involve multinational funding agencies and years of planning—where the financial stakes are huge. But these are not typical for scientific collaborations in the way that they are when it comes to those professionals operating in the financial markets. Note too that financial traders are typically directly financially rewarded for their successful trades with commissions and bonuses. In contrast, this is rare in the case of scientific collaboration (though it does sometimes occur in private sector scientific research).

¹⁷ An earlier version of this paper was presented at the 'Market Mind Hypothesis' international conference held in Edinburgh in May 2022. A subsequent version of the paper was presented in April 2024 as part of the 'Webinars on Digital Economy and Artificial Systems' series, hosted by IDEAS ('Institute for Digital Economy and Artificial Systems'), which is co-run by Moscow State University and Xiamen University. I am grateful to the audiences on both occasions for their feedback and, especially, the organisers of these events, which were, respectively, Patrick Schotanus and Mirko Farina. Thanks also to an anonymous reviewer for detailed comments on an earlier version of this paper.