Forthcoming in Allori, V. (ed.) *Quantum Mechanics and Fundamentality*

All flash, no substance?

Elizabeth Miller

Abstract: The GRW dynamics propose a novel, relevantly "observer"-independent replacement for orthodox "measurement"-induced collapse. Yet the tails problem shows that this dynamical innovation is not enough: a principled alternative to the orthodox account demands some corresponding ontological advancement as well. In fact, there are three rival fundamental ontologies on offer for the GRW dynamics. Debate about the relative merits of these candidates is a microcosm of broader disagreement about the role of ontology in our physical theorizing. According to imprimitivists, the GRW dynamics directly describe (only) some (element's) undulation in an unfamiliar high-dimensional physical field. Primitivists resist this GRW0 proposal on the grounds that it fails to secure comprehensible contact with our data about macroscopic objects in ordinary low-dimensional space-time. They expect an adequate fundamental ontology to include at least *some* spatiotemporally localized entities—intuitively, concrete constituents of our familiar macroscopic landscape. The most compelling case goes by way of *distributional basing*: minimally, primitivists expect a theory's predictions immediately about spatiotemporal *distributions* of fundamental entities to provide a supervenience base for data about configurations of macroscopic objects. But while the background intuition is familiar, the distributional model is surprisingly subtle. Lack of clarity about its details generates serious confusion for both sides of our debate.

While textbook quantum theory provides a remarkably successful recipe for predicting the outcomes of experiments, its standard interpretation rests on gerrymandered dynamics in which "observers" making "measurements" play a starring role. This is the price orthodox quantum mechanics pays to solve—or, more accurately, evade—the measurement problem (§1).[1] GRW's dynamical innovation promises a more principled solution, but a reincarnation of the measurement problem persists for GRW (§2). On one view, the tails problem proves that our initial innovation is not enough: we need some ontology to go along with our new dynamics.[2] In fact, there are three candidates on offer: GRW0, GRWm, and GRWf propose three rival fundamental ontologies for the GRW dynamics. Debate about

---

[1] I am deeply grateful to Ned Hall for encouraging me to develop these ideas and for sharing many invaluable insights in conversation about them. For further discussion, I am grateful to David Albert and Nina Emery, as well as to the editor of and referees for this volume. My thanks also to David Baker, Michael Della Rocca, Heather Demarest, Barry Loewer, Tim Maudlin, Mark Maxwell, Michaela McSweeney, Alyssa Ney, Jill North, Zeynep Soysal, Scott Sturgeon, Elanor Taylor, and Alastair Wilson; to members of my philosophy of physics seminars at Brown and Yale; and to audiences at the Metro Area Philosophy of Science Association (at NYU), the University of Vermont, and USC.

[2] For Maudlin, an ontology is "a collection of items taken to be physically real" by a theory, which can include familiar macroscopic objects featuring explicitly in our data but only implicitly in the theory's own predictions (2013, 143). A theory is stated "in terms of" some *subset* of its ontology so understood. I follow, among others, Allori (2013) and Emery (2017) in describing this subset as the theory's *fundamental* ontology; for simplicity, though, I drop the 'fundamental' where possible. Some instead use the label 'primitive ontology' for this subset, but since others reserve that for fundamental elements that *also* meet some further condition, I avoid this term here; cf. Dürr, et al. (1992), Maudlin (2013). Still my own label 'primitivism' alludes to discussions of primitive ontology—as well as to discussions of "spatial primitivism" elsewhere in metaphysics; cf. Chalmers (2012) 325ff.

their relative merits is a microcosm of broader disagreement about the role of ontology in our physical theorizing.

One point of dispute is whether an adequate fundamental ontology must or should include some "local beables"—physical entities localized "at definite places and times in the real world" (§3).[3] Local beables play a starring role in our data. Since our antecedent vision of the world features familiar macroscopic objects in space-time, any comprehensible theory about *our* world will be at least *indirectly* about these. Minimally, predictions *directly* about the features and behaviors of its fundamental ontology will provide a supervenience base for facts about "the arrangement of things in ordinary 3-dimensional space".[4] But Bell imposes a further continuity constraint: we must recognize, implicit in these predictions, a familiar "image of our physical world". Since this image features objects in space-time, *primitivists* expect our base itself to include at least some fundamental local beables (or *flobs*)—intuitively, concrete *constituents* of our macroscopic landscape. On the other side are *imprimitivists*, who maintain that a theory can do just as well, if not better, without any flobs at all.

In our microcosm, imprimitivists champion GRW0, which takes the GRW dynamics to directly describe (only) some (element's) undulation in a high-dimensional physical field. Primitivists resist GRW0 on the grounds that it fails to secure adequate contact with our data (§4). The most compelling resistance goes by way of prior commitment to *distributional basing*, which—in Bell's words—links a "piece of matter" from our data to some "galaxy" of constituents within our theory.[5] Interestingly, though, recent arguments for imprimitivism begin from the premise that distributional basing is a non-starter: even for primitivists, (i) a minimal supervenience base must include something *besides* an occurrent distribution of flobs). Imprimitivists focus on showing that (ii) once we add requisite functional or dynamical factors to the base, distributional inputs are superfluous: GRW0 can do just as well without any flobs.

---

[3] Bell (1987) 45.
[4] Bell (1987) 44.
[5] Bell (1987) 45.

Such a focus obscures two concerns. One is that imprimitivists may prove too much: if (i) and (ii) are true, we risk undermining some motivation (reviewed in §§1-2) for turning to ontology in the first place. The other, and my primary, concern is that imprimitivists lack an adequate defense of (i). In advancing (i), they conflate distributional supervenience with an implausible claim of geometric sufficiency. But imprimitivists are not alone in their mistake. A lack of clarity about the details of distributional basing generates serious confusion on *both* sides of our dispute (§5). When it comes to recovering Bell's "image" of our data, it is not enough for primitivists to know it when they see it: before we can hope for any further progress, we need careful investigation into their—familiar yet surprisingly subtle—model of contact between macroscopic objects and "galaxies" of localized constituents. My aim is to highlight the distributional model's subtlety—and, in light of that, to suggest a path forward in disagreement about quantum ontology.

1 The GRW dynamics

Start with Schrödinger's hapless cat, "penned up in a steel chamber":

> [I]n a Geiger counter there is a tiny bit of radioactive substance, *so* small, that *perhaps* in the course of one hour one of the atoms decays…; if it happens, the counter tube discharges and through a relay releases a hammer which shatters a small flask of hydrocyanic acid. If one has left this entire system to itself for an hour, one would say that the cat still lives *if* meanwhile no atom has decayed.[6]

Schrödinger seals the chamber at t_0, expecting one of two outcomes after a tense hour. The textbook recipe promises a cat-flask system with one of two wavefunctions, $|\psi_{\text{dead}}\rangle$ or $|\psi_{\text{alive}}\rangle$, when he opens the door at t. In the interim, it predicts something even more "diabolical". Between t_0 and t, the theory ascribes neither $|\psi_{\text{dead}}\rangle$ nor $|\psi_{\text{alive}}\rangle$ to Schrödinger's system; instead, it evolves the system's initial state to $\frac{1}{\sqrt{2}}|\psi_{\text{dead}}\rangle + \frac{1}{\sqrt{2}}|\psi_{\text{alive}}\rangle$ just before t. On one hand, it is not clear what $\frac{1}{\sqrt{2}}|\psi_{\text{dead}}\rangle + \frac{1}{\sqrt{2}}|\psi_{\text{alive}}\rangle$ amounts to physically. On the other, it seems clear enough what it does *not*: anything we actually find in our world. Schrödinger himself suggests "the living and the dead cat (pardon the expression) mixed or smeared out in equal

---

[6] Trimmer (1980) 328, translating Schrödinger (1935).

parts", and, in essence, the orthodox interpretation agrees: Schrödinger's *mathematical* evolution indicates a *physical* "indeterminacy" only "resolved" by observation.

More generally, since the textbook recipe frequently yields pre-measurement functions in superpositions of relevant eigenstates, its standard interpretation says that systems *themselves* frequently evolve away from familiar physical states. Nevertheless, it does not follow that we should expect to *observe* any unfamiliar happenings. Given pre-measurement state $\frac{1}{\sqrt{2}}|\psi_{\text{dead}}\rangle + \frac{1}{\sqrt{2}}|\psi_{\text{alive}}\rangle$, the recipe promises *one* of two familiar eigenstates, each with prior probability $|\frac{1}{\sqrt{2}}|^2=\frac{1}{2}$, at *t*. Orthodox quantum mechanics tailors its physical dynamics to underwrite this guarantee: when Schrödinger opens the chamber, his system's character transforms into either a dead state or a live one. By design, systems evolve away from familiar states only when, and indeed *because*, *we* are not looking—luckily, though, *our* data only concerns what the world is like when we *are*.

An alternative from Ghirardi, Rimini, and Weber targets this "measurement"-induced collapse, proposing a novel stochastic mechanism in its stead.[7] The resulting, *unified* GRW dynamics yield, as a *theorem*, something close to the orthodox story, thereby cohering with the same empirical data. Crucially, though, there is nothing special about observers or measurements *per se*. Instead, the GRW proposal starts with a division of the world into units (*particles*), each assigned some probability per unit time of undergoing a *hit* event. A hit on a particle transforms the wavefunction representing its position: roughly, we can think of a hit centered around some point in space-time as prompting our particle to jump there. A larger system's probability of undergoing a hit-induced transformation is then a function of the number of particles interacting within it: a hit on one particle transforms the wavefunction of the whole.

Schrödinger's possible outcomes differ with respect to the positions of macroscopically many particles: a system in state $|\psi_{\text{dead}}\rangle$ has flask particles strewn across region R_{dead} of the chamber floor, rather than arranged in R_{alive}. While any single particle will undergo a hit, on average, only once every 10^8 years, we can expect a hit *somewhere* in our flask—on the order of 10^{23} particles—every 10^{-8} seconds.[8] The

---

[7] Girardhi, Rimini, and Weber (1986); cf. Albert (1992) 80-116, Albert and Loewer (1996), and Lewis (2006).

[8] Maudlin (2011) 226-8.

matched coefficients in $\frac{1}{\sqrt{2}}|\psi_{\text{dead}}\rangle + \frac{1}{\sqrt{2}}|\psi_{\text{alive}}\rangle$ indicate that the next hit is just as likely to be centered in R_{dead} as in R_{alive}, amplifying one of two components. Wavefunctions still spread out between hits, but Schrödinger's system is large enough—and so hits frequent enough—that its spreadings are imperceptibly brief. Thanks to our first hit, one of $|\psi_{\text{dead}}\rangle$ and $|\psi_{\text{alive}}\rangle$ stands amplified at *t*, and subsequent hits are highly likely to reiterate that choice.

2 The tails problem and fundamental ontology

GRW's insight is that anything recognizable as an orthodox "measurement" shares key physical characteristics with Schrödinger's case. Whenever these characteristics are present, stochastic hits naturally rein in smeared-out wavefunctions. Importantly, though, these hits do not transform such wavefunctions into *strict* eigenstates. Mathematically, a hit on a particle multiplies its wavefunction by a Gaussian, amplifying some component over the rest. But thanks to non-zero "tails" on our multiplying curve, such amplification does not eliminate the others. In Schrödinger's case, a hit takes $\frac{1}{\sqrt{2}}|\psi_{\text{dead}}\rangle + \frac{1}{\sqrt{2}}|\psi_{\text{alive}}\rangle$ to some function retaining *multiple* components, say $\sqrt{0.999}|\psi_{\text{dead}}\rangle + \sqrt{0.001}|\psi_{\text{alive}}\rangle$, at *t*. If our pre-hit function was troubling because it was not a plain eigenstate, this *tailed* post-hit function should be just as bad. The natural fix, curtailing the tails, is not a live option: our measurement problem is back with a vengeance, just reincarnated in a subtler form.[9]

Strictly, the tails problem proves only that GRW's initial innovation is not enough. We find our tailed function troubling because we think it amounts to neither a definitely dead nor definitely live cat. Yet this conclusion rests on a further assumption about the relationship between representation and reality: a system has a definite "value" for some given "observable" *only if* its wavefunction is in a strict eigenstate of the mathematical operator associated with that observable. We might avert the tails problem by revising this link, providing a more sophisticated way of extracting physical outcomes from formal representations. Specifically, we need a sort of sophistication that counts $\sqrt{0.999}|\psi_{\text{dead}}\rangle + \sqrt{0.001}|\psi_{\text{alive}}\rangle$ as *close*

---

[9] In brief, "tailless" collapses threaten the conservation of energy, since they reduce uncertainty about position without any compensatory increase in uncertainty about (velocity and so) energy; cf. Albert (1992) 78.

enough to $|\psi_{\text{dead}}\rangle$. One option is stipulation: we could propose a mapping between *almost* eigenstates and physical outcomes.[10] To provide a *principled* alternative to the orthodox account, however, GRW needs to *motivate* our close enough verdict.

A different approach starts by getting serious about quantum *ontology*.[11] Maybe, as Schrödinger suggests, mathematical evolution to $\frac{1}{\sqrt{2}}|\psi_{\text{dead}}\rangle + \frac{1}{\sqrt{2}}|\psi_{\text{alive}}\rangle$ represents *some worldly* spreading out or other, but what are the *entities doing* this spreading, what sort of physical *processes* are involved, and how does *any* of this *bear on* our data? The eigenstate-eigenvalue link—like any stipulative revision to it—attempts to map directly from mathematical representations to the recognizable macroscopic outcomes at issue in our data. In doing so, it allegedly skips a crucial step, obscuring the mediating role of ontology in our theorizing. Compare a classical case: a candidate theory assigns formal state C to the current contents of my office. It also issues predictions about those contents later, assigning mathematically distinct C* to them at *t*. I leave my office, planning to return at *t*. Should I expect to find the furniture rearranged when I do?

To answer, it is not enough to know that our theory's predictions have *some* bearing on the contents of my office. We need to know what its mathematical states *directly* represent, and how these underwriting physical happenings relate to our observable outcomes. In fact, C depicts an *arrangement of particles* in my office. C* depicts another, differing with respect to, say, the relative positions of three constituents. Equipped with this ontological insight, we can opine: I should not expect anything new when I return. Our distinct mathematical states represent different physical happenings, but so slight a difference at the level of particles will not show up in our macroscopic data. In this sense, C* is *close enough* to C: we can recognize Bell's image of my desk in both. To dissolve the tails problem, we need to issue the same *sort* of verdict: $\sqrt{0.999}|\psi_{\text{dead}}\rangle + \sqrt{0.001}|\psi_{\text{alive}}\rangle$ is close enough to $|\psi_{\text{dead}}\rangle$, since they depict happenings with the same (unfortunate) import for our data. For this, we need some new and improved—or at least *some*, reasonably clear—ontology for the GRW dynamics.

---

[10] Albert and Loewer (1996) and Lewis (2006) discuss this strategy.

[11] I learned the term 'serious ontology' from John Heil; cf. Heil (2012).

3 GRW0 and primitivism

According to Bell, orthodox disregard for ontology proves especially troubling once we note that the wavefunction for my desk, Schrödinger's flask, or any other macroscopic system "lives in a much bigger space" than the one we ordinarily take ourselves to inhabit.[12] Perhaps, then, any complete interpretive backstory for the textbook recipe will include some *physically* high-dimensional ingredient(s). Regardless, Bell and likeminded primitivists expect it to include at least *some* flobs in "ordinary 3-dimensional space" as well. Since science promises us deeper, even revisionary, insight into our world, characters in the scientific story may be antecedently unfamiliar to us. Still, they cannot be *so* unfamiliar that we fail to recognize the story as about *our* world—Bell's familiar "arrangement of things"—at all.

GRW0 proposes a thoroughly non-local fundamental ontology for the GRW dynamics. On the GRW0 proposal, these dynamics depict (only) some (element's) undulation in a high-dimensional field, a physical counterpart to the mathematical space in which the universal wavefunction resides. Ordinary space-time is not a straightforward subspace of this realm: concrete happenings are aspects of, projections from, or patterns in some fundamental affairs "outside" of our familiar milieu. GRW0 furnishes a complete supervenience basis for our data: we can find *some* mapping between various high-dimensional affairs and familiar macroscopic outcomes.[13] According to primitivists, however, this is not enough: GRW0 does not afford the right *sort* of contact with our data: at the very least, some alternative including flobs can—and in fact does—do better.

Before we ask whether $\sqrt{0.999}|\psi_{\text{dead}}\rangle + \sqrt{0.001}|\psi_{\text{alive}}\rangle$ and $|\psi_{\text{dead}}\rangle$ amount to the *same* outcome, we need to recognize both as concerning fundamental happenings that bear, *somehow* or other, on outcomes *here* in the lab. But why should *these* particular undulations tell us about the contents of Schrödinger's chamber *here*, rather than about, say, the furniture in my office, or some distant configuration of elm trees? Not because the undulations are *here* too: GRW0 has *no* fundamental happenings here—or, indeed, *anywhere* in

---

[12] Bell (1987) 44.

[13] It is "informationally" complete in the sense of Maudlin (2007a).

space-time. Minimally, we need to say *which* high-dimensional happenings bear on *which* pieces of data, but we started down the road of ontology looking for something better than stipulation.

Rather than turning back, primitivists urge us to keep going. GRW0 shows that not *just any* ontological ingredients will do: we need at least some flobs in the mix. But not just any localized additions will do either. Consider (fictional) GRW0$^+$, which supplements GRW0's high-dimensional happenings with arbitrary flob *o*. Like GRW0, GRW0$^+$ furnishes at least one supervenience basis for our data. Unlike GRW0, it furnishes at least one basis *also* comprising a localized ingredient(s). Yet *o*'s addition does not secure any relevant advantage over GRW0. Intuitively, *o* is a mere dangler: while all relevant facts supervene on a base including it, they do so *only* because they *already* supervene on GRW0's high-dimensional happenings alone. Primitivists expect their flobs to play a more essential, distinctive role in securing contact with our data.

Whatever else it may comprise, that data includes facts about the gross configuration of macroscopic objects within, or the gross distribution of matter across, space-time. Primitivists expect such facts to be fixed *entirely* by their theory's predictions just about some underwriting flobs. Our theory may bring other—perhaps non-local, even non-spatiotemporal—ingredients along for the ride. Still, these play a comparatively *indirect* role in "determining" our data: in Maudlin's terms, their effects are "screened off" by our "[p]rimary" mediating flobs.[14] Intuitively, other fundamental ingredients can causally or physically influence these flobs, but they alone *directly* ground our data.

For this intuitive distinction to have any bite, primitivists must circumscribe those predictions relevantly *about* their candidate flobs. Otherwise, GRW0$^+$ could count as screening off *all* our data by ascribing some complicated but nominally "local" state—merely *represented*, in our theory, by the *universal* wavefunction—to *o* alone. A "Democritean" fix blocks this move: the relevant theoretical predictions about primary local beables ascribe (only) spatiotemporal *positions* to otherwise *qualitatively indistinguishable*

---

[14] According to Maudlin, "if we imagine keeping the behavior of the Primary Ontology fixed but altering the behavior of the Secondary Ontology, the data would remain the same" (2013, 144).

elements.[15] To make contact with familiar facts about my office, our theory need not explicitly mention desks. Still, it should describe at least *something* here in $\mathrm{R}_{\mathrm{office}}$, *where* we antecedently take my desk to be. Whatever *features* this something may or may not turn out to have on our scientific story, it at least should be distinguishable from what the theory describes over *there* in $\mathrm{R}_{\mathrm{empty}}$, where we find no macroscopic objects at all.

Minimally, then, our theory should depict some elements in $\mathrm{R}_{\mathrm{office}}$ without corresponding counterparts in $\mathrm{R}_{\mathrm{empty}}$. We can extend this reasoning to smaller scales. Consider the edge of my desk here: some materially occupied subregions of $\mathrm{R}_{\mathrm{office}}$ abut comparatively vacant neighbors. This contrast at the macroscopic level should show up, somehow or other, as some difference in corresponding contents at the level of fundamental ontology. What results is a primitivist paradigm of *distributional basing*. Rather than stipulating some link between *this* data about the macroscopic configuration of my office and *those* discontinuous fundamental happenings in GRW0, primitivists hope to rely on some antecedently familiar continuity between configurations of macroscopic objects and underwriting distributions of concrete constituents. Primitivists can disagree amongst themselves about the nature of this continuity, and so about further metaphysical or epistemological constraints on the base.[16] Minimally, though, they expect data about macroscopic configurations to supervene on the theory's distribution of flobs in space-time.

4 GRWm and GRWf

Primitivists are spoiled for choice. There are two alternatives to GRW0, each proposing some flobs for the GRW dynamics. But this choice also generates controversy, since even for primitivists, not just any localized additions guarantee a relevant advantage over GRW0. Maudlin argues that GRWm does not secure the right sort of contact with our data: it does not provide even a minimal distributional supervenience base for "the arrangement of things in ordinary 3-dimensional space". Behind both Maudlin's original argument and Albert's criticism of it is common confusion about what exactly the

---

[15] Oppenheim and Putnam (1958) 16 n. 18; cf. Ney (2013). Alternatively, we can think of these predictions as ascribing fundamental states of "occupation" to points or small regions of space-time itself.

[16] For discussion of candidate constraints, Allori (2013), Ney (2013), and Emery (2017).

distributional model demands.

According to GWRm, reality is not ultimately particulate. Instead, each "particle" from our initial presentation of the dynamics gets associated with some portion of continuous mass density. Our dynamics now depict evolving distributions of mass density in ordinary space over time: smooth temporal evolution of a particle's wave function corresponds to a spreading out of mass density. What we originally described as a particle undergoing a hit centered within a region involves consolidation of associated mass density there. As before, a hit on one part of the systems transforms the wavefunction for the whole: Schrödinger's system in state $\sqrt{0.999}|\psi_{\text{dead}}\rangle + \sqrt{0.001}|\psi_{\text{alive}}\rangle$ has much more of its mass density in R_{dead} than in R_{alive}. Given our transformed coefficients, subsequent hits are likely to consolidate mass density in R_{dead}, leaving R_{alive}'s shadow to spread and thin. What about the lingering $|\psi_{\text{alive}}\rangle$? Our Gaussian tails guarantee some thin smear of mass density in R_{alive}, but, allegedly, that smear does not show up at the level of our data: our tailed function is close enough to $|\psi_{\text{dead}}\rangle$ to yield a familiar outcome.

Maudlin illustrates his doubts about this last step with the case of a pointer on a dial. If we like, we can imagine the pointer's possible orientations, A and B, representing dead and alive states of Schrödinger's system, but for now just consider the pointer itself: its wavefunction spreads from an initial ready state, splitting briefly into $\frac{1}{\sqrt{2}}|\text{in } R_A\rangle + \frac{1}{\sqrt{2}}|\text{in } R_B\rangle$, before yielding $\sqrt{0.999}|\text{in } R_A\rangle + \sqrt{0.001}|\text{in } R_B\rangle$ at *t*. Then:

> *On the assumption that…*[*R_B's*] *small (and ever shrinking) mass density can be safely neglected*, the post-hit state [$\sqrt{0.999}|\text{in } R_A\rangle + \sqrt{0.001}|\text{in } R_B\rangle$] would be as satisfactory in accounting for our beliefs about the outcome as it is in the case of predictive certainty [$|\text{in } R_A\rangle$].
>
> But on what basis, exactly, can the small mass density be neglected? After all, that mass density is *something*, and it has the same shape and…behavior and dispositions to behave as it would have had if the hit had left it with the lion's share of mass density.[17]

Our data includes one pointer at A, with no macroscopic counterpart at B, but GRWm depicts two portions of mass density.

According to Maudlin, these portions match in all relevant respects, even at *t*. Indeed, the actual distribution of mass density across *both* R_A and R_B matches, in all such respects, what we would find in the

---

[17] Maudlin (2010) 135; his emphasis.

case of $\sqrt{0.001}|\text{in } R_A\rangle + \sqrt{0.999}|\text{in } R_B\rangle$. Of course, R_B's actual mass density is supposed to be too *thin* to amount to a macroscopic object, but on what "basis" can we discount it? In principle, we might introduce some concentration threshold: only R_A contains any recognizable distribution of portions *with* thickness greater than ρ at t. If we stipulate some value for ρ, however, we risk undermining our motivation for introducing flobs in the first place: GRW0 can just as well stipulate some mapping between its discontinuous fundamental happenings and macroscopic outcomes.

Albert and likeminded imprimitivists share Maudlin's risk assessment, but they embrace this consequence for GRW0. They take Maudlin's case to point to a deeper problem with primitivism itself. According to Albert, familiar objects are distinguished, in part, by characteristic behaviors and dispositions. As a result, everyone, even primitivists, must accept: (i) an adequate basis for our data will specify *more* than the occurrent distribution of fundamental ontology. Albert takes Maudlin's reference to "behavior and dispositions to behave" to mark his own *explicit* acceptance of (i). On Albert's diagnosis, a confused fixation with distributional factors gets in the way: Maudlin *implicitly* expects "whatever is *shaped* like" a pointer to *be* a pointer, and so mistakenly assumes that his similarly shaped portions of mass density have the same dispositions at t.[18] Thanks to the GRW dynamics, only Maudlin's thicker portion is actually disposed to consolidate—its thinner counterpart is likely to spread and fade.

According to Albert, the same sort of confusion lying behind Maudlin's error in this case is also what motivates primitivists' attachment to fundamental local beables in the first place. Primitivists fail to appreciate (i), because they overemphasize distributional considerations. If we take (i) to heart, though, we find that GRW0 can do just as well as its rivals. That is: (ii) once we have functional information in the base, distributional facts are, at best, superfluous.[19] It is precisely because of (ii), however, that primitivists should not, and Maudlin plausibly *does* not, grant (i) to begin with: primitivists hope to secure an advantage over GRW0 by providing a *distributional* basis.

As Maudlin observes, even if we need further factors to distinguish those distributions that

---

[18] Albert (2015) 151.

[19] Cf. Rubenstein (forth.).

amount to genuine *pointers* from mere pointer-*like* objects, such functionalism is beside the point in his case: pointer or not, GRWm depicts "*something*" macroscopic in R_B. Muddying the waters throughout Albert and Maudlin's exchange are their differing uses of the term 'disposition'. Albert and his fellow imprimitivists are concerned with causal or functional dispositions. Yet Maudlin instead seems to have in mind the "arrangement, order; [or] relative position of the parts or elements of a whole".[20] In comparing "dispositions" in *his* sense, then, Maudlin is quite plausibly *affirming* his commitment to distributional basing. After all, he rejects GRWm because he thinks it fails to mark any requisite difference in the arrangements of flobs across R_A and R_B.

Maudlin's own response is to swap GRWm's continuous mass density for GRWf's discrete event ontology. The physical correlate of a hit centered around some point in space-time is now a "flash" there. Our matched weights in $\frac{1}{\sqrt{2}}|\text{in } R_A\rangle + \frac{1}{\sqrt{2}}|\text{in } R_B\rangle$ signify that R_A and R_B are equally likely to host the next flash. In fact, a hit in R_A takes us to $\sqrt{0.999}|\text{in } R_A\rangle + \sqrt{0.001}|\text{in } R_B\rangle$ at *t*. Our transformed coefficients indicate that the next flash is highly likely (probability .999) to be in R_A as well. The GRW dynamics ensure that further hits are likely to reiterate this choice. Crucially, no shadows haunt GRWf: where GRWm looks at $\sqrt{0.999}|\text{in } R_A\rangle + \sqrt{0.001}|\text{in } R_B\rangle$ and posits a *thin but sure* smear of mass density in R_B, GRWf posits only some *probability* of a flash—and then a cascade of flashes—there. But probabilities of flashes are neither flashes nor pointers nor shadows of these. Our tailed function is close enough to an eigenstate: $\sqrt{0.999}|\text{in } R_A\rangle + \sqrt{0.001}|\text{in } R_B\rangle$ depicts a galaxy in R_A—and none in R_B.

5 Massy gluts and flashy gaps

Since GRWf has "no 'low density' or 'low intensity' flash-sequences to be ignored or discounted or argued away", Maudlin takes it to avoid the "interpretive problems" that stymie GRWm.[21] Yet GRWf's comparative sparseness also ensures that, sometimes, GRWf has mere probabilities where

---

[20] *OED* online (2020); thanks to Maudlin (p.c.) for confirming. Compare Maudlin's characterization of the Humean mosaic as "determined by nothing more than the values of the individual pixels plus their spatial disposition relative to one another…" (2007b, 51), and his claim that (for our primitivists) "the disposition of the local ontology screens off the quantum state from the data" (2013, 148).

[21] Maudlin (2010) 138-9.

GRWm has sure *non*-shadowy mass. On the GRWf account, wavefunctions underdetermine occurrent distributions of flashes. Given $\sqrt{0.999}|\text{in } R_A\rangle + \sqrt{0.001}|\text{in } R_B\rangle$ at *t*, we can expect another flash in R_A *very* shortly. Still, its actual occurrence is determined by GRW's stochastic dynamics. In the interim, however brief, we have *no* flashes at all in R_A. Indeed, if flashes are truly instantaneous, then GRWf promises many *entirely* empty moments, amid some *barely* flashy accompaniment. GRWf virtually guarantees some unfolding cascade in R_A, but it does not promise a multi-flash galaxy *at* any moment. Where GRWm has some sure, thick mass density in R_A, GRWf offers—*if* we are lucky—a single flash.

Flashy gaps generate a challenge for GRWf analogous to the challenge shadowy mass density generates for GRWm. One potential objection to GRWm starts from $\frac{1}{\sqrt{2}}|\text{in } R_A\rangle + \frac{1}{\sqrt{2}}|\text{in } R_B\rangle$, before our initial hit in Maudlin's case: we have two exactly matched portions of mass density at time *t*-. But this cannot be a basis on which to prefer GRWf, because it is in the same boat: $\frac{1}{\sqrt{2}}|\text{in } R_A\rangle + \frac{1}{\sqrt{2}}|\text{in } R_B\rangle$ depicts two equally *vacant* distributions of flashes at *t*-. In fact, the situation for GRWf is worse: $\sqrt{0.999}|\text{in } R_A\rangle + \sqrt{0.001}|\text{in } R_B\rangle$ at *t* likely depicts two vacant distributions *after* our hit as well.

Since Maudlin takes GRWf to avoid the "interpretive problems" that plague GRWm, the image of our data need not comprise some single pointer-shaped distribution of flobs at *t*. Our data concerns humanly detectable happenings, so perhaps what we need is a recognizable distribution across some suitably extended temporal interval. One option is to link R_A's momentary state to some broader distribution of flobs. Another is to deny that there are any momentary macroscopic states at all: our data concerns configurations spread over space *and* time. Either way, once we shift our attention to an extended interval, we can find a difference in the decorations of R_A and R_B. But we *also* find at least *some* difference between GRWm's portions of mass density: though $\frac{1}{\sqrt{2}}|\text{in } R_A\rangle + \frac{1}{\sqrt{2}}|\text{in } R_B\rangle$ depicts a *momentary* match, the contents of R_B soon spread and fade.

For Maudlin, of course, this is not the right *sort* of difference to underwrite our macroscopic data: thick or thin, GRWm still predicts "*something*" in R_B, even after our hit. His remark might seem to suggest that GRWf, unlike GRWm, satisfies the following condition: a region of space is macroscopically vacant *only if* there are *no* flobs in it during some relevant interval. Even in the classical case, though, we learn that "empty" space is not devoid of occupants. In a large enough region, a single flash, or even some sparse scattering of flashes, does not suffice for anything macroscopic at all. Instead, perhaps the claim is that for any relevant region of space-time R, R is macroscopically vacant *only if* there is no recognizably *shaped* distribution of flobs there. Even after our first hit, GRWm retains two recognizably pointer-*shaped* distributions, only one of which is thick enough to be a genuine macroscopic object. According to Maudlin, GRWm needs some basis on which to discount some recognizably *shaped* distributions as macroscopically inconsequential. In contrast, perhaps any recognizably pointer-*shaped* galaxy of flashes amounts to *something* macroscopic for GRWf.

While our choice between continuous and discrete ontologies can obscure this point, GRWf's appearance of an advantage arises from ambiguity in such talk of shape. Trace the outline of Maudlin's distribution of mass density in R_A over some interval. Now imagine filling in this outline with flashes. How many flashes does it take to yield a recognizably pointer-*shaped* distribution? Perhaps *any* distribution "consistent" with our outline—even just one flash within this boundary—counts as pointer-shaped in *some* minimal sense. But not every distribution that counts as pointer-shaped in *this* sense does—or should—amount to a genuine macroscopic object. Perhaps Maudlin has some more robust notion in mind: a *relevantly* pointer-shaped distribution of flashes is not only minimally consistent with our outline but *also* sufficiently clustered to show up in our data. But then GRWf and GRWm are in the same boat: both must distinguish some special, sufficiently *concentrated* candidates from among all *minimally* pointer-shaped ones.

Perhaps this result offers Albert a way of shoring up his case for (i): GRWf and GRWm are in the same boat, because neither has any advantage over GRW0. On Albert's diagnosis, implicit commitment to geometric sufficiency motivates primitivists' attachment to local beables. As Maudlin himself points out, however, GRWm violates this condition. Since GRWf is in the same boat, it does too: On what *basis* can *either* one discount insufficiently concentrated distributions? According to Albert, the right response is to give up on the expectation of distributional basing: even primitivists must agree that we need *more* than the occurrent distribution of fundamental ontology in the base. But then GRW0 can do just as well without any localized fundamental ontology at all.

Alternatively, maybe GRWf and GRWm are in the same boat because *both* offer a relevant advantage over GRW0. According to Maudlin, GRWf does better than GRW0. If GRWm is in the same boat, then—despite his own suspicions to the contrary—it does too. To defend this alternative, primitivists must separate distributional basing from any implausible claim of geometric sufficiency. Recall our primitivists' motivation for introducing fundamental local beables in the first place: roughly, they hope to link macroscopic configurations to arrangements of localized constituents (§3). Our data includes a single macroscopic pointer, at A, over some relevant interval. The tip of our pointer extends through small region R_{A1}, while some neighbor R is (comparatively) empty. Primitivists expect our theory to predict at least some flobs in R_{A1} without any counterparts in R. Still, more flobs in R_{A1} is consistent with *some*—sparser—flobs in R.

In Maudlin's case, though, GRWm does mark *this* sort of distributional difference, even between his similarly *shaped* portions of mass density. To make the parallel with GRWf vivid, imagine GRWm's mass density packaged in thin layers: mass density gets "painted on" a region in a series of discrete strokes. We can have *some* difference in distributions—whether massy "strokes" or flashes—across R_A and R_B, without any difference in "shape" *per se*. Neither GRWf nor GRWm counts every *minimally* pointer-shaped distribution as a macroscopic object—but nor should they. The key question for primitivists is whether this result demotes both to the level of stipulation or, alternatively, whether both preserve continuity that secures their alleged advantage over GRW0. Primitivists expect at least minimal ("global") supervenience between our data and spatiotemporal distributions of flobs.[22] Intuitively, they expect a special link between our data about one part of the world and some distribution in that part. To make progress, though, primitivists need to clarify what exactly, beyond minimal supervenience, their model demands.

GRWf's sparseness underscores this need. For whatever exactly distributional continuity comes to, the distribution of flobs *just* within a region does not, in general, metaphysically suffice for its

---

[22] Cf. Kim (1987). A referee asks how distributional supervenience relates to Humean supervenience (HS). There is much more to say, but *prima facie*, primitivists can remain neutral on HS, which claims supervenience for *all facts* about the world, including causal and nomological facts, not just facts about macroscopic configurations. Cf. Maudlin (2007b), Miller (2018).

macroscopic state on the GRWf account. For small regions, this result is unsurprising: it is just the spatial analogue of our earlier point about temporal sparseness. While there is a pointer in R_A through some interval, subregion R_{A*} may well contain *no* flashes at all. So, if R_{A*} counts as macroscopically occupied at all, it is because of some distribution extending beyond its bounds. But now consider larger R_A: this hosts the broader distribution D that, intuitively, makes up our pointer at A. Perhaps within *our* actual world, any flashy duplicate of D amounts to something macroscopic. Yet consider another (near) world *w*, which hosts some duplicate of D in much flashier surroundings. Does this—indeed, *should* this—duplicate amount to *any* macroscopic object? Surely, it depends: not if even "empty" space is densely decorated in *w*. In that case, after all, D is not even recognizably pointer-*shaped* in anything but our minimal sense. Plausibly, in fact, any relevant further sense is essentially *contrastive*: what it *takes* to be sufficiently dense, in any given case, may itself depend on some more global distribution of flobs.

References

Albert, D. (1992) *Quantum Mechanics and Experience* (Cambridge, MA: Harvard UP).

Albert, D. (2015) *After Physics* (Cambridge, MA: Harvard UP).

Albert, D. and B. Loewer. (1996) "Tails of Schrödinger's Cat" R. Clifton (ed.) *Perspectives on Quantum Reality: Non-Relativistic, Relativistic, and Field-Theoretic* (Dordrecht: Kluwer) 81-92.

Albert, D. and A. Ney (2013) The Wave Function: Essays on the Metaphysics of Quantum Mechanics (Oxford: Oxford UP).

Allori, V. (2013) "Primitive Ontologies and the Structure of Fundamental Physical Theories" in Albert and Ney (eds.) 58-75.

Bell, J. (1987) "Are there quantum jumps?" in C. Kilmster (ed.) *Schrödinger: Centenary Celebration of a Polymath* (Cambridge: Cambridge UP) 41-52.

Chalmers, D. (2012) *Constructing the World* (Oxford: Oxford UP).

Dürr, D., S. Goldstein, and N. Zanghì (1992) "Quantum Equilibrium and the Origin of Absolute Uncertainty" *Journal of Statistical Physics* 67: 843-907.

Emery, N. (2017) "Against Radical Quantum Ontologies" *Philosophical and Phenomenological Research* 95(3): 564-91.

Ghirardi, D., A. Rimini, A. and T. Weber (1986) "Unified Dynamics for Microscopic and Macroscopic Systems" *Physical Review D* 34: 470-91.

Heil, J. (2012) *The Universe As We Find It* (Oxford: Oxford UP).

Kim, J. (1987) "'Strong' and 'Global' Supervenience Revisited" *Philosophy and Phenomenological Research* 48: 315-26.

Lewis, P. (2006) "GRW: A case study in quantum ontology" *Philosophy Compass* 1(2): 224-44.

Maudlin, T. (2007a) "Completeness, supervenience and ontology" *Journal of Physics A: Mathematical and Theoretical* 40: 3151-71.

Maudlin, T. (2007b) "Why Be Humean?" *The Metaphysics Within Physics* (Oxford: Oxford UP) 51-77.

Maudlin, T. (2010) "Can the World be Only Wavefunction?" S. Saunders, et al. (eds.) *Many Worlds?: Everett, Quantum Theory, and Reality* (Oxford: Oxford UP) 121-43.

Maudlin, T. (2011) *Quantum Non-Locality and Relativity*, 3rd ed. (Malden, MA: Wiley-Blackwell).

Maudlin, T. (2013) "The Nature of the Quantum State" in Albert and Ney (eds.) 126-53.

Maudlin, T. (2015) "The Universal and the Local in Quantum Theory" *Topoi* 34: 349-58.

Miller, E. (2018) "Local qualities" *Oxford Studies in Metaphysics, Vol. 11:* 224-42.

Ney, A. (2013) "Ontological Reduction and the Wave Function Ontology" in Albert and Ney (eds.) 168-83.

OED online (2020) “disposition, n.” URL= <https://www.oed.com/view/Entry/55123?redirectedFrom=disposition> (accessed September 19).

Oppenheim, P. and H. Putnam (1958) “Unity of Science as a Working Hypothesis” H. Feigl, et al. (eds.) *Minnesota Studies in the Philosophy of Science*, vol. 2 (Minneapolis: Minnesota UP) 3-36.

Rubenstein, E. (forth.) “Grounded Shadows, Groundless Ghosts” *The British Journal for the Philosophy of Science.*

Schrödinger, E. (1935) “Die gegenwärtige Situation in der Quantenmechanik” *Naturwissenschaften* 23: 807-12.

Trimmer, J. (1980) “The Present Situation in Quantum Mechanics: A Translation of Schrödinger’s “Cat Paradox” Paper” *Proceedings of the American Philosophical Society* 124: 323-38.