Maudlin's Challenge Revisited

Peter J. Lewis

October 13, 2011

Abstract

In 1994, Maudlin proposed proposed an objection to the transactional interpretation (TI), involving an absorber that changes location depending on the trajectory of the particle. Maudlin considered this objection fatal. However, the TI did not die; rather, a number of responses were developed, some attempting to accommodate Maudlin's example within the existing TI, and others modifying the TI. I argue that none of these responses is fully adequate. The reason, I submit, is that there are two aspects to Maudlin's objection; the more readily soluble aspect has received all the attention, but the more problematic aspect has gone unnoticed. I consider the prospects for developing a succesful version of the TI in light of this second aspect of the objection.

1 Pseudotime explanation

The central explanatory mechanism of Cramer's transactional interpretation (TI) is the pseudotime sequence; offer waves are sent forwards in time from the particle source to the potential absorbers, *and then* each absorber returns a confirmation wave backwards in time to the source, *and then* a transaction forms along one offer-confirmation pair with a probability given by the amplitude of the returning confirmation wave. But this explanatory mechanism has always been regarded as problematic due to the backwards-causal link embodied by the confirmation wave. The locution "and then" in the pseudotime sequence cannot be understood temporally—but then how should we understand it? Indeed, Cramer himself describes the pseudotime sequence as "a semantic device" and "a pedagogical convention" (1986, 661). But as

Maudlin complains, “all of the seeming illumination provided by the account depends on the pseudotime narrative”, so without the narrative there is no explanation (1994, 198).

It seems to me that the pseudotime problem is a pseudoproblem, resting on the false assumption that all explanations are temporal explanations. Another perfectly good form of explanation is the constraint problem; the effect we see is the only solution to the relevant equations, subject to some applicable constraint. To pick a relevant example, the explanation of a standing wave on a string can be given this form; the standing wave patterns we see are the solutions to the wave equation subject to the constraint that the solution is stable over time. The reason this example is relevant is that, as Cramer notes, the formation of a standing wave provides a good analogy to the formation of a transaction; “an equally valid interpretation of the process is that a four-vector standing wave has been established between emitter and absorber” (1986, 663). In the TI case, the solutions of the fundamental equations of quantum mechanics are retarded and advanced waves, and the constraint that the result be a standing wave has a straightforward justification in terms of consistency; anything but a standing wave would ascribe inconsistent properties to spacetime points.

This understanding of pseudotime explanation puts the above narrative on a perfectly secure footing; the locution “and then” can be understood as the “then” of explanation rather than the temporal “then”. The explanation of the quantum statistics we see is that only certain offer-confirmation pairs are consistent (the standing wave solutions), and nature chooses among the consistent pairs according to the amplitudes of the confirmation waves. Maybe the pseudotime sequence is a mere semantic device if regarded as a temporal narrative, but it is not a mere semantic device if regarded as an explanatory narrative; one should take it literally as the explanation of what we see. Maudlin is right to note that all the illumination provided by the TI depends on the pseudotime narrative, but wrong that the narrative is incoherent.

2 Maudlin’s challenge

But taking the pseudotime narrative literally (as I think the adherent of the TI must) lays one open to Maudlin’s main objection to the TI. Consider the experiment shown schematically in fig. 1. A particle is emitted at time t_0.

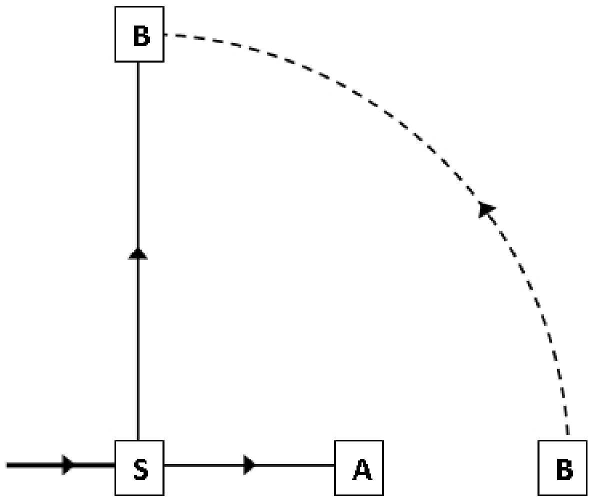

Figure 1: Maudlin's experiment

If it follows the lower path it is detected by absorber A at time t_1. If it follows the upper path then A does not detect a particle at t_1, which triggers an absorber B to be swung from its initial position behind A to a point on the upper path, where it detects the particle at a later time t_2. Now what happens if (by means of the beam-splitter S) a particle is produced in a symmetric superposition of following the lower and upper paths? Standard quantum mechanics tells us that there is a 50% chance that the particle is detected at A and absorber B stays put, and a 50% chance that absorber B swings round and the particle is detected at B.

How can TI account for these statistics? The puzzle here is that the equiprobability of the two outcomes demands that equal amplitude confirmation waves be received from absorber A and absorber B. But if B sends back a confirmation wave, then it must have been struck by an offer wave, so it must have swung round. And if it has swung round, then the particle is not detected at A. So when a confirmation wave is received from absorber B, the particle must always go to absorber B, despite the fact that the amplitude of the confirmation wave is 1/2. That is, the recipe for recovering the standard quantum probabilities at the heart of the TI is inconsistent, and Maudlin concludes that "the TI collapses" (1994, 200).

But the TI has not collapsed; a variety of responses have been developed, either claiming to accommodate Maudlin's example within the existing TI, or modifying the TI to accommodate it. I run through these responses in the first half of this paper, and argue that none of them is satisfactory. The reason, I contend, is that there are two elements to Maudlin's challenge that have not been clearly distinguished. Distinguishing them allows us to evaluate TI much more clearly.

3 Extant responses

Of the extant responses to Maudlin's challenge, Berkovitz (2002), Kastner (2006) and Marchildon (2006) attempt to accommodate Maudlin's example within the existing TI, and Cramer (2005), Kastner (2010) and Chiatti (2011) work with modified versions of the TI. I will start with the former.

Berkovitz (2002) models Maudlin's example as a causal loop, since the confirmation wave returning backwards in time from B is necessary (and sufficient) for the particle to be emitted towards B. Berkovitz then notes that in causal loops the relative frequencies of events can differ significantly from their objective chances. So in this case, the fact that the particle always takes the upper path when B swings round is not inconsistent with the upper path having an objective chance of 1/2. However, it is not clear that Berkovitz intends his analysis as a defense of the TI; he notes that since the link between objective chance and long-run frequency is broken, the TI fails to predict the long-run frequencies of outcomes. This is not a way forward for the advocate of the TI.

Kastner (2006) identifies the following as a problematic aspect of Maudlin's example; if absorber B does not swing round, then the offer wave on the upper path heads into space, and no confirmation wave is received from the upper path at the source. Typically in the TI, the confirmation waves returning from a complete set of absorbers cancel out to the past of the initial emission event. But in this case, the absence of a confirmation wave from the upper path means that the confirmation wave from the lower path propagates into the past, prior to t_0. On the other hand, if B does swing round, then there is a complete set of absorbers, and the confirmation waves *do* cancel out prior to t_0.

The lesson Kastner takes from this is that the emission event is not suitable as a starting-point for the TI analysis, since this event and its past depend on whether B swings round or not, i.e. on the outcome of the experiment. Instead, Kastner takes the offer-confirmation wave pair on the lower path as the starting point of her analysis, since this pair exists whether B swings round or not. The confirmation wave here has an amplitude of 1/2, and so Kastner reasons that the corresponding transaction—in which the particle takes the lower path where it is absorbed by A—has a probability of 1/2. But the only other possibility is that the particle takes the upper path where it is absorbed by B, so this also must have a probability of 1/2. This recovers the standard quantum mechanical probabilities.

But this account is problematic in two regards. First, events prior to the emission event do not depend on the outcome of the experiment as Kastner claims; everything up to the emission event is exactly the same whether the particle is absorbed by A or by B. So something must be amiss if her analysis entails this about the past. Second, Kastner's analysis self-consciously rejects the usual pseudotime analysis. I argued above that the pseudotime analysis is a necessary part of the TI; without it, as Kastner herself admits "we don't as yet even have a heuristic way to understand this process" (2006, 9).

Marchildon (2006) corrects Kastner's point about the past. He notes that when B does not swing round, the offer wave along the upper path is still absorbed *somewhere*; a particle taking this path would be absorbed by something eventually. Due to the retrocausal nature of the pseudotime sequence, it makes no difference how far in the future this absorption event lies; the confirmation wave is still received by the particle source at t_0. One may as well assume that there is a third absorber C, situated on the upper path beyond the point B swings to. Hence a complete set of confirmation waves is received at the particle source whether or not B swings round, and in either case the confirmation waves cancel out prior to t_0.

Maudlin's example remains problematic, however; if B does not swing round, the new absorber C returns a confirmation wave with amplitude 1/2, and yet no transaction is ever completed with C. Marchildon attempts to dissolve this problem by appealing to the four-dimensional space-time implicit in the TI; "the future, though not predictable, is well defined" (2006, 427). The four-dimensional "blockworld" is subject to consistency conditions, and "in the present case, these conditions are that B absorbs the particle if and only if A does not absorb it" (2006, 427). Given this consistency condition, there are just two possible trajectories, and each is ascribed the probability corresponding to the amplitude of its respective confirmation wave, i.e. 1/2.

But to appeal to the existence of two possible trajectories here is, again, to bypass the pseudotime sequence entirely. The pseudotime sequence describes two possibilities if B swings around (absorption by A and by B) and two possibilities if B does not swing around (absorption by A and by C). The probability rule of the TI fails in such cases, because it ascribes probability 1/2 to each of these four outcomes, which is inconsistent (Wharton, Cyber Roundtable). It is certainly true that only two of these are genuine possibilities, but one cannot appeal to this fact to rescue the TI, since this is the conclusion that the TI is supposed to yield, not additional input to the TI. If one wants to rescue the TI as a genuine explanatory theory, one has to find some way for the pseudotime sequence and its associated probability rule to apply to Maudlin's example.

Cramer (2005) attempts to do just that, by modifying the way that pseudotime explanations are constructed. In the standard TI, all possible transactions have the same status; what Cramer proposes is a hierarchy of possible transactions, ordered by the spacetime interval between the beginning and end of the transaction. That is, transactions with shorter spacetime interval are given the opportunity to form or not form "before" (in pseudotime) those with longer spacetime interval. This takes care of Maudlin's example, since the possible transaction with absorber A is decided first. The pseudotime sequence now goes like this: The offer-confirmation pair to absorber A is treated first, and since it has amplitude 1/2, there is a probability of 1/2 of a transaction forming with A. If it fails to form, then absorber B swings round, and the offer-confirmation pair to absorber B can be considered; it too has amplitude 1/2, corresponding to a probability 1/2 of this transaction forming.

But the hierarchical pseudotime sequence envisioned here is not a fully general solution to the problem; it will not work for Maudlin-type experiments involving photons, since in that case the spacetime interval for every possible transaction will be zero (Miller, Cyber Roundtable). Furthermore, it is hard to motivate the hierarchy from within the TI; the confirmation waves from closer absorbers do not arrive at the emission point first, so the resultant explanation seems to rest on a fiction. The hierarchy looks suspiciously ad hoc when applied to cases that do not have the Maudlin contingent-absorber structure; why should shorter transactions be decided first in such contexts (Marchildon, Cyber Roundtable)? That is, the relevant structure here is causal structure, not spacetime structure, and Cramer's hierarchy fails to capture this aspect of Maudlin's example.

Kastner (2010) suggests that the way to retain the status of pseudotime histories as genuine explanations is to adopt a modal realist ontology; all possible transactions exist in a (real) space of possibilities. She calls the resulting modification of the theory "possibilist TI" (PTI). But it is not clear that reifying the possible transactions helps with Maudlin's challenge—in fact, it brings his criticism into sharper focus. If the pseudotime sequence describes a genuine process occurring in a space of possibilities, then the inconsistencies in the pseudotime sequence for the Maudlin example cannot be waved away.

Chiatti (2011) regards the TI pseudotime sequence as superfluous, since he thinks that the structure of transactions can be found within the formalism of standard quantum mechanics. The way to do this, he argues, is to adopt an ontology of particle creation/destruction events; these are the only events that physically exist, and the TI offer and confirmation waves are "a mathematical fiction suited to calculate the statistics of the connection" between a pair of events (2011, 2). The physical universe is regarded as a network of particle creation/destruction events, with the TI offer and confirmation waves serving only to coordinate the statistical relations among these events.

Chiatti claims that in this context, Maudlin's argument is unfounded because the two possible transactions have different events at their extremities; absorption by A is a different event than absorption by B. Different extremal events entail different TI-style analyses, and hence we have two simple applications of the theory, not one paradoxical one. But this decoupling of the problem fails to show how the event of B-absorption is related to that of A-absorption. However, Chiatti has a suggestion here, namely that "the event (not A)$^+$(not A)$^-$ is itself a transaction termination (null interaction)", so that "the second transaction assumes as its input the output state of the first" (2011, 28). The suggestion is that the event of a particle *not* being absorbed by A can serve as the terminal event of a transaction, with probability 1/2, and this event can also serve as the initial event of a second transaction that ends with a particle absorbed by B. The trouble with this suggestion is that a null interaction—a particle failing to be destroyed—is not a physical event in Chiatti's ontology, so it is hard to see how it can serve as the terminal event in a transaction. There are not two back-to-back transactions here, but a single transaction with a complex structure. As a mathematical fiction Chiatti's proposal may be acceptable; the main problem is that Chiatti, like several others, rejects the pseudotime sequence as genuinely explanatory, and hence robs the TI of its explanatory mechanism.

4 The first challenge

It looks like none of the extant responses to Maudlin's challenge are adequate, at least if one insists on finding a role for the pseudotime sequence in the TI. The feature of Maudlin's argument to which all commentators have been attempting to respond is the contingent nature of the absorber structure—the fact that the locations of the absorbers depend on the trajectory of the particle. This is not surprising, since Maudlin himself describes his challenge in these terms: "This picture depends crucially on the idea that the absorbers are somehow just sitting out there in the future, waiting to absorb... But there is no reason for the absorbers to be fixed in the future, unaffected by everything that happens in the present" (1994, 199). However, I think the contingent structure of the absorbers is only one element of Maudlin's challenge, and not the most problematic one.

First, note that Maudlin's experiment involves a mixed quantum/classical system; the particle is a quantum system, but the movable absorber is a classical (macroscopic) system. This, though, looks like an expendable feature of Maudlin's example; one could make all the subsystems quantum mechanical, and still retain the contingent absorber structure. For example, consider the following variant of the experiment (fig. 2). If a particle is emitted along the lower path, it collides with a carefully-timed incoming anti-particle (dotted line) at t_1. The two particles annihilate, and the resulting photon is detected at A. If a particle is emitted along the upper path, then clearly the anti-particle doesn't encounter it at t_1; instead, the antiparticle travels on, and collides with the particle at a later time t_2 on the upper path. This time the resulting photon is detected at B.

The moving parts in this version of Maudlin's experiment are all quantum systems, but the absorber structure is just as in the original version, with the anti-particle playing the role of the moving absorber. There is a slight difference, namely that in the original version there are two absorbers, one of which moves, whereas in the new version there is just one moving absorber. But this is unimportant; in the original version, it could just as well be absorber A that moves to the upper path when it fails to detect a particle on the lower path. So if it is the contingent absorber structure that is the problematic aspect of Maudlin's example, it ought to be just as problematic in this version. But it does not appear so; one can construct a fairly straightforward pseudotime narrative in this case that yields the right probabilities.

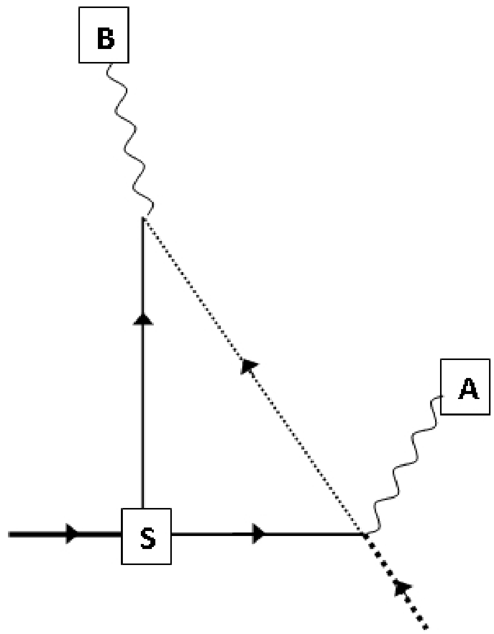

Figure 2: Maudlin experiment: Quantum version

Here is how it goes. A particle and an anti-particle are emitted at the beginning of the experiment, so let us write the initial offer wave as $|O\rangle_p\,|O\rangle_a$. The offer wave for the particle is split by the beam-splitter into two equal components, following the lower path and upper path respectively; at this stage, we can write the total offer wave as

$$\frac{1}{\sqrt{2}}\left(|L\rangle_p + |U\rangle_p\right)|O\rangle_a\,. \tag{1}$$

The offer wave for the anti-particle travels towards the lower path, where it is split into two terms depending on whether it meets the particle on the lower branch or not. That is, the term $|L\rangle_p\,|O\rangle_a$ in (1) evolves to $|L\rangle_p\,|L\rangle_a$, where $|L\rangle_a$ is an offer wave for the anti-particle that encounters the particle on the lower branch. The term $|U\rangle_p\,|O\rangle_a$ in (1) evolves to $|U\rangle_p\,|U\rangle_a$, where $|U\rangle_a$ is an offer wave for the anti-particle that travels onwards through the lower path and encounters the particle on the upper path. Hence the superposition state (1) as a whole evolves to the "later" (in pseudotime) superposition state

$$\frac{1}{\sqrt{2}}\left(|L\rangle_p\,|L\rangle_a + |U\rangle_p\,|U\rangle_a\right). \tag{2}$$

Each term returns a confirmation wave from the point at which the particle and anti-particle annihilate; the first term returns a confirmation wave $\frac{1}{2}\langle L|_p \langle L|_a$, and the second term returns a confirmation wave $\frac{1}{2}\langle U|_p \langle U|_a$. Hence the (dual) source receives two confirmation waves, each with amplitude 1/2, and a transaction forms along one of the offer-confirmation pairs with probability 1/2 each. This pseudotime narrative hence successfully ascribes probabilities of 1/2 each to the two possible transactions—the one in which the particles annihilate on the lower path, and the one in which they annihilate on the upper path.

Hence TI can account for the standard quantum probabilities in contingent-absorber experiments without any new difficulties or any modifications. Why, then, has Maudlin's example been taken to be so problematic?

5 The second challenge

The reason is that the contingent absorber structure in Maudlin's example is not instantiated by quantum systems, but by classical (macroscopic) objects. As explained in the previous section, the straightforward and natural response to the contingent absorber problem just sketched involves incorporating the absorber (the anti-particle) into the TI pseudotime analysis, rather than treating it as part of the environment. But such incorporation is far more problematic if we need to incorporate the state of a macroscopic object into our TI analysis. There are two tricky issues that arise when one contemplates treating a macroscopic object as falling within the scope of a TI analysis. The first issue concerns whether the proposed analysis even makes sense; it depends on the way in which TI practitioners conceive of their theory. The second issue concerns the nuts and bolts of constructing such an analysis, given that it makes sense.

To understand the first issue, a little broad-brush history will help. In broad terms, there are two traditions in the interpretation of quantum mechanics. The first, harkening back to Bohr and Heisenberg, takes the distinction between the quantum world and the classical world as basic; call this the Copenhagen tradition. According to this tradition, the world (or at least, the part of it relevant to a given experiment) divides into system and apparatus. The apparatus belongs to the world of experience, and behaves classically. The apparatus delivers results which we take to be produced by an unseen micro-world. Quantum mechanics does not describe the workings of this micro-world directly, but rather describes what we should expect to see when the apparatus interacts with the system. On this view, the system/apparatus divide is built into the theory of quantum mechanics itself; the observables of the theory correspond to the operations of applying various pieces of macroscopic measuring equipment to the quantum system. Quantum mechanics applies to the interaction between classical measuring devices and quantum systems, and so any attempt to apply quantum mechanics to the measuring equipment itself is fundamentally misguided.

The TI is naturally understood as part of this Copenhagen tradition—as a way of making it more precise. Emitters and absorbers are part of the classical world; we arrange them in constructing an experiment, and we observe the results they present. TI quantum mechanics generates probabilistic predictions based on the particular arrangement of emitters and absorbers we have constructed. On this understanding of the TI, it is about the interaction of an unseen quantum world with a *given* arrangement of emitters and absorbers. To attempt to incorporate the absorber itself into a TI analysis on this view is like trying to incorporate the concert hall into an acoustic analysis; the concert hall, like the arrangement of emitters and absorbers, constitutes the environment in which the analysis takes place. If we understand the TI in this way, of course, then the straightforward route to resolving Maudlin's first challenge is blocked as a matter of principle; there is no way round.

But there is a second tradition in the interpretation of quantum mechanics, stretching from Einstein and Schrödinger through Bell to Maudlin; call this (without too much prejudice, I hope) the realist tradition. According to this tradition, quantum mechanics really does describe the workings of the micro-world, not just system/apparatus interactions. The observables of quantum mechanics should not be conceived in terms of measurement operations, but in terms of actual properties of the quantum system; they are *be*ables rather than *observ*ables, in Bell's memorable phrase (Bell 1987, 174). On this view, there is no significant distinction between the micro-world and the macro-world; quantum mechanics applies just as much to the latter as to the former, since macroscopic systems are built out of microscopic ones. Hence quantum mechanics can be applied unproblematically to the measuring devices themselves. There is no reason in principle why the TI can't be thought of in these terms; perhaps some of its practitioners do so. If so, the first tricky issue is a non-issue; there is nothing in principle to prevent the incorporation of a macroscopic absorber into a TI analysis.

It is worth noting here that the realist tradition faces a problem that doesn't arise in the Copenhagen tradition; quantum mechanics treats measurements differently from non-measurements (only the former trigger collapse), but if measurements are themselves quantum processes this distinction cannot arise. Various responses to this measurement problem have been developed, and it is interesting that all of them (or at least all the major ones) treat Maudlin's experiment by including absorber B in the system to be analyzed. This is most obvious in the Everett (many-worlds) theory. Applying Everett to Maudlin's original experiment yields the following final state:

$$\frac{1}{\sqrt{2}}\left(|L\rangle_p\,|Y\rangle_A\,|N,L\rangle_B + |U\rangle_p\,|N\rangle_A\,|Y,U\rangle_B\right), \tag{3}$$

where $|Y\rangle_A$ and $|N\rangle_A$ are states of A in which it does and does not absorb a particle, and $|N,L\rangle_B$ and $|Y,U\rangle_B$ are states of B in which it doesn't absorb a particle and remains on the lower branch, and does absorb a particle after swinging to the upper branch. The existence of these two terms, according to the Everettian, explains the observed results; in one branch of reality the particle takes the lower path and is absorbed by A, and in the other branch of reality the particle takes the upper path and absorber B swing round to absorb it there. Bohm's hidden variables theory tells essentially the same story, except that one branch is associated with the Bohmian particles, and hence corresponds to the actual result. The GRW collapse theory appeals to the instability of (3) under its collapse dynamics to account for the fact that the final state is (close to) one of the two terms. The lesson (to generalize a little) seems to be that one *has* to incorporate absorber B into the quantum mechanical analysis if one is to give any adequate account of the Maudlin experiment.

As mentioned above, provided that the TI is seen as a direct account of the micro-world in the spirit of the realist tradition, there is nothing to prevent it doing just that. But the details of how this is to be accomplished remain somewhat murky; this is the second tricky issue advertised above. The immediate problem is that the TI analysis (as it stands) admits only two kinds of end-points—emission events and absorption events. To be subject to TI analysis, particles must be followed from birth to death. But the particles that make up absorber B are not emitted at the start of the experiment or absorbed at the end. If we have to know the full life-history of all the particles that make up absorber B before we can apply the TI, the theory becomes impossible to apply.

However, the details of the evolution of the particles in B before and after the experiment seem irrelevant to the analysis at hand. So perhaps the TI can avail itself of the following harmless myth; *pretend* that all the particles involved in the analysis are created at the beginning of the experiment and destroyed at the end. This myth might be justified by appealing to the fact that wherever and whenever the particles in B are actually created and destroyed, the offer and confirmation waves over the course of the experiment will correspond to their mythical counterparts. If that justification works, and the myth is adopted, then a pseudotime narrative for Maudlin's original experiment can be given along the lines of the previous section.

Here is how it goes. According to the myth, the test particle and all the particles in B are emitted at the beginning of the experiment. Let us write the initial offer wave as $|O\rangle_p |N, L\rangle_B$ (using the same notation as above), since B is initially in a state in which it is on the lower path and has not absorbed a particle. The offer wave for the particle is split by the beam-splitter into two equal components, following the lower path and upper path respectively; at this stage, we can write the total offer wave as

$$\frac{1}{\sqrt{2}}\left(|L\rangle_p + |U\rangle_p\right)|N, L\rangle_B\,. \tag{4}$$

The term $|L\rangle_p |N, L\rangle_B$ in (4) evolves to $|L\rangle_p |N, L\rangle_B$, and the term $|U\rangle_p |N, L\rangle_B$ evolves to $|U\rangle_p |Y, U\rangle_B$; hence the superposition state (1) as a whole evolves to the "later" (in pseudotime) superposition state

$$\frac{1}{\sqrt{2}}\left(|L\rangle_p |N, L\rangle_B + |U\rangle_p |Y, U\rangle_B\right). \tag{5}$$

Note (5) is exactly the same as the Everettian final state (3), except that absorber A has not been included in the current analysis. Applying the myth again, we assume the test particle and all the particles in B are destroyed at this point, so that each term in (5) returns a confirmation wave. from the point at which the particle and anti-particle annihilate; the first term returns a confirmation wave $\frac{1}{2}\langle L|_p \langle N, L|_B$, and the second term returns a confirmation wave $\frac{1}{2}\langle U|_p \langle Y, U|_B$. Hence the source receives two confirmation waves, and a transaction forms along one of the offer-confirmation pairs with probability 1/2 each. This pseudotime narrative successfully ascribes probabilities of 1/2 each to the two possible transactions—the one in which the particles takes the lower path and B stays put, and the one in which the particle takes the upper path and B swings round.

6 Transactions and trajectories

So has the Maudlin puzzle been solved? The pseudotime sequence just sketched constitutes what I take to be the most natural, and perhaps the only, available solution. But at what cost have we arrived at this solution? We have had to find a way to incorporate the analysis of pieces of measuring equipment into the TI analysis. Some may find this price too high, and with good reason.

At issue is the uniqueness of classical trajectories in the TI. It is important to note that TI transactions are not always particle trajectories. Two-slit interference is a case in point; the completed transaction goes through both slits to a point on the screen, so the transaction is not a determinate particle trajectory, but a superposition of such trajectories. In fact, this is a generic feature of interference; whenever two or more distinct offer waves contribute to the amplitude at the absorption point, the resulting transaction incorporates the trajectories corresponding to all the offer waves to that point. So the TI does not always recover determinate particle trajectories as e.g. Bohm's theory does. But provided such indeterminacy is kept confined to the micro-world, this is not a problem.

The worry about incorporating macroscopic objects into the TI analysis is that it opens the door for macroscopic objects to have indeterminate trajectories. Thanks to decoherence, interference effects are tiny for macroscopic objects, but they do not go away entirely. So in my treatment of the Maudlin experiment above, the amplitudes of the terms in (5) are affected by the presence of anomalous interference terms. For example, there are low-amplitude terms in which the particle takes the upper path and absorber B swings round, but then at the last instant they veer back to the lower path to coincide with the first term in (5). So in reality, the wave amplitude at the lower-path location for B includes tiny contributions from offer waves that have not travelled via the lower path. According to the TI, then, the completed transaction contains a large contribution following the standard trajectory, but also minor contributions following the anomalous trajectories. This might seem to threaten the determinacy of trajectories for macroscopic objects, either rendering the TI empirically inadequate or turning it into a baroque version of Everett.

But perhaps such worries can be deflected; the additional terms are, after all, very small. One might quite reasonably insist that a transaction in which one trajectory is so dominant simply is, for all practical purposes, that trajectory. Provided that no *significant* interference occurs involving macroscopic objects, the determinacy of macro-trajectories is (arguably) safe. However, there is a further worry that cannot be dealt with in this way. If the TI is interpreted in the realist tradition, there is nothing to prevent it being applied to the state of the universe as a whole (just as Everett, Bohm and GRW can be so applied). Indeed, this is one of the touted advantages of the realist tradition. One wouldn't need the myth for such an application; we really would be following every particle from birth to death. But there is no guarantee that interference between macroscopically distinct terms can be suppressed in such an application; it depends on the global structure of the universe. For example, in highly symmetric universes, branching in the initial stages of the universe might be matched by "reverse branching" (i.e. interference) in the final stages. In that case, there would be no unique trajectory corresponding to the evolution of the macroscopic objects in the universe, and again the TI would either be empirically inadequate or reduce to a version of Everett.

One might be tempted to dismiss such cosmological speculation, and trust that the universe will have a structure amenable to the application of the TI. But it is at least unclear whether the TI can be unproblematically applied to macroscopic objects. Without such application, though, Maudlin's challenge retains its power against the TI.

Acknowledgments

I would like to thank the participants in the "TI Cyber-Roundtable" (March–August 2011) for illuminating discussions on many aspects of the TI, and especially Ruth Kastner for organizing and moderating the event.

References

Bell, J. S. (1987), *Speakable and Unspeakable in Quantum Mechanics.* Cambridge: Cambridge University Press.

Berkovitz, J. (2002), "On causal loops in the quantum realm", in T. Placek and J. Butterfield (eds.), *Non-locality and Modality.* Dordrecht: Kluwer, pp. 235–257.

Chiatti, L. (2011), "The transaction as a quantum concept", draft.

Cramer, J. G. (1986), "The Transactional Interpretation of Quantum Mechanics", *Reviews of Modern Physics* 58: 647–687.

——— (2005), "The quantum handshake: a review of the transactional interpretation of quantum mechanics", presented at *Time-Symmetry in Quantum Mechanics* conference, Sydney, Australia, July 23, 2005. Available at: http://faculty.washington.edu/jcramer/PowerPoint/Sydney_20050723_a.ppt

Kastner, R. E. (2006), "Cramer's transactional interpretation and causal loop problems", *Synthese* 150: 1–14.

——— (2010), "The quantum liar experiment in Cramer's transactional interpretation", *Studies in History and Philosophy of Modern Physics* 41: 86–92.

Marchildon, L. (2006), "Causal loops and collapse in the transactional interpretation of quantum mechanics", *Physics Essays* 19: 422–429.

Maudlin, T. (1994), *Quantum Non-Locality and Relativity.* Oxford: Blackwell.