Counterfactuals and the Past Hypothesis

Mathias Frisch†

Albert (2000) provides a sketch of an entropy account of the causal and counterfactual asymmetries. This paper critically examines a proposal that may be thought to fill in some of the lacunae in Albert's account.

1. Introduction. When I pour milk into my coffee in the morning, milk and coffee mix, yet we never observe milk and coffee separating spontaneously. Your setting the toaster to 'dark' caused my toast to burn—if you hadn't adjusted the toaster my toast would not have burned—while the burning of the toast did not cause the prior adjusting of the heat. These are examples of three pervasive temporal asymmetries: on the one hand, the asymmetry of thermodynamics, according to which the entropy of a closed system never decreases and, on the other hand, a causal and a counterfactual asymmetry, according to which the future is both causally and counterfactually dependent on the present, while the past is not. What if anything is the relation between these asymmetries?

Albert (2000) suggests that all three asymmetries share the same micro-statistical foundations, which involve crucially what he calls *the past hypothesis*—the assumption that the universe began its life in a state of extremely low entropy. In fact, Albert appears to believe that the time asymmetric assumption of a past hypothesis without a corresponding 'future hypothesis' is, as it were, the Holy Grail of the philosophy of time and provides us with a single 'master arrow' from which other temporal asymmetries can be derived. I criticize Albert's account of the causal and counterfactual asymmetries in Frisch (forthcoming). Here I want to examine a variant of the view Albert proposes in his book—an account suggested by recent work of Barry Loewer (see Loewer, forthcoming)—that might be thought to fill in some of the lacunae in Albert's account.

After briefly summarizing the micro-statistical account of the thermodynamic asymmetry endorsed by both Loewer and Albert, I will sketch

†To contact the author, please write to: Department of Philosophy, University of Maryland, College Park, MD 20742; e-mail: mfrisch@umd.edu.

Philosophy of Science, 72 (December 2005) pp. 739–750. 0031-8248/2005/7205-0008$10.00

how this theory might be invoked to account for the temporal asymmetry of certain kinds of counterfactuals. In Section 3 I will raise worries about different steps of the argument and will conclude that the account does not successfully reduce the causal and counterfactual asymmetries to thermodynamic considerations. I am not here interested in challenging Albert and Loewer's account of the thermodynamic asymmetry and will, for the sake of argument, take it for granted that what happens in my coffee cup in the morning can indeed be explained by appealing to an assumption about the early universe. Rather, my focus here will be on the claim that the micro-statistical account can also explain a counterfactual asymmetry, which then might account for the asymmetry of causation.

2. The Past-Hypothesis and Counterfactuals. The thermodynamic asymmetry that the entropy of a closed macroscopic system never decreases is explained, according to the view endorsed by Albert and Loewer, by appealing to a time-symmetric micro-dynamics and an asymmetric constraint on initial conditions. If we assume an equiprobability distribution of micro-states compatible with a given macro-state of nonmaximal entropy, then it can be made plausible that, intuitively, 'most' micro-states will evolve into states corresponding to macro-states of higher entropy. However, if the micro-dynamics governing the system is time-symmetric, then the same kind of considerations also appear to show that the system evolved *from* a state of higher entropy. This undesirable retrodiction, which is at the core of what is known as the reversibility objection, can be blocked, if we conditionalize the distribution of micro-states not only on the present macro-state but also on a low entropy initial state of the system. Albert and others argue that since the reversibility objection arises for all times in the past, we are ultimately led to postulate an extremely low entropy state for the early universe.

Thus, the assumptions of the micro-statistical account (STAT MECH) are the following:

i. time-symmetric, deterministic dynamical micro-laws;
ii. the *past hypothesis* (PH) which characterizes the initial macro-state of the universe as a low-entropy condition satisfying certain further symmetry conditions;
iii. a *probability postulate*, according to which all micro-states compatible both with the universe's macro-state at time t, $M(t)$, and PH are equiprobable.

How can this account explain a time-asymmetry of counterfactuals? Loewer (forthcoming) argues that it is a consequence of the statistical mechanical account that possible evolutions of the universe form a *tree-structure* and that small differences in the micro-conditions at t that do not involve changes in the macro-state can give rise to very different macro-futures but not to different macro-pasts (TREE). While the underlying micro-dynamics is deterministic, Loewer suggests macro-histories form a tree structure branching toward the future. That is, according to TREE there are different macro-futures with non negligible conditional probabilities, conditionalized on the present macro-state, while any actual past macro-state has a probability close to 1, given the present macro-state, where the probabilities are those induced by the statistical mechanical probability distribution.

One alleged consequence of this tree-structure is that the present contains much more information about the past than about the future; and that much of this information consists in localized *records* of the past (RECORD). That is, the statistical mechanical account is taken to explain an asymmetry characterizing our knowledge of the world at times other than the present—namely, the fact that there exist records of past events but not of future events. The records claim plays a prominent role in Albert (2000), and I criticize it in detail in Frisch (forthcoming), but since it plays a less important role in Loewer (forthcoming), I will have less to say about it here.

At the heart of the proposal I wish to consider here is the claim that the tree structure can explain a time-asymmetry of two kinds of counterfactuals. The first kind includes *decision counterfactuals* of the form "If I were to decide D then the probability of E would be *p*"; and the second includes what one might call 'paradigmatically causal' counterfactuals postulating small macroscopic changes to the world. An example of the latter kind is "If the cue ball had struck the eight ball, the eight ball would have had a very high chance of having gone into the corner pocket." I say that counterfactuals such as this are causal, since their truth is closely tied to the truth of certain causal claims—in the present example the claim that the cue ball's striking the eight ball *caused* the eight ball to go into the corner pocket. Both Loewer and Albert apparently believe that the truth of the counterfactual claim underwrites the truth of the causal claim, but even those opposed to some kind of counterfactual account of causation would presumably agree that there is *some* close relation between the causal and counterfactual claims.

According to the proposal, given appropriate truth conditions, TREE implies that forward looking probabilistic decision counterfactuals and counterfactuals positing small macro-changes may be true, but backtracking counterfactuals of these kinds are false in general (ASYM).

How, then, are we to evaluate probabilistic counterfactuals of the form "If event *C* were to occur, then the probability of macro-event *E* would be *p*"? In contrast to Lewis's well-known semantics for counterfactuals (Lewis 1986), Loewer and Albert's proposals for evaluating such counterfactuals does not involve macro-miracles. Rather, in the case of forward looking counterfactual conditionals postulating some non-actual event $C(t_1)$, the counterfactual is evaluated by postulating "micro-miracles" at the latest time t_0 possible such that (i) the macro-state $M(t_0)$ in a 'miracle world' is identical to the actual macro-state and (ii) the micro-states evolve in accord with the micro-dynamics in such a way that $C(t_1)$ occurs. That is, forward looking counterfactuals are evaluated by looking at micro-histories that are identical to the actual micro-history up to t_0, when they diverge from the actual micro-history, and the divergence of micro-evolutions occurs in such a way that macro-states at t_0 are identical to the actual macro-state at t_0. Let us then consider the following truth conditions for probabilistic counterfactuals:

> (PC) "If event C were to occur, then the probability of macro-event E would be p" is true exactly if $P(E(t_2)/C(t_1)\ \&\ M(t_0)) = p$.[1]

In the case of paradigmatically causal counterfactuals, C will usually be a macro-event. In the case of decision counterfactuals, we might tentatively assume that C will be a micro-event, if we assume that changes in decisions correspond to micro-changes in a person's brain state. But if we wish to avoid to be committed to a particular neurophysiological model of decisions, the account ought to be able to recover the asymmetry for both kinds of antecedent events.

If $C(t_1)$ is a *macro*-state, then the micro-miracle needs to be introduced at some time strictly prior to the counterfactual state $C(t_1)$ in which we are interested: The miracles need to be introduced at the last time t_0 when the macro-future at t_1 is still 'open' with respect to the macro-state at t_0. Thus, the times in condition PC satisfy the strict inequality $t_0 < t_1 < t_2$.

By contrast, if the counterfactual antecedent state $C(t_1)$ is a *micro-state* compatible with the actual macro-state, then t_1 itself is the latest time at which a micro-miracle can be introduced that is compatible with the macro-state at that time and results in $C(t_1)$. Thus, $t_1 = t_0$ and, in the case of decision counterfactuals, $M(t_0)$ is the macro-state at the time of decision $D(t_0) = C(t_1)$. Thus, "If I were to decide D at t_0, then the probability of macro-event $E(t_2)$ would be p" is true exactly if

$$P(E(t_2)/M(t_0)\ \&\ D(t_0)) = p. \tag{1}$$

So far I have discussed only *forward* looking counterfactuals, which in the second case are evaluated by postulating a micro-miracle at a time prior to the occurrence of the counterfactual's antecedent. The truth conditions for *backtracking* counterfactuals ought to be analogous, with the

1. This probability is also conditionalized on the past-hypothesis, which is left implicit.

temporal order reversed such that $t_0 \geq t_1 > t_2$. Any other procedure for assessing backtrackers would risk smuggling in a temporal asymmetry illegitimately.[2] That is, in the backtracking case we have to examine micro-histories that have the same micro-*future* as the actual world and converge with the micro-history of the actual world through a miracle that occurs at some time t_0 *no earlier* than the time of *C*, which is postulated to occur after *E*. The temporal asymmetry of counterfactuals consists in the purported fact that there are many pairs of non-actual events $E(t_2)$ and $C(t_1)$ with $t_0 \leq t_1 < t_2$ such that $P(E(t_2)/C(t_1)\ \&\ M(t_0))$ can differ appreciably from zero; but for (almost) all pairs of non-actual (suitably small and localized) events $E(t_2)$ and $C(t_1)$ with $t_0 \geq t_1 > t_2$ $P(E(t_2)/C(t_1)\ \&\ M(t_0))$ is negligibly small.

3. Critical Discussion.

3.1. Does TREE Imply ASYMM? In assessing whether TREE, in conjunction with PC, implies ASYMM, we need to distinguish between the case of micro-miracles and that of macro-miracles. In the case of counterfactuals that postulate *micro*-changes, ASYMM does indeed appear to be a direct consequence of TREE. According to TREE the future is open—that is, micro-changes compatible with the present macro-state can lead to different macro-futures. The past, by contrast, according to TREE, is insensitive to most changes in the current micro-state compatible with the present macro-state. That is, the probability *p* in (1) can differ appreciably from zero, if *E* is in the future of *D*, and will generally be negligibly small if *E* is in the past of *D*.

TREE does not, however, imply that backtracking probabilities are negligible in the case of counterfactuals with antecedents postulating *macro*-changes. Even if we assume it to be the case that the probability of any non-actual macro-event at t_1 is small, given the actual macro-state at some later time t_0—the time at which we imagine the micro-miracle to take place through which the counterfactual micro-history converges with the actual history—it does not follow that the probability of all non-actual events E occurring earlier than *C* is negligible given *both* $C(t_1)$ and

2. If forward looking counterfactuals are assessed by determining the consequences of microscopic *divergence* miracles, then backtracking counterfactuals have to be assessed by considering microscopic *convergence* miracles. If in the backtracking case the time of the relevant macro-state *M* was assumed to be *earlier* than the time of the decision at issue, the account would commit an error similar to one committed by Lewis in his argument for why divergence requires less of a miracle than convergence. Instead of considering pure convergence miracles, Lewis stacks the deck by considering *re*-convergence miracles (see Frisch 2005 for a criticism of Lewis's account).

$M(t_0)$. Clearly, the inequality

$$P(E(t_2)/M(t_0)) \leq \varepsilon, \quad \text{with } \varepsilon \approx 0, \tag{2}$$

as postulated by TREE, does not imply

$$P(E(t_2)/C(t_1) \ \&\ M(t_0)) \leq \delta, \quad \text{with } \delta \approx 0, \tag{3}$$

where $t_0 > t_1 > t_2$, as required in the case of backtracking counterfactuals. Thus, at the very least an additional argument would be needed to support the move from TREE to ASYMM.

We can ask what condition would have to be added to (2) to allow us to infer (3). Expanding the left-hand side of (2), and suppressing the time-dependence, we obtain

$$P(E/M) = P(E/C \ \&\ M) \times P(C/M) + P(E/{\sim}C \ \&\ M) \times P({\sim}C/M) \leq \varepsilon. \tag{4}$$

Since both summands are nonnegative, it follows that

$$P(E/C \ \&\ M) \times P(C/M) \leq \varepsilon. \tag{5}$$

But from (5) we can infer (3) only under the additional assumption that P(C/M) is not too small:

$$P(C/M) \gg \varepsilon. \tag{6}$$

Yet one might worry that (6) is in tension with TREE. According to TREE, the past is fixed by the present macro-state, and hence

$$P(C(t_1)/M(t_0)) \leq \gamma, \quad \gamma \approx 0, \tag{7}$$

for any counterfactual event C in the past of M, i.e., for $t_0 > t_1$.

Now, it is possible for both (6) and (7) to hold, for γ could be very small and yet many orders of magnitude larger than ε. But in order for the argument from (2) to (3) to go through, we must assume that the probability of any event conditional on the macro-state in its future, such as $P(C(t_1)/M(t_0))$, is *extremely* sensitive to the time difference $t_0 - t_1$—to ensure that $P(C(t_1)/M(t_0))$ can be orders of magnitude larger than $P(E(t_2)/M(t_0))$—and, furthermore, that the *earlier* a counterfactual past state is, the *higher* is the probability that it could *not* have occurred, given the complete actual macro-state M at some time in its future. Yet intuitively, this gets things backward: the closer an actual past event A is to the time of M, the more unlikely it seems that 'things could go wrong' and the more nearly fully determined A should be by M.

Consider the following counterfactual: If the ball had not rolled into the corner pocket a second ago, it could not have been struck by the cue ball slightly earlier. And take M to be the current macro-state of the world, including the billiard table with a ball at rest in the corner pocket. The present proposal forces us to accept that the probability of the ball *not* having rolled into the pocket a second ago, given that it currently is in the pocket, is orders of magnitude *larger* than the probability of its *not* having been struck by the cue ball slightly earlier, given that it currently is in the pocket. Again, this seems to conflict with our intuitive assessment that the state of the ball as it is rolling into the corner pocket is more fully determined by the current state of the world (with the ball at rest in the pocket) than is the state of the ball at some earlier time when it was struck by the cue ball. Moreover, it is not clear how the additional assumption of a low-entropy state billions of years ago might affect this assessment. Albert (2000) assumes that without positing the past-hypothesis, all our retrodictions will be radically mistaken. But it is not obvious how it might follow from this that once we do assume the past-hypothesis, past states are more fully determined by the present the further distant these states are.

Thus, the proposal under consideration fails for counterfactuals with antecedents positing small macroscopic changes. The assumption that the past (unlike the future) is fixed by the present macro-state does not imply that backtracking counterfactuals positing small macro-changes are false, if such counterfactuals are evaluated according to their prescription. At best, the account can cover decision counterfactuals, and only if decisions can be modeled as pure *micro*-events. It is unclear how the account might be extended to paradigmatically causal counterfactuals concerning possible *macro*-changes to the world.

Loewer (forthcoming) proposes that the counterfactual asymmetry is fixed by counterfactuals whose antecedents are decision micro-events and that this asymmetry is then projected onto counterfactuals involving macro-events. But this suggestion presupposes that some plausible account can be given of why (micro-)*decision* counterfactuals, rather than, for example, counterfactuals linking possible human (macro-)*actions* to their consequences ought to be viewed as the ultimate source of our notion of causal dependence.

3.2. Does TREE Imply RECORD? According to TREE, most micro-states compatible with the present macro-state that result from the present micro-state through small changes will be associated with the same macro-past as is the actual current micro-state. Thus, the probability of any past macro-state of the world, given the present macro-state and the past hypothesis, will be large, and perhaps close to 1:

$$P(M(t_2)/M(t_1)\ \&\ PH) \approx 1. \tag{8}$$

This is not yet, however, the same as the claim that there are *localized* records of the past. According to (8), the *entire* present macro-state determines the entire past. If we assume a relativistic theory, the part of the present macro-state that will determine any localized substate of a past macro-state will be restricted to the forward light cone of that substate. What is still missing is an argument from (8) to the claim that there can be genuinely local records of past events, for which

$$P(R(t_2)/C(t_1)\ \&\ PH) \approx 1, \tag{9}$$

where $R(t_2)$ is a localized record at t_2 of some earlier event $C(t_2)$, such that R does not fix the macro-state in the entire forward light cone of C.

The puzzle concerning the existence of records is this: In the case of a Lorentz-invariant theory that allows us to determine possible time evolutions through a pure initial or final value problem, the state on some finite subregion of a hypersurface is dynamically determined by the cross section of its future light cone with a future hypersurface. Yet we are all familiar with the existence of much more local records of the past. I know that Caesar was killed on the Ides of March, because there are present day records of the event. And I could not possibly know this, if the only form of inference about the past available to me was an inference based on the dynamical laws and appropriate initial or final conditions. How then are such local records possible?

It is not clear to me how an appeal to TREE can answer this problem. TREE says that local *changes* to the present micro-state compatible with the present macro-state do not affect past macro-states. But this on its own does not entail that local present macro-states determine local past macro-states, for the past half of TREE simply is equivalent to the claim that the world is (near) backward macro-deterministic. What remains to be shown is that a world that is backward macro-deterministic, as TREE stipulates, must contain *local* records of its past.

3.3. Does STAT MECH Imply TREE? Central to both Loewer's and Albert's accounts is the assumption that the past hypothesis constrains the past evolution of the universe while the future is not similarly constrained. And it is supposed to follow from this that there are many possible micro-evolutions compatible with the present macro-state that evolve into nonactual macro-states, while the past is not similarly sensitive to changes to the present micro-state.

There are two separate questions we need to consider in assessing the move from STAT MECH to TREE: First, 'Does STAT MECH imply that the future is open?'; and second, 'Does STAT MECH imply that the past is fixed by the present macro-state and PH?'

To address the second question first, it is important to note that while the past hypothesis *constrains* possible initial macro-states of the universe, it is not prima facie obvious how it would determine a *unique* initial macro-state. The past hypothesis constrains all possible micro-histories to those that have originated in *a* low entropy macro-state with the right sort of symmetry properties, whatever they may be. What is still missing is an argument for why this constraint implies the stronger constraint that non-actual micro-histories also would have to have originated in the actual initial state. That is, what still has to be shown is that there could not be micro-histories compatible with the current macro-state that originated in a low entropy initial state *distinct* from that of the actual initial state. Of course, most micro-histories compatible with the present macro-state will not have originated in any low entropy initial state. But what still has to be shown is that of those 'few' histories that do originate in a low entropy state, the overwhelming majority originated in the actual initial state.

PH constrains most micro-histories to result in thermodynamically normal macro-evolutions. That is, macro-evolutions compatible with most micro-histories are from states with lower entropy to states with higher entropy. But high entropy states occupy larger regions of phase space: there are many more micro-states compatible with a given high-entropy macro-state than are compatible with a low-entropy macro-state. Thus, as far as a comparison of sizes of regions of phase space is concerned, it would seem quite possible that different micro-states compatible with the present macro-state could have evolved from different lower entropy micro-states.

This last observation also calls into doubt the claim that STAT MECH on its own implies that the future is open in ways in which the past is not. For future macro-states will occupy much larger regions of phase space than the present and, thus, there are many more possible changes to the current micro-state that will leave the future (higher entropy) macro-state unchanged than there are changes to the current micro-state that will leave the past (lower entropy) macro-state unchanged. Now, this is not yet enough to show that TREE is *inconsistent* with STAT MECH, since TREE is restricted to possible changes in the micro-state that are compatible with the present macro-state. But it does put the missing step in the argument into sharper focus. It is possible, given the respective phase space volumes alone, that all micro-states compatible with the present macro-state evolve into the actual future macro-state at some future time *t*. Thus, what still would need to be shown is that it follows from STAT MECH that many or most small changes to the micro-state compatible with the actual macro-state and with a low entropy initial state are such that they evolve into different future macro-states, even though the phase space volumes associated with future macro-states are large enough to accommodate all micro-states compatible with the current macro-state.

While there clearly are many systems that exhibit an openness toward the future of the kind posited, it is not clear that this openness is a consequence of STAT MECH alone. And neither is it clear how general this feature is. There are also many systems that do *not* exhibit this kind of openness toward the future. Take, for example, the many kinds of system treated in classical mechanics or electrodynamics textbooks, that are modeled as deterministic macroscopic systems. The very fact that there are systems that can be treated deterministically on the macro-level seems to show that there are systems that are not open toward the future in the sense postulated by TREE. If a system is nearly macro-future-deterministic, as many physical systems appear to be, then its future macro-evolution is not sensitive to small changes in the system's current micro-state in the way suggested by them.

Consider the paradigm example of a thermodynamic system: a body of gas that is confined to one half of a container until a partition is removed and that then spreads through the entire container. Almost all small changes to the initial micro-state of the gas will not affect the macro-evolution of the gas and will not change the fact that the gas will spread through the entire container. In fact, it appears to be a consequence of the macro-evolution toward higher entropy states that small differences in micro-states will generally wash out. Thus, the equilibrium state of the gas when it is spread evenly through the entire container is insensitive even to small differences in the gas's macroscopic initial state. Similarly, the macro-state of an absorber of electromagnetic radiation is generally insensitive even to differences in the macro-state of radiation sources. Thus, the macro-state of the walls of a building is to a large extent insensitive to whether and when the lamps in the building were turned on in the past.

At the very least, then, there must be an additional hidden assumption, in addition to STAT MECH in the account, namely an instability assumption according to which macro-systems behave chaotically toward the future. Could the account be saved by simply adding such an assumption? I believe the answer is 'no', for the following two reasons.

First, chaotic systems will generally be chaotic toward both the future and the past. It remains to be shown that constraining the remote past of chaotic systems to states of extremely low entropy (with certain further symmetry properties) is enough to remove the backward chaotic behavior.

Second, proponents of entropy accounts intend their accounts to be rather general. The statistical accounts are meant to be alternatives to Lewis's account of the counterfactual and causal asymmetries and, hence, ought to explain the asymmetry that characterizes *all* counterfactual and causal reasoning concerning relatively small macro-events. But not all actual macro-systems satisfy an instability assumption. Intuitively, my flipping the light switch *causes* the light to go on just as the famous flap of the butterfly's wings *causes* a storm many thousand miles away. Yet arguably the first system will not in general satisfy an instability assumption. Due to the fact that the light is absorbed by the walls (and due to my forgetfulness) whether or not the light is on makes no difference to the future macro-state of the world. Our intuitive counterfactual and causal assessments of the two examples are the same—we do not think that the truth of causal claims depends on whether a given system is chaotic or not—yet the account we have been considering can at most cover the case of the chaotic system.

4. Conclusion. I have argued that a certain proposal for accounting for the causal and counterfactual asymmetries by appealing to entropic considerations is unsuccessful. Central to the account is the claim that a micro-statistical account of the thermodynamic asymmetry implies that possible histories of the universe form a certain tree structure, in that the future but not the past is 'open', given the present macro-state and the past hypothesis. The proposal relies on the additional assumption, however, that the constraint that the universe originated in *a* low-entropy state is sufficient to constrain possible histories to have originated in *the* actual big bang state. Moreover, it appears that the openness of the future is not implied by the statistical mechanical account alone and only follows if an instability assumption is added. But once such an assumption is added, the account loses its putative generality and can no longer explain the time-asymmetry characterizing our counterfactual and causal reasoning in the case of non-chaotic systems.

I have also shown that it is not a consequence of the assumption that the past is nearly deterministically 'fixed' by the present, together with a plausible proposal for the truth conditions for probabilistic counterfactuals, that backtracking counterfactuals positing small macro-changes to the present come out false. Putting the two criticisms together, it follows that the proposal can at most account for a time asymmetry of counterfactuals positing *micro*-changes in *chaotic* systems. As a general account of a time-asymmetry of counterfactuals, it fails.

Existing non-causal accounts of the counterfactual asymmetry have not fared well. In Frisch (2005) I show that Lewis's account of the counterfactual asymmetry is problematic. In Frisch (forthcoming) I argue that Albert's original statistical account of the counterfactual and causal asymmetries is unsuccessful. Perhaps it is time, then, to try a different approach and *begin* with a causal asymmetry that can then help to account for why in certain contexts our counterfactual reasoning exhibits a temporal asymmetry.

REFERENCES

Albert, David Z. (2000), *Time and Chance*. Cambridge, MA: Harvard University Press.

Frisch, Mathias (2005), *Inconsistency, Asymmetry and Non-locality: A Philosophical Investigation of Classical Electrodynamics*. New York: Oxford University Press.

——— (forthcoming), "Causation, Counterfactuals and the Past-Hypothesis", in Huw Price and Richard Corry (eds.), *Russell's Republic: The Place of Causation in the Constitution of Reality*. Oxford: Oxford University Press.

Lewis, David (1986), "Counterfactual Dependence and Time's Arrow", in *Philosophical Papers*. Oxford: Oxford University Press.

Loewer, Barry (forthcoming), "Counterfactuals and the Second Law", in Huw Price and Richard Corry (eds.), *Russell's Republic: The Place of Causation in the Constitution of Reality*. Oxford: Oxford University Press.