Episteme, 14, 2 (2017) 233–253 © Cambridge University Press
doi:10.1017/epi.2016.7

COHERENCE AND CORRESPONDENCE IN THE NETWORK DYNAMICS OF BELIEF SUITES

PATRICK GRIM, ANDREW MODELL, NICHOLAS BRESLIN,
JASMINE MCNENNY, IRINA MONDESCU, KYLE FINNEGAN,
ROBERT OLSEN, CHANYU AN AND ALEXANDER FEDDER
patrick.grim@stonybrook.edu

ABSTRACT

Coherence and correspondence are classical contenders as theories of truth. In this paper we examine them instead as interacting factors in the dynamics of belief across epistemic networks. We construct an agent-based model of network contact in which agents are characterized not in terms of single beliefs but in terms of internal belief *suites*. Individuals update elements of their belief suites on input from other agents in order both to maximize internal belief coherence and to incorporate 'trickled in' elements of truth as correspondence. Results, though often intuitive, prove more complex than in simpler models (Hegselmann and Krause 2002, 2006; Grim *et al.* 2015). The optimistic finding is that pressures toward internal coherence can exploit and expand on small elements of truth as correspondence is introduced into epistemic networks. Less optimistic results show that pressures for coherence can also work directly against the incorporation of truth, particularly when coherence is established first and new data are introduced later.

INTRODUCTION

What is it for a belief to be true? Sketched in the broadest of strokes, the classical debate is between theories of coherence and correspondence. On the coherence theory of truth, whether something is true is simply a matter of whether it coheres with other beliefs. Truth is a quasi-logical and distinctly internal characteristic of a belief set. On the correspondence theory, whether something is true is a matter of whether it correctly represents or corresponds to some fact, a feature of the world entirely distinct from any belief or belief set. Truth is a distinctly external relation between elements of a belief set and the world beyond them.

If we are forced to choose between classical coherence and classical correspondence as theories of what truth is, or as linguistic theories as to what 'truth' means, correspondence clearly triumphs. The whole point of the coherence theory is to analyze truth in terms of some quasi-logical relation internal to a belief set. But given any internal quasi-logical relation as 'coherence,' the core problem for coherence theories is that there can be a set of fully coherent beliefs many of which, or perhaps all of which, are false. Given any measure of coherence, it is possible for any set of beliefs to construct an alternative and contradictory set that is equally coherent (Russell 1910, 1912; Rescher 1973 [1982]; Young 2013). Given that possibility, 'truth' cannot mean and truth cannot be what the classical coherence theory says it is (see also McGinn 2002; Thagard 2007).

It is when we consider coherence and correspondence as epistemic theories of justification – as theories of how we can decide what is true and what is not – that coherence comes into its own (Rescher 1973 [1982]; Haig and Borsboom 2012). We cannot gauge each of our beliefs individually against the belief-independent facts. Indeed it isn't clear how we can access the facts at all independently of our beliefs concerning them. Belief always mediates our contact with the world, and our contact with the world pays off only by propagating through the broadly internal logical relations of our sets of belief (Hempel 1935; Neurath 1983; Davidson 1986). Epistemically, coherence is at least as important as is correspondence (BonJour 1985; Harman 1986; Lehrer 2000; Thagard 2007).

This is merely a broad-stroke characterization of just the two classical contenders in the theory of truth. An expanded sketch might be laid out in terms of realism and anti-realism. Coherentism, constructive empiricism, social constructivism, 'internal realism,' and certain forms of pragmatism would be rendered as an anti-realist cluster. Correspondence theory and still other forms of pragmatism would be rendered as realist. Deflationary, redundancy, and other non-substantive theories would be sketched as a third grouping. For our purposes here, however, a broad-stroke characterization in terms of the two classical contenders will be all we need.

Our aim is not to refine either theory or its variants, nor to decide between them. We seek a better appreciation of the role of something like correspondence and coherence construed as partners rather than as competitors in the social network dynamics of belief and belief change.

In contrast to standard philosophical approaches, our focus will be the impact of coherence and correspondence, acting together, across belief sets embedded in networks of social communication. The great bulk of epistemic network modeling, phrased either explicitly or implicitly in terms of correspondence, has been structured in terms of single beliefs percolating across the agents of a social network (Hegselmann and Krause 2002, 2006, 2009; Olsson 2011; Grim *et al.* 2013; Vallinder and Olsson 2014). Coherence theory, in contrast, is structured essentially in terms of internal relations within sets of multiple beliefs. A model that incorporates coherence as well as correspondence will call for networked agents envisaged in terms of *suites* of beliefs adapting internally as well as in contact across networks. Our guiding intuition, substantiated in the results that follow, is that attention to the coherence dynamics within belief suites can be expected to make an important difference in the social dynamics of belief.

At a number of points we attend to available psychological work as a gauge to how realistic our model assumptions are. The goal is not a model designed to match psychological data, however. What we are after is a better grasp of the epistemological effects of basic cognitive mechanisms: differences in dynamics and outcomes given different parameters for coherence and correspondence across networks. Though our long-range goals are ultimately both descriptive and normative, both aspects remain here on the abstract plane of an attempt to understand general dynamic principles.

How do the forces of coherence and correspondence affect the dynamics of belief change across networks of believers? In section 1 we outline a family of network belief suite models, gauging the impact of coherence pressures alone in section 2. In section 3 we 'trickle in' elements of correspondence truth in limited sections of belief suites, tracking effects on both average coherence and over-all truth across the network. Section 4 highlights important temporal effects: results vary greatly between (a) a scenario in which coherence and correspondence mechanisms are introduced together and (b) a scenario in which correspondence input is added only after a measure of coherence has been achieved.

1. COHERENCE: A SIMPLE NETWORK MODEL

We incorporate two basic features into our model of belief dynamics: *networks* of agents and *suites* of beliefs. Much of the work in belief revision in a logical tradition (Alchourrón *et al.* 1985; Gärdenfors 1988) is geared to this normative question: what change in a single agent's belief set is rational in the face of new evidence? Related work on database updates in artificial intelligence can be interpreted in precisely the same light (Fagin *et al.* 1983; Rodrigues *et al.* 2011). Descriptive dynamics of belief across networks of multiple agents is not yet part of those traditions. On the other hand, work on belief dynamics that explicitly addresses descriptive questions of what can be expected from networks of agents with a particular structure has almost always dealt with what is thought of as a single belief across the network. Change in that single belief may be modeled either as a binary 'yes-no' or as a continuous value (Hegselmann and Krause 2002, 2006, 2009; Olsson 2011; Grim *et al.* 2013, 2015; Vallinder and Olsson 2014; interesting exceptions are Riegler and Douven 2009 and especially Axelrod 1997). In contrast to both of these lines of work, but capturing simple aspects of each, the model used here is explicitly designed to explore the role of coherence and correspondence among *suites* of beliefs of agents interacting in structured social *networks* of epistemic contact.

We model suites of beliefs for our agents as ordered sets of five values between 0 and 9. A single belief suite might therefore be 25891, for example, or 44662. Our aim is an abstract and minimal model, with the benefit of broad applicability that comes at the cost of low specificity. There is nothing in the model that specifies what the five values or their numerical range represents. The numerical values might be taken to stand for a range of estimated values for five scientific measurements. Alternatively, they might be taken to represent a range of opinions from 'strongly agree' to 'strongly disagree' on five political propositions.

A major focus of the work that follows is on the dynamics of belief coherence within suites of beliefs across a network. Despite a long tradition in coherence theories of truth, 'coherence' has been taken to mean very different things and has often been left underspecified (Ewing 1934; Lewis 1946; Price 1950; Blanshard 1964; Rescher 1973 [1982]). Logical consistency, linked probabilities, and explanatory connections might all be taken as elements of coherence, though a balance between different aspects or some further measure may be required as well (BonJour 1985; Lehrer 2000; Thagard 2000; Lycan 2012). Here again our model carries little commitment. We model a measure of coherence across a suite of beliefs without attempting to specify what that measure represents.

For our purposes, coherence within a suite of beliefs is measured in terms of numerical sameness across the set. The set 77777 will represent a perfectly coherent suite of beliefs, as will 22222: in each case all beliefs are at the same 'level.' Less than fully coherent belief suites will be those with uneven levels, recognizing here again that any of a number of very different assignments may have the same degree of (in)coherence. A similar modeling for consistency appears in Bednar *et al.* (2010).

We might use either of two variations on such a measure. The first is an *adjacent* measure, in which coherence is measured in terms of the closeness of each term to its adjacent neighbors. The adjacent *in*coherence of a set of beliefs is the sum of absolute differences between value 1 and 2, 2 and 3, 3 and 4, and 4 and 5. Neither 76567 nor 09090 are fully coherent sets in this sense, but the first has a much higher coherence measure than the second. Taking our values as v1 through v5, and normalized to a range in which 0 is the lowest value for adjacent coherence and 1 is the highest, our measure for adjacent coherence is $[1 - ((|v1 - v2| + |v2 - v3| + |v3 - v4| + |v4 - v|)/36)]$. On such a measure, the set of beliefs 76567 comes in with an adjacency coherence measure of 0.8888; 09090 has an adjacency coherence measure of 0.

For adjacent coherence, the order or a suite of beliefs will matter; the suite 09090 will have a radically different coherence value than will 00099. For modeling some aspects of human cognition such a measure might be entirely appropriate; some of our beliefs may be close enough that incoherence between them is easily noticed, with major pressure for alignment. Other beliefs may be far enough removed that incoherence is less noticeable, with less pressure for alignment. One's beliefs regarding policy toward Israel and policy toward Palestine may be closely related in this sense. Beliefs regarding policy toward Israel and toward North Korea might be less closely related, though in the long run a consistent foreign policy may demand coherence between these two as well.

The second measure of coherence we term *aggregate* coherence: the sum of pairwise differences between *any* two values in a belief suite. Normalized to a range between 0 and 1, with 0 as the lowest value and 1 as the highest, this second measure is determined by the sum of differences between value v1 and v2, v1 and v3, v1 and v4, v1 and v5, and so forth. Aggregate coherence is calculated by subtracting from 1 the sum of all of the pairwise differences divided by 54 – the maximum absolute incoherence that a belief set can possess. Here 76567 comes in with an aggregate coherence measure of 0.8148. 09090 comes in with an aggregate coherence of 0. Because order is not important here, 77665 and 00099 will have the same aggregate coherence as the rearrangements 76567 and 09090: 0.8148 and 0, respectively. Because a measure that is not tied to order is arguably more plausible as a general model for coherence, it is aggregate coherence that will play a primary role in much of the work that follows.

Although the model is formal and abstract, the intention is to incorporate a simple mechanism that mirrors some very real psychological forces toward conceptual coherence in a social environment (Thagard and Kunda 1998; Gawronski and Strack 2012). A core element of coherence, we have noted, is consistency. A long tradition of work on cognitive dissonance details the mental stress on inconsistent beliefs and the efforts made to alleviate or avoid that stress (Festinger 1957; Cooper 2007; Gawronski and Strack 2012).[1] In practice, when given a new input that contradicts a held belief, people choose to update their beliefs in a way that will preserve consistency (Walsh and Sloman 2004). Input could be

---

[1] One criticism of post-cognitive dissonance studies is in interpretation of results. Perhaps what is indicated is not that individuals have altered previous beliefs or preferences but that making a choice may help solidify previously ambiguous preferences (Chen and Risen 2010). Although resolution of ambiguity is not part of the current model, we believe this too can be interpreted as a move toward consistency – at least toward explicit consistency of beliefs.

new empirical observation, but could also reflect the social environment modeled in our networks: interaction with people who hold similar beliefs can strengthen an agent's confidence in a belief, while interaction with those who hold opposing beliefs can weaken confidence (DeMarzon *et al.* 2003). Low-balling is a manipulative sales technique which exploits both social pressures toward coherence. A customer is led to make a public commitment which can then be used as a consistency lever, leading him to accept a deal at a higher price (Guégen *et al.* 2002; Burger and Cornelius 2003; Cialdini and Goldstein 2004).

One element of our simple model is coherence within belief suites of individuals, influenced by social contacts. A second aspect is the structure of contact between individuals, represented as links between nodes in a network, and the mechanism of belief updating. In what follows we consider a range of possible network structures, starting with random networks of different mean degrees – the average number of links to each agent or node. Individuals update on their contacts by choosing a random element of their belief suite – the value in one of their five slots – and comparing their current over-all coherence with the coherence they would obtain were they to adopt their neighbor's value in that slot. They change their belief value in that slot to that of their neighbor if it represents an improvement in belief consistency over all.

With this form of updating, an individual with belief suite 88288, comparing the randomly chosen third belief with a neighboring agent's third belief in the suite 99799, will change his belief suite to 88788, thereby increasing his internal coherence on either an adjacent or aggregate measure. An individual with belief suite 88288, comparing a randomly chosen second belief with that of a neighboring agent 77777, will not change an 8 to a 7 since that would decrease his coherence on both of our measures.

All updating is done for agents asynchronously in a random order. In other words, in a given round it might be agent 25 who updates first, followed by 14, 2, 35 … Once agent 25 has updated, any agents who interact with 25 are then comparing belief suites with agent 25's newly updated beliefs. Simultaneous as opposed to asynchronous updating would demand instead that updating be done at discrete ticks of a global clock, with a set of beliefs 'frozen' for all agents at the beginning of each updating round. Classic work with cellular automata demonstrates that simultaneous as opposed to asynchronous updating can bias results unrealistically, a problem that can be expected to carry over to updating on network structures in general (Huberman and Glance 1993). In virtually all social systems, agents act and influence each not at discrete intervals and in synch but in uncorrelated series of randomly ordered influence, a pattern better captured by the asynchronous updating we employ throughout.

2. THE DYNAMICS OF COHERENCE

We begin with random networks of 50 nodes, in which each agent starts with a belief suite composed of five slots, each containing a random digit from 0 to 9. We follow the updating algorithm above, having each agent update on a random belief with a random neighbor on each round. Where updating is geared to adjacent coherence, individuals change their belief value to that of their neighbor if such a change would increase their adjacent coherence, reducing the total absolute distance between neighboring beliefs within their suite. Where updating is geared to aggregate coherence, individuals change the targeted belief to that of their neighbor if such a change would increase aggregate coherence, decreasing the total absolute distance between all beliefs taken pairwise.

Figure 1 shows results after 500 rounds for 10 runs of increasingly dense networks when agents update in terms of adjacent coherence. For each pair of nodes, edge probability measures the probability of a connection between them. For an edge probability of 20% or higher, it's clear that adjacent coherence peaks at a value of approximately 0.94. This is the same as the average value in runs of a complete network, in which all nodes are connected to all other nodes.

Figure 2 shows results for 10 runs of increasingly dense networks when agents update in terms of aggregate coherence. Here final coherence is notably higher – results for a complete network come in at approximately 0.98 – though results hit 0.96 at an edge probability of 20% and are essentially stable above an edge probability of 45%.

It is perhaps not surprising that agents maximizing individual aggregate coherence produce a higher average measure than those operating in terms of adjacent coherence. An agent updating in terms of adjacent coherence will change a given belief only if there is some improvement measured in terms of that belief and the two adjacent to it in the suite. An agent updating on another in terms of aggregate coherence will change that belief given any improvement in terms of pairwise measures with any of the other four beliefs. Updating in terms of aggregate coherence is thus more sensitive to possibilities for improvement across an entire suite of beliefs.

What may be surprising is that in neither case do we achieve full average coherence across a network nor convergence to a single belief suite. Not even when network density reaches that of a complete network does coherence-driven updating result in a single belief suite across the population. Figure 3 shows the final number of distinct belief suites in each case, once again measured in terms of increasingly dense random networks. For adjacent coherence, the density of a random networks makes almost no difference: slightly less than 40 belief suites remain after 500 iterations regardless of the probability of connection between pairs of nodes. Runs of a complete network give us an average of 37 distinct belief suites.

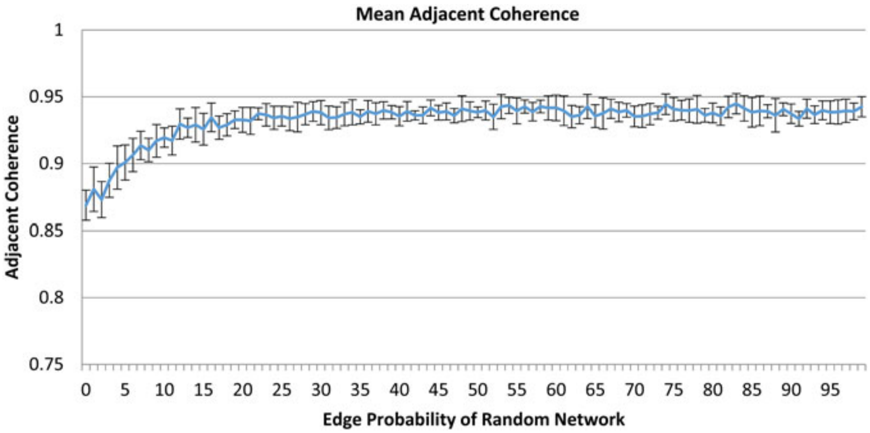

Fig. 1. Adjacent coherence of increasingly dense random networks after 500 rounds. Bars mark standard deviation.

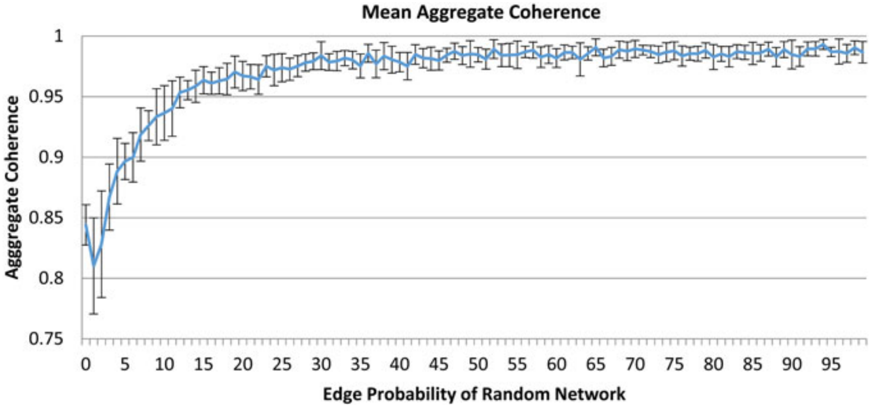

Fig. 2. Aggregate coherence of increasingly dense random networks after 500 rounds. Bars mark standard deviation.

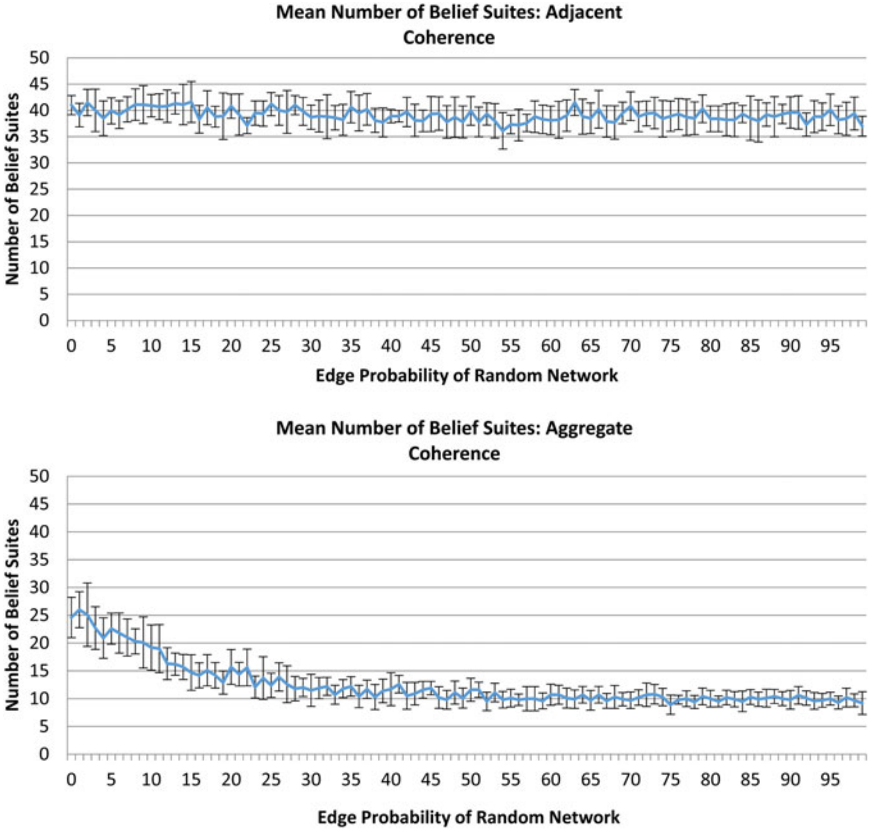

Fig. 3. Final number of belief suites for each form of coherence.

The pattern for aggregate coherence shows a different result, with lower final numbers of belief suites and a major decline with increasing density of the network. Above a probability of 35% for links between node pairs, the number of belief suites after 500 iterations settles to approximately 10. Average results for complete networks come in at 9.2 distinct belief suites.

Although time to convergence is not our main focus here, it should be noted network density does make a difference in terms of the speed with which a network settles into its final belief suites. For adjacent coherence, at an edge probability of 25%, the average number of generations until numbers of belief suites fall below 45 comes in at 51.09 over 100 runs, excluding three cases in which belief suites remained above 45 through 500 generations. At an edge probability of 65%, the average is 50.74 over 100 runs, excluding a single case in which belief suites remained above 45 through 500 generations. Standard deviation for adjacent coherence is high in each case, however: 23.98 at edge probability 25%, 26.86 at edge probability 65%. Aggregate coherence shows a similarly small time difference, with a much tighter standard deviation. At edge probability 25% it takes an average of 22.87 generations for belief suites to fall below 20 for aggregate coherence, with standard deviation 4.76. At edge probability 65% the average is 22.32, with the same standard deviation.

The reason our networks neither achieve complete coherence nor settle on a single belief suite becomes clear when we consider the final belief suites to which networks converge.

For a typical run of aggregate coherence on a network with edge probability of 95% – effectively a complete network – we might end up with the following 9 belief suites:

33323 11122 77777 55555 88887 22222 33343 66666 44444

Agents with none of these, in contact with any other, will increase coherence by updating any single belief. Consider the simplest cases:

77777 55555 22222 44444

Each of these is perfectly coherent by our measure – there is no absolute distance between any pair of beliefs – and thus none will update on any others. Consider also cases such as 33323 and 11122. The only change that would make the first suite more coherent would be changing a fourth ‘2’ to a 3. But none of the other belief suites listed above has a 3 in that position. Because we are updating beliefs one at a time, in comparison with beliefs in the same position, none of the other suites allows a change to greater coherence. For less dense networks the lack of connection between nodes can be expected to result in even lower aggregate coherence.

The higher number of final belief suites for adjacent coherence is understandable in similar terms. Among the many final belief suites in a network set generated through adjacent coherence might be the following:

44999 44666 55666 23444 44444 22111 44422 22333

The suite 44444, of course, is perfectly adjacent coherent and will not change. 44999 will not adopt a 4 in its third position, as might be suggested by interaction with 44444, because it would then reduce the distance between its second and third place from 5 to 0, but would increase its distance between its third and fourth place from 0 to 5. The result would be no improvement in total adjacent coherence. 44666 will not adopt a 4 in its third place from interaction with 44422, because its total adjacent incoherence of 2 would merely be shifted from the difference in its second and third position to the difference in its third and fourth – again no improvement in total adjacent coherence. A similar analysis applies to other pairs.

The message in all of this is that the dynamics and final result of coherence as a force for belief change across networks is more complex than might be supposed. Here we have emphasized different network degrees and measures of coherence, but it is clear that specific network structure and order of updating can also make a difference on final outcomes. Belief development on a network can be expected to be path dependent (Page 2006; Boas 2007).

Perhaps the clearest lesson is a very simple one. Different belief suites can each be internally coherent. Because of that, the dynamics of belief updating on even complete networks may result in high average coherence but not in a single consistent belief suite across the population. Despite strong network linkage and a dynamics driven by coherence, a significant number of internally coherent but distinct and mutually incompatible belief suites may remain (see also Feldman and Warfield 2010; Garbayo 2014).

Despite the differences noted, the central qualitative lesson of incomplete convergence holds for both adjacent and aggregate measures of coherence. In order to maintain simplicity when we add in correspondence, we will focus on aggregate coherence as our measure.

3. THE EFFECT OF CORRESPONDENCE: TRICKLING IN TRUTH

The only measure applied to our belief suites at this point has been coherence: a suite is coherent to the extent that the difference between numerical values on our five slots is minimized, either in linear pairwise or full pairwise comparison. But we can also use the features of the model to add something like correspondence to the picture

Here we emphasize the linear ordering of values in each belief slot, from 0 to 9. We will set a value of 9 as ranking higher than 8 on correspondence, 8 higher than 7, and so on. With this interpretation we can read each of our beliefs as having either a higher or lower truth value or as representing a more or less accurate approximation to objective fact. We have noted that we need have no commitment to what ‘coherence’ means beyond the model. We can similarly allow for alternative interpretations of our correspondence measure. For present purposes we will speak simply of ‘truth,’ though ‘empirical adequacy’ or even ‘pragmatic efficacy’ might be offered as equally legitimate readings.

A belief suite will be judged true to the extent that it approaches 99999. Each belief in the suite is untrue to the extent it differs from 9, with the sum of those differences a measure of untruth for the suite as a whole. That summed value is then normalized to a range between 0 and 1. Although coherence and truth are distinct, in the sense that the summed distance between spots in a belief suite is distinct from the level between 0 and 9 of those spots, there is one formal relation between coherence and truth that is built into the formal structure of the model. Although a fully coherent belief suite need have no particular relation to truth (it may, after all, be 00000), a belief suite each member of which is fully true will also necessarily be fully coherent (99999).

As noted in the introduction, coherence and correspondence are classically presented as rival theories of the nature of truth. Here we want to explore their interrelation as contributing factors in belief dynamics. We focus on networks generated with an edge probability of 10% because the graphs of the previous section indicate 10% as a convenient middle case for coherence effects. We start by 'trickling in truth' at just one spot. We suppose that just one of the elements of our belief suites is directly susceptible to influence by something like correspondence; just one element of a belief suite is 'empirically sensitive.' Will pressures for coherence carry over the effect of that one sensitive area to the rest of an individual's beliefs? For the sake of simplicity, and because it offers an arguably more general model for belief suite coherence, we focus on results for aggregate coherence. Qualitative trends for adjacent coherence follow much the same pattern.

It is worth noting aspects of realism in even this very abstract model. The psychological literature makes it clear that there is pressure for cognitive coherence given input that conflicts with held beliefs (DeMarzon *et al.* 2003). Input at one point in a web of beliefs will influence the rest as individuals try to make their other beliefs cohere with the new data (Walsh and Sloman 2004; Cooper 2007). It is the effect of pressures for cognitive coherence given small percentages of external 'correspondence' data – a trickling in of truth – that we attempt to model here across social networks.

There is one further measure we build into the model: 'incoherence tolerance,' measuring the extent to which individuals are willing to accept new information that does not cohere with existing beliefs. It is clear that the obviousness of a demonstration, the 'hardness' of facts, or the undeniability of the 'given' (Lewis 1929) can overpower the comfort of existing coherence. Psychological data shows evidence of a "truthfulness bias"; statements made by others are treated *prima facie* as true (Kam 2012). But it is also clear that the drive for coherence can lead an individual to question, ignore, or reinterpret data that conflicts with existing beliefs. The classic work of Lord, Ross, and Lepper claims two effects of exposure to mixed bodies of evidence: (a) attitude polarization and (b) biased assimilation (Lord *et al.* 1979). Attitude polarization is the tendency of people exposed to mixed evidence to strengthen their confidence in initial beliefs regarding the death penalty, for example, regardless of whether those initial beliefs were pro- or contra-. Biased assimilation is the tendency to devalue or reject presented information that is in conflict with current beliefs, while easily accepting and rating more favorably information that supports current beliefs. Further psychological work has qualified conclusions regarding attitude polarization (Kuhn and Lao 1996). On the other hand, numerous studies have made it clear that biased assimilation is a robust effect (Miller *et al.* 1993; Kuhn and Lao 1996; Kahan *et al.* 2012). Resistance to new information given conflicting background beliefs also forms a major part of Zaller's Receive-Accept-Sample model of belief change. Although acceptance of new information remains the default, those with prior beliefs in conflict with new input will reject that input in order to maintain the coherence of previous beliefs (Zaller 1992).

It is clear that there are different degrees to which incoherence within a belief suite may be tolerated for the sake of new input. For some individuals, or in some cases, new information may be accepted despite a significant initial clash with background beliefs. In other cases, new information may be resisted given even a slight conflict with beliefs already held (Russel and Jones 1980; McRae 1996; Onraet *et al.* 2011; Schwartzstein 2014). When exposed to empirical data that would generate incoherence if added to one's current beliefs, how willing is one to adopt that data despite the resulting incoherence? Eventually, other beliefs in a suite can be adjusted to regain coherence. But when there is trade-off between coherence across beliefs and empirical adequacy in a specific area, how much of the former is one willing to sacrifice for the latter?

'Trickling in truth' at just one empirically sensitive spot, spot 5, in belief suites, we incorporate 'incoherence tolerance' as a parameter of where individuals will adopt a 'truer' belief in slot 5 despite compromising their aggregate coherence. In some cases we can suppose that an individual with an aggregate coherence measure of 0.9 will opt for truth – a 'truer' value in slot 5 from a network neighbor – but only if adoption of that value in that slot will not drop his coherence measure below 0.8 overall. Since our coherence measures range from 0 to 1, we can think of a drop from 0.9 to 0.8 as a 0.1 drop in coherence, and can measure the effects of truth in a network of agents willing to make that much of a compromise in coherence but no more. In other cases we gauge the effect of an incoherence tolerance of 0.2, 0.3, or more; a 'truer' belief in slot 5 will be adopted even if the incoherence compromise across beliefs is severe. The greatest change in coherence resulting from the adoption of 9 in a single slot, it turns out, is the change from 0 0 0 0 0 to 0 0 0 0 9, representing an aggregate coherence change of 0.66.

In this version of the model, our agents adopt a neighbor's belief in a randomly selected slot if that increases their coherence, but also adopt a neighbor's belief in a particular slot – an empirically sensitive slot 5 – as long as the adoption does not reduce coherence beyond a set tolerance level. For different incoherence tolerances, Figure 4 shows the effect on: (a) average aggregate coherence of agents, (b) average aggregate truth, and (c) the number of distinct belief suites after 500 iterations within such a model.

In the charts of the previous section, random networks with a 10% edge probability, operating in terms of coherence alone, generated an aggregate coherence of approximately 0.94. The same result appears here as approximately 0.94 coherence in a network of agents who have an 'incoherence tolerance' approaching 0 (i.e. agents who are not willing to compromise coherence for the sake of truth at all). Given any willingness to compromise coherence in order to pursue truth, coherence itself drops, reaching a low point that is still at 0.77 for incoherence tolerance levels of 0.3 or greater. At that same tolerance agents achieve an average truth of only approximately 0.75, unable to score higher even at higher rates of incoherence tolerance. Although coherence pressures will push other beliefs to higher truth levels, truth trickled in at just one position does not take over entirely – trickled in truth does not force belief suites to 99999, a result that holds even if agents tolerate any level of incoherence for the sake of truth.

For random networks of the degree indicated, acceptance of 'trickled-in' truth at a single position in a belief suite comes at the price of reduced coherence but does not achieve truth across belief suites. This contrasts dramatically with the 'funnel theorem' and results from Hegselmann and Krause (2006), which show a strong tendency for group convergence on the truth given even a small tendency toward the truth among individuals. The major reasons for the difference are the two characteristics of the current study that distinguish it from its predecessors: (a) the fact that random networks are being used here, in contrast with the effectively complete networks of Hegselmann and Krause, and (b) the fact that we are using belief *suites* rather than single beliefs that characterize individual agents. Results here indicate that both the complexities of internal conceptual organization represented by belief suites and the complexities of social

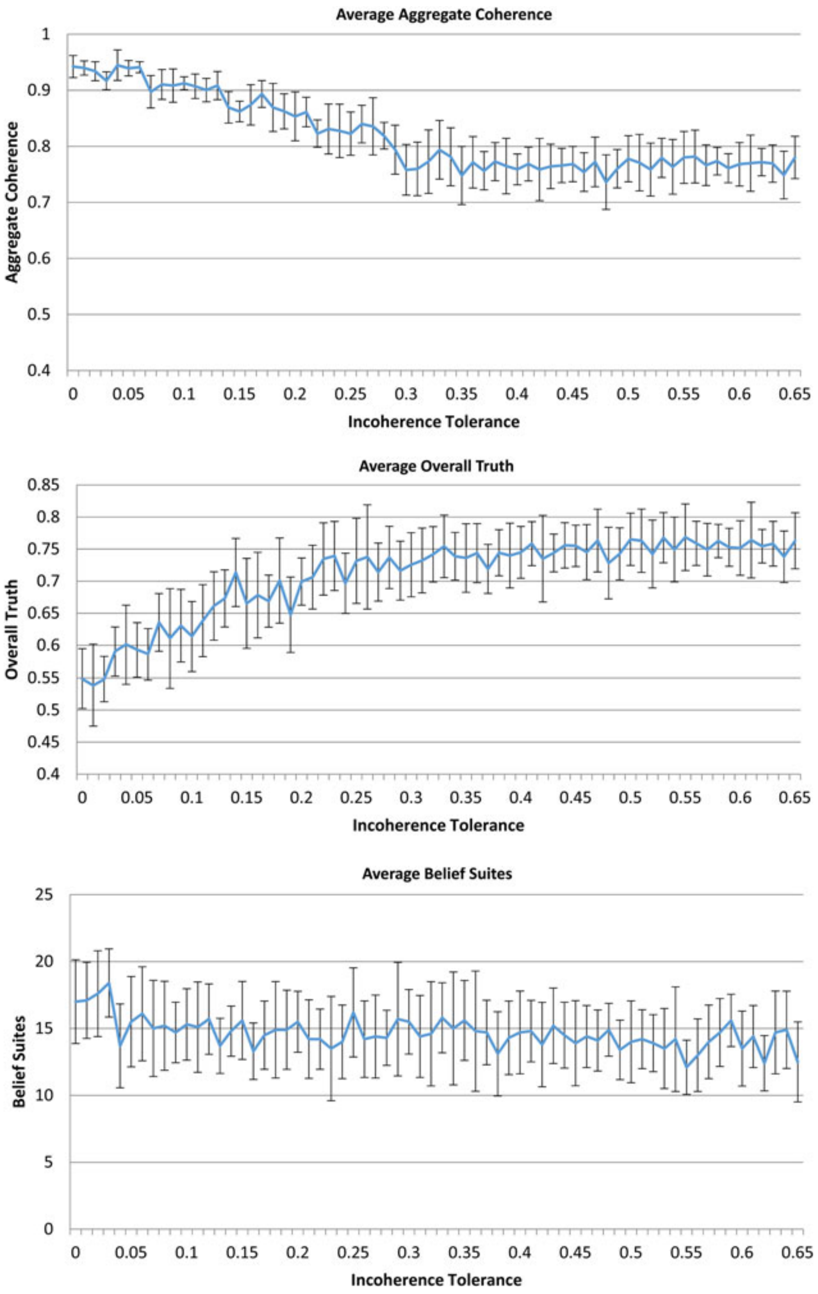

Fig. 4. The effects of truth 'trickled in' in slot 5 for networks with a 10% edge probability. Agents switch to a 'truer' option in slot 5 if their incoherence tolerance is within the bounds indicated on the x-axis. The greatest possible coherence change with a new belief in one position is 0.66.

communication represented by random networks can block the path to truth. Both effects are evident in the psychological data (Taber and Lodge 2006; Lodge and Taber 2013; Kraft *et al.* 2015).

Time differences at different incoherence tolerances are evident but relatively small. At an incoherence tolerance of 0.25, for example, over 100 runs it takes networks with truth trickled in at a single position an average of 63.01 generations to fall below 22 distinct belief suites, with a standard deviation of 18.70. At an incoherence tolerance of 0.65, the average is 61.54 with a standard deviation of 22.05.

What if agents' beliefs are more open to 'correspondence' data: what if truth is trickled in at more than one position? The graphs in Figure 5 show the more optimistic results in that case.

When belief suites are sensitive to truth in two positions rather than one, their overall truth achievement increases dramatically. With an incoherence tolerance of 0.3 or higher, average truth comes very close to 1. One might think that such an increase in truth is purchased at the price of coherence. But for any incoherence tolerance at that level or higher the value for aggregate coherence comes in at 0.94, essentially the same as the pure coherence result for networks of this degree. When truth is trickled in at two positions rather than one, coherence can be as high as when they seek coherence alone. Empirical sensitivity at a single position in a belief suite does not preserve coherence and does not achieve full truth. But empirical sensitivity at two of five – a minority of positions within a belief suite – obtains both equivalent coherence and near unanimity on truth.[2]

The biggest surprise, perhaps, is a significant drop in the number of distinct belief suites across a network when truth is trickled in at two positions. With coherence alone, the number of distinct belief suites for networks of this density is approximately 20. With truth sensitivity at one position, the number of belief suites drops to 14. With truth sensitivity in two positions, that number drops to merely 4, representing a network of near unanimity across agents with both high internal coherency and a high level of overall truth.

Although speed is not our main focus, time differences are again worthy of note. For an incoherence tolerance of 0.25, over 100 runs it takes networks with truth trickled in at two position an average of 45.96 generations to fall below 22 distinct belief suites, with a standard deviation of 16.79. At an incoherence tolerance of 0.65, the average is 40.09 with a standard deviation of 8.08.

4. COHERENCE AND CORRESPONDENCE: THE TEMPORAL EFFECT

In the studies above, we have always begun with a random array of beliefs, introducing both forces toward coherence and forces toward a 'trickled in' correspondence value *simultaneously*. But what if these forces do not act simultaneously on an initial random array

---

[2] In these runs, agents adopted the belief in slot 5 of a neighboring agent with a higher 'truth-value' as long as that did not compromise their coherence beyond the threshold noted. Movement was total and immediate: within bounds of coherence, an agent would adopt a full 9 of a neighbor. We also ran tests in which agents only went halfway toward a truer belief in slot 5 on a single turn. Results for both one-slot trickled in truth and two-slots, though somewhat slower, seemed perfectly consistent with these results in final form.

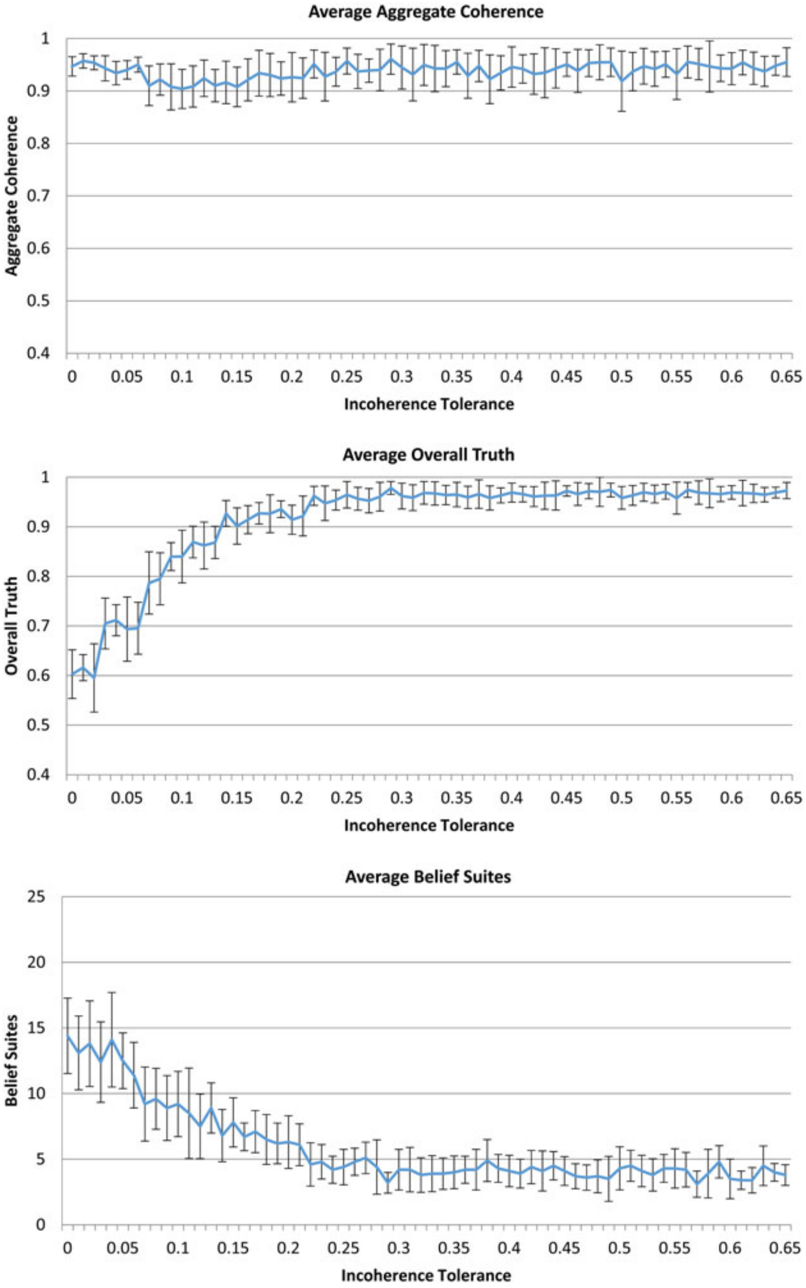

Fig. 5. The effects of truth 'trickled in' in slots 4 and 5 for networks with a 10% edge probability. Agents switch fully to a 'truer' option in slots 4 or 5 if their incoherence tolerance is within the bounds indicated on the x-axis. The greatest possible coherence change with a new belief in one position is 0.66.

of beliefs? A more realistic scenario may be one in which new information is trickled into a community in which individuals have already established some measure of belief coherence. What happens if coherence is established first, with truth trickled in later?

The result in this case is quite different than that presented above. Figure 6 shows a standard pattern, exhibited for a network with 10% edge probability and using an incoherence tolerance at the extreme of 0.66. Agents operate solely in terms of coherence up to generation 100, updating on network contacts so as to increase their own coherence. Only at iteration 100 are agents' updating rules changed so that they become truth-sensitive in slot 5. Here we show development over 250 generations.

The final result is quite different than when coherence and correspondence operate together from initial random belief suites. In this case, average truth for agents begins to rise at the point it is trickled in, but only achieves an overall average of approximately 0.57, as opposed to approximately 0.75 when coherence and correspondence operate together. Overall coherence, although initially rising above 0.9, drops to approximately 0.7, as opposed to the value of 0.77 achieved when the two operate together. Over 1000 runs, the average final values for aggregate truth and coherence are 0.60 and 0.69, respectively.

Figure 7 shows the result of trickling in truth with a similar time delay but in both slots 4 and 5.

When both coherence and correspondence operate from the beginning, truth achieves a level of approximately 0.98. Here it reaches only approximately 0.72. Coherence initially rises to slightly above 0.90, but falls to below 0.60 with the introduction of truth. It then rises as truth does, but still reaches a value of only 0.70. Over 1000 runs, average truth reaches 0.75; average coherence reaches 0.68.

Here the lesson regarding coherence and correspondence is less optimistic than above. Coherence can indeed benefit overall truth as a secondary force: when coherence and correspondence operate together from initially random belief suites, truth trickled into two slots of five can raise the overall truth across all slots to very close 1. When coherence is established first, however, it blocks the overall truth that can be obtained. Here again

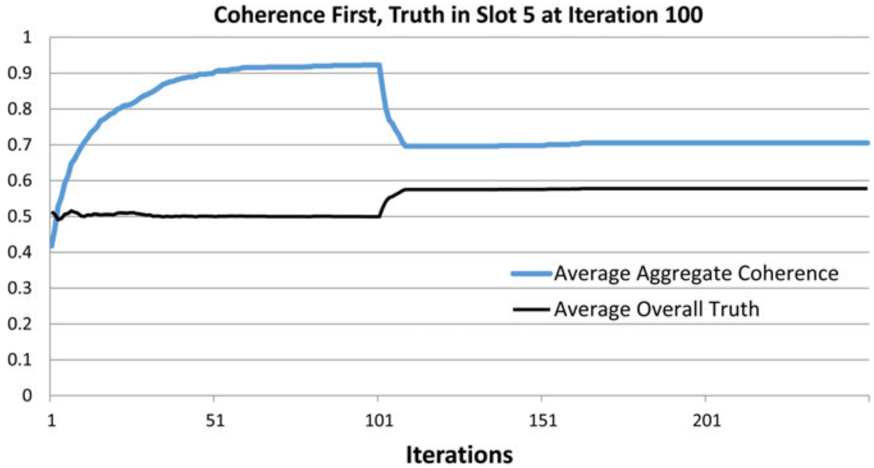

Fig. 6. The temporal effect with coherence established first and truth trickled in at slot 5 only after iteration 100. Results for a random network with 10% edge probability and incoherence tolerance 0.66.

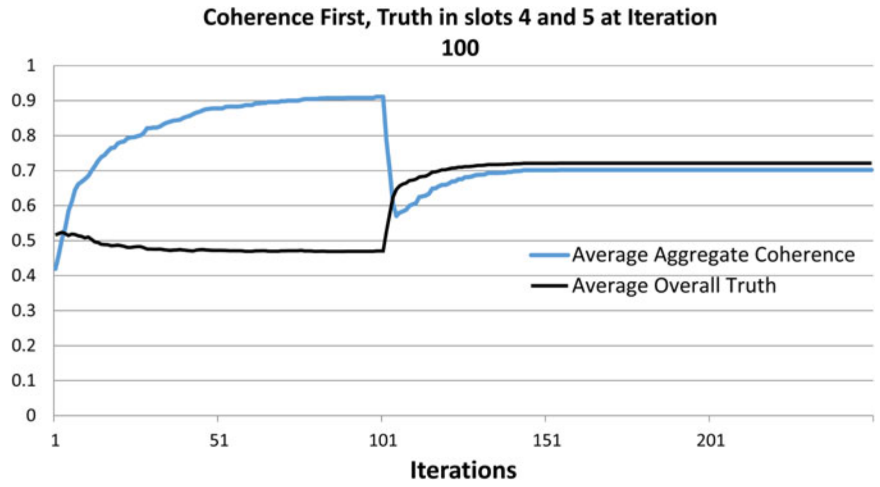

Fig. 7. The temporal effect with coherence established first and truth trickled in at slots 4 and 5 only after iteration 100. Results for a random network with 10% edge probability and incoherence tolerance 0.66.

our results are in contrast to earlier modeling results in which truth tends to triumph, such as Hegselmann and Krause (2006), in which truth-seeking and updating on neighbors are pursued simultaneously. Our results are very much in accord with psychological evidence indicating that established beliefs can block the ability to update on new input (Lord *et al.* 1979; Zaller 1992; Miller *et al.* 1993; Kuhn and Lao 1996; Kahan *et al.* 2012).

What if we reverse the order, establishing correspondence before introducing coherence mechanisms? Figure 8 shows the result of having all agents update solely on truer beliefs of neighbors in the single slot 5 up to iteration 100. Only at that point do agents also start to update in terms of coherence as well.

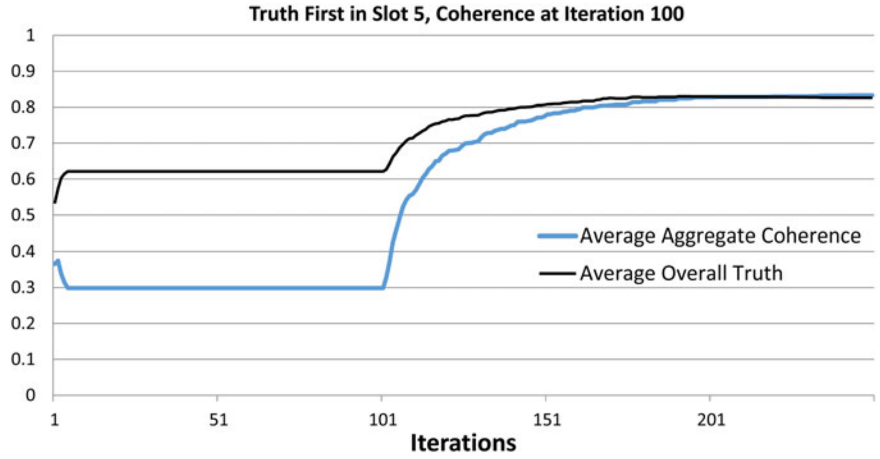

Fig. 8. The temporal effect with truth influence in slot 5 first and coherence updating activated only after iteration 100. Random network with 10% edge probability, incoherence tolerance 0.66.

Here results are importantly different. Truth-updating in a single slot alone brings average overall truth to only slightly above 0.60, which is not particularly surprising. When the force of coherence is added, both coherence measures and overall truth rise dramatically, and rise to levels significantly above those achieved when the order was reversed. In that case, truth-sensitivity in slot 5 resulted in a value of 0.57, far lower than the approximately 0.84 achieved here. Coherence reached a height of 0.70, significantly below the level here, also approximately 0.84. Over 1000 runs, the average for both truth and aggregate coherence are 0.79, slightly higher even than when both coherence and correspondence are in play from the beginning.

Figure 9 shows the results when truth-updating is done first in both slots 4 and 5, then followed by coherence considerations.

Truth consideration alone achieves a level of 0.68 or so, with coherence dropping to a low 0.30. When coherence updating mechanisms are added, both truth and coherence rise to approximately 0.98. Over 1000 runs, coherence averages 0.97, with truth at 0.98 – slightly better than when both operate together from the beginning.

The simple lesson is that the temporal effect of correspondence and coherence is asymmetric. If coherence is the first consideration, the results can negatively impact both truth and coherence achieved, even if truth is ‘trickled in’ at two of five slots in belief suites. The same does not hold if truth is the first consideration. There we again see the mutually reinforcing effect of the two operating together. Though only a minority of slots are directly truth-sensitive, the drive for coherence spreads truth across belief suites. The rise in overall truth brings with it a rise in overall aggregate coherence as well.

The optimal order, it appears, is an initial consideration of correspondence input, followed later by coherence adjustments across agents’ belief suites as a whole. Data first, theory second. That is the good news. The bad news is that neither research nor intuition indicates that information processing in this order is something that people are inclined to or even able to do.

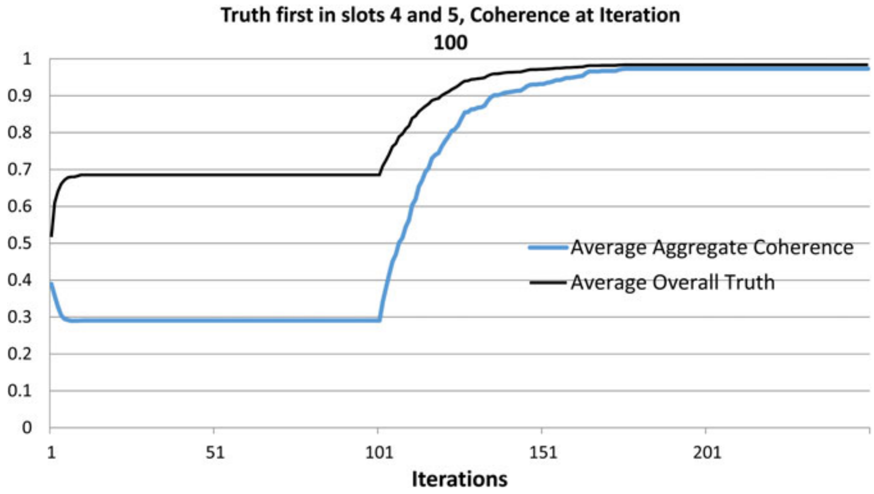

Fig. 9. The temporal effect with truth influence in slot 5 first and coherence up-dating activated only after iteration 100. Random network with 10% edge probability, 0.66 inconsistency tolerance.

5. CONCLUSION

Should coherence and correspondence be thought of as rival forces within belief dynamics?

With an eye to the dynamics of change in belief suites across networks, our results indicate that coherence and correspondence *can* in some cases be thought of as rivals. For belief suites that are sensitive only at a small number of positions – one in five slots in our model – it is only with compromise in coherence that truth as correspondence can be achieved. In such a case, there is also a limit to the level of truth across the belief suites that can be reached.

With an eye to the temporal introduction of coherence and correspondence the two can also be seen as competing forces. Specifically, a priority on the establishment of coherence can result in lower rates both for truth acquisition and for coherence itself.

But coherence and correspondence do not function as rival forces in all cases. For belief suites that are truth sensitive at a larger number of positions – two of five slots in our model – a very high level of truth can be achieved with effectively no compromise in coherence. Even when agents are willing to sacrifice a great deal of coherence for the sake of trickled in truth, coherence is not ultimately compromised and in a broad range of cases is clearly enhanced. If the goal is correspondence, then, coherence is in many cases an important means rather than a competitive threat. Even truth trickled in at a minority of spots in belief suites, and incorporated only within a relatively small range of incoherence tolerance, can result in a very high level of truth across networked belief suites as a whole.

The temporal introduction of the two can also show this effect. A priority on the establishment of correspondence, followed with mechanisms of coherence, can result in both high coherence and high correspondence.

We consider our model an early step toward tracking correspondence and coherence as aspects of network belief dynamics. Results, though perhaps intuitive, are here instantiated with the formalism of agent-based modeling; in the search for truth, there must be some element of correspondence, and that element must have priority. Given a small influence of correspondence on belief suites, coherence can play a major role, raising correspondence with it across a configuration of linked beliefs, both individually and within a social epistemic community. What this opens up, we hope, are further prospects for computational modeling targeting the social dynamics of change within belief suites across social contact networks.

REFERENCES

Alchourrón, C. E., Gärdenfors, P. and Makinson, D. 1985. 'On the Logic of Theory Change: Partial Meet Contraction and Revision Functions.' *Journal of Symbolic Logic*, 50: 510–30.

Axelrod, R. 1997. 'The Dissemination of Culture: A Model with Local Convergence and Global Polarization.' *Journal of Conflict Resolution*, 41: 203–26.

Bednar, J., Bramson, A., Jones-Rooy, A. and Page, S. 2010. 'Emergent Cultural Signatures and Persistent Diversity: A Model of Conformity and Consistency.' *Rationality and Society*, 22: 407–44.

Blanshard, B. 1964. *Reason and Analysis*. La Salle, IL: Open Court.

Boas, T. C. 2007. 'Conceptualizing Continuity and Change: The Composite-Standard Model of Path Dependence.' *Journal of Theoretical Politics*, 19: 33–54.

BonJour, L. 1985. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.

Burger, J. M. and Cornelius, T. 2003. 'Raising the Price of Agreement: Public Commitment and the Lowball Compliance Procedure.' *Journal for Applied Social Psychology*, 33: 923–34.

Chen, M. K. and Risen, J. L. 2010. 'How Choice Affects and Reflects Preferences: Revisiting the Free-choice Paradigm.' *Journal of Personality and Social Psychology*, 99: 573–94.

Cialdini, R. B. and Goldstein, N. J. 2004. 'Social Influence: Compliance and Conformity.' *Annual Review of Psychology*, 55: 591–621.

Cooper, J. M. 2007. *Cognitive Dissonance: Fifty Years of a Classic Theory*. London: Sage.

Davidson, D. 1986. 'A Coherence Theory of Knowledge and Truth.' In E. Lepore (ed.), *Truth and Interpretation*, pp. 307–19. Oxford: Blackwell.

DeMarzon, P. M., Zwiebel, J. and Vayanos, D. 2003. 'Persuasion Bias, Social Influence, and Unidimensional Opinions.' *Quarterly Journal of Economics*, 118: 909–68.

Ewing, A. C. 1934. *Idealism: A Critical Survey*. London: Methuen.

Fagin, R., Ullman, J. D. and Vardi, M. Y. 1983. 'On the Semantics of Updates in Databases.' In *Proceedings of the ACM SIGACT-SIGMOD Symposium on Principles of Database Systems*. Atlanta, GA: ACM.

Feldman, R. and Warfield, T. A. (eds) 2010. *Disagreement*. New York, NY: Oxford University Press.

Festinger, L. 1957. *A Theory of Cognitive Dissonance*. Stanford, CA: Stanford University Press.

Garbayo, L. 2014. 'Epistemic Considerations on Expert Disagreement, Normative Justification, and Inconsistency Regarding Multi-criteria Decision Making.' In M. Ceberio and V. Kreinovich (eds), *Constraint Programming and Decision Making*, pp. 35–45. New York, NY: Springer.

Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Cambridge, MA: MIT Press.

Gawronski, B. and Strack, F. (eds) 2012. *Cognitive Consistency: A Fundamental Principle in Social Cognition*. New York, NY: Guilford Press.

Grim, P., Singer, D. J., Fisher, S., Bramson, A., Berger, W. J., Reade, C., Flocken, C. and Sales, A. 2013. 'Scientific Networks on Data Landscapes: Question Difficulty, Epistemic Success, and Convergence.' *Episteme*, 10: 441–64.

—— and Fisher, S. 2015. 'Germs, Genes, & Memes: Functional and Fitness Dynamics on Information Networks.' *Philosophy of Science*, 82: 219–43.

Guégen, N., Pascal, A. and Diagot, L. 2002. 'Low-Ball and Compliance to a Request: An Application in a Field Setting.' *Psychological Reports*, 91: 81–4.

Haig, B. D. and Borsboom, D. 2012. 'Truth, Science, and Psychology.' *Theory and Psychology*, 22: 272–89.

Harman, G. 1986. *Change in View: Principles of Reasoning*. Cambridge, MA: MIT Press.

Hegselmann, R. and Krause, U. 2002. 'Opinion Dynamics and Bounded Confidence Models.' *Journal of Artificial Societies and Social Simulation*, 5. http://jasss.soc.surrey.ac.uk/5/3/2.html

—— and —— 2006. 'Truth and Cognitive Division of Labour: First Steps towards a Computer Aided Social Epistemology.' *Journal of Artificial Societies and Social Simulation*, 9. http://jasss.soc.surrey.ac.uk/9/3/10.html

—— and —— 2009. 'Deliberative Exchange, Truth, and Cognitive Division of Labour: A Low-Resolution Modeling Approach.' *Episteme*, 6: 13–144.

Hempel, C. 1935. 'On the Logical Positivists' Theory of Truth.' *Analysis*, 2: 49–59.

Huberman, B. A. and Glance, N. S. 1993. 'Evolutionary Games and Computer Simulations.' *Proceedings of the National Academy of Sciences USA*, 90: 7716–18.

Kahan, D. M., Peters, E., Wittlin, M., Slovic, P., Ouellette, L. L., Braman, D. and Mandel, G. 2012. 'The Polarizing Impact of Science Literacy and Numeracy on Perceived Climate Change Risks.' *Nature Climate Change*, 2: 732–5.

Kam, C. D. 2012. 'The Psychological Veracity of Zaller's Model.' *Critical Review*, 24: 547–67.

Kraft, P. W., Lodge, M. and Taber, C. S. 2015. 'Why People 'Don't Trust the Evidence': Motivated Reasoning and Scientific Beliefs.' *Annals of the American Academy of Political and Social Science*, 658: 121–33.

Kuhn, D. and Lao, J. 1996. 'Effects of Evidence on Attitudes: Is Polarization the Norm?' *Psychological Science*, 7: 115–20.

Lehrer, K. 2000. *Theory of Knowledge*, 2nd edition. Boulder, CO: Westview.

Lewis, C. I. 1929. *Mind and the World Order*. New York: Charles Scribner [Reprinted New York: Dover 1956].

—— 1946. *An Analysis of Knowledge and Valuation*. LaSalle, IL: Open Court.

Lodge, M. and Taber, C. S. 2013. *The Rationalizing Voter*. New York, NY: Cambridge University Press.

Lord, C. G., Ross, L. and Lepper, M. 1979. 'Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence.' *Journal of Personality and Social Psychology*, 37: 2098–109.

Lycan, W. G. 2012. 'Explanationist Rebuttals (Coherentism Defended Again).' *Southern Journal of Philosophy*, 50: 5–20.

McGinn, C. 2002. 'The Truth about Truth.' In R. Schantz (ed.), *What is Truth?*, pp. 194–204. Berlin: Walter de Gruyter.

McRae, R. R. 1996. 'Social Consequences of Experiential Openness.' *Psychological Bulletin*, 129: 323–37.

Miller, A. G., McHoskey, J. W., Bane, C. M. and Dowd, T. G. 1993. 'The Attitude Polarization Phenomenon: Role of Response Measure, Attitude, Extremity, and Behavioral Consequences of Reported Attitude Change.' *Journal of Personality and Social Psychology*, 64: 561–74.

Neurath, O. 1983. *Philosophical Papers 1913–1946*. R. S. Cohen and M. Neurath (eds). Dordrecht: D. Reidel.

Olsson, E. J. 2011. 'A Simulation Approach to Veritistic Social Epistemology.' *Episteme*, 8: 127–43.

Onraet, E., Van Heil, A., Roets, A. and Cornelis, I. 2011. 'The Closed Mind: 'Experience' and 'Cognition' Aspects of Openness to Experience and Need for Closure as Psychological Bases for Right-wing Attitudes.' *European Journal of Personality*, 25: 184–97.

Page, S. E. 2006. 'Essay: Path Dependence.' *Quarterly Journal of Political Science*, 1: 87–115.

Price, H. H. 1950. *Perception*. London: Methuen.

Rescher, N. 1973 [1982]. *The Coherence Theory of Truth*. New York, NY: Oxford University Press 1973; University Press of America 1982.

Riegler, A. and Douven, I. 2009. 'Extending the Hegselmann-Krause Model III: From Single Beliefs to Complex Belief States.' *Episteme*, 6: 145–63.

Rodrigues, O., Gabbay, D. and Russo, A. 2011. 'Belief Revision.' In D. Gabbay and F. Guenthner (eds), *Handbook of Philosophical Logic*, 2nd edition, vol. 16, pp. 1–114. Dordrecht: Springer.

Russel, D. and Jones, W. H. 1980. 'When Superstition Fails: Reactions to Disconfirmation of Paranormal Beliefs.' *Personality and Social Psychology Bulletin*, 6: 83–8.

Russell, B. 1910. 'On the Nature of Truth and Falsehood.' In *Philosophical Essays*. London: Routledge.

—— 1912. *The Problems of Philosophy*. New York: Oxford

Schwartzstein, J. 2014. 'Selective Attention and Learning.' *Journal of the European Economic Association*, 12: 1423–52.

Taber, C. S. and Lodge, M. 2006. 'Motivated Skepticism in the Evaluation of Political Beliefs.' *American Journal of Political Science*, 50: 755–69.

Thagard, P. 2000. *Coherence in Thought and Action*. Cambridge, MA: MIT Press.

—— 2007. 'Coherence, Truth, and the Development of Scientific Knowledge.' *Philosophy of Science*, 74: 28–47.

—— **and Kunda, Z.** 1998. 'Making Sense of People: Coherence Mechanisms.' In S. J. Read and L. C. Miller (eds), *Connectionist Models of Social Reasoning and Social Behavior*, pp. 3–26. New York, NY: Lawrence Erlbaum.

Vallinder, A. and Olsson, E. J. 2014. 'Trust and the Value of Overconfidence: A Bayesian Perspective on Social Network Communication.' *Synthese*, 191: 1991–2007.

Walsh, C. R. and Sloman, S. A. 2004. 'Revising Causal Beliefs.' In *Proceedings of the 26th Annual Conference of the Cognitive Science Society*, pp. 1423–7. Mahwah, NJ: Lawrence Erlbaum Associates.

Young, J. O. 2013. 'The Coherence Theory of Truth.' In E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, http://plato.stanford.edu/entries/truth-coherence/

Zaller, J. R. 1992. *The Nature and Origins of Mass Opinion*, pp. 40–52. New York, NY: Cambridge University Press.

The authors constitute a research team within the Center for Study of Complex Systems at the University of Michigan. Patrick Grim's home position is with the Department of Philosophy at Stony Brook. The other members of the group represent undergraduate strengths in Departments of Philosophy, Political Science, Economics, Sociology, Mathematics, and Computer Science at the University of Michigan.
