

Why Pursue Unification? A Social-Epistemological Puzzle

Randall HARP and Kareem KHALIFA

Received: 31/07/2014
Final Version: 29/04/2015

BIBLID 0495-4548(2015)30:3p.431-447
DOI: 10.1387/theoria.12647

ABSTRACT: Many have argued that unified theories ought to be pursued wherever possible. We deny this on the basis of social-epistemological and decision-theoretic considerations. Consequently, those seeking a more ubiquitous role for unification must either attend to the scientific community's social structure in greater detail than has been the case, and/or radically revise their conception of unification.

Keywords: unification, social epistemology, decision theory, science, theory pursuit.

RESUMEN: Son muchos los que han defendido que deberían buscarse teorías unificadas siempre que sea posible. Nosotros lo negamos a partir de consideraciones socio-epistemológicas y de teoría de la decisión. En consecuencia, aquellos que busquen un papel más omnipresente para la unificación han de prestar un mayor grado de atención a la estructura social de la comunidad científica del que se le ha prestado, y/o revisar radicalmente su concepción de la unificación.

Palabras clave: unificación, epistemología social, teoría de la decisión, ciencia, búsqueda de teorías.

1. Introduction

Unification is hotly debated in the philosophy of science. Much of the debate concerns its status as a desideratum of inquiry: is unification a fruit that all of our scientific pursuits ought to bear? Most parties in these debates assume that unification has been fruitful in some domains, and less fruitful in others. “Unificationists” claim that the latter inquiries should nevertheless treat unification as a desideratum. Non-reductionists and pluralists deny this claim¹.

Such debates rarely countenance the *social interactions* between scientists required to produce a scientific theory. This paper argues that, unless some important social-epistemological questions are answered, it can be rational to forgo the pursuit of unification, even when unification promises the greatest cognitive benefit. After setting the stage (Sections 2 and 3), we argue that a unified theory promising greater cognitive benefits but also requiring cooperation between scientists with different specializations can be less rational to pur-

¹ Compare: “Monists might admit that a plurality of approaches and models meet appropriate scientific standards (or satisfy the corresponding epistemic values) but insist that this is only because today's science is incomplete. But we do not believe that the plurality in today's science is necessarily a temporary state of affairs” Kellert, Stephen H., Helen E. Longino, and C. Kenneth Waters. 2006. *Scientific Pluralism*. Minneapolis, MN: University of Minnesota Press..



sue than a less unified theory promising lower cognitive benefits but requiring no such cooperation (Section 4).

While we provide no solution to this “puzzle,” our view entails that the value of unified theories has a distinctively social-epistemological element that has been largely overlooked. We regard our challenge as an invitation for unificationists (and interested social epistemologists) to examine the social conditions that would make unification a more robust desideratum of science. In this regard, our challenge is not intended as decisive, but rather as placing new burdens of proof upon those who would seek a more prominent role for unification.

2. Preliminaries About Theory Pursuit and Unification

We begin by clarifying four features of our account: theory-pursuit (§2.1), fields and research units (§2.2), unification (§2.3), and the concept of interfield ‘hooks’ and ‘latches’ (§2.4). These ideas will inform our argument in Section 4, where we present our social-epistemological concerns.

2.1. Theory Pursuit

In discussing unification as a desideratum or aim of scientific inquiry, we take the fundamental issue to be whether or not scientists should *pursue* unified theories, not whether they should *believe* such theories. Believing a theory means taking it as true; pursuing a theory means allocating resources towards its further development. Pursuit might involve testing the theory, refining its hypotheses, or applying for grants to undertake such tasks.

Belief and pursuit are distinct stances towards a theory. Scientists can believe theories without pursuing them, e.g. if they think the theory is true and requires no further development. Scientists can also pursue theories without believing them. This might happen simply because pursuit keeps them gainfully employed; less cynically, it may also occur if they are undecided about the veracity of a theory, and the pursuit promises to shed light on this issue. Our model is compatible with agents pursuing a theory while adopting any doxastic state towards that theory.

Consequently, the rational standards for believing a theory differ from the rational standards for pursuing a theory. Roughly stated, the standards of rational belief concern *evidence*, while the standards of rational pursuit concern *utility*. Of course, scientists derive high utility from evidential considerations, so pursuit does not proceed entirely independently of epistemic considerations. In the interests of generality, we will not take a substantive stance on how scientists should assign utility. Thus, readers are allowed to make these assignments as epistemically high-minded or grubbily pragmatic as they see fit; regardless of these considerations, our conclusion follows.

2.2. Fields and Research Units

By and large, it is fruitful to view the decision-makers in our model as slightly larger units of analysis than individual scientists; theories are pursued not only by individual researchers

but also by labs and teams of scientists. To that end, we will use Darden and Maull's (1977, 44) notion of a scientific field, defined as follows:

... a central problem, a domain consisting of items taken to be facts related to that problem, general explanatory factors and goals providing expectations as to how the problem is to be solved, techniques and methods, and, sometimes, but not always, concepts, laws and theories which are related to the problem and which attempt to realize the explanatory goals [... and often,] a special vocabulary.

Darden and Maull distinguish their conceptual account of fields from sociological or institutional accounts. So construed, fields are abstract entities, or at the very least, conceptual products that can be distinguished from the humans that produce them. Since we are interested in both the conceptual and the social aspects of fields, we will follow Darden and Maull's use of the term 'field' and introduce the term 'research unit' to denote the group of scientists associated with a field who engage in cooperative activities with respect to theory pursuit. These activities involve shared deliberation, shared values, or shared intentions. Throughout this paper, we use the terms 'scientist' and 'research unit' interchangeably. However, our interest is primarily with research units, which consist of scientists cooperating in the pursuit of theories, rather than with individual scientists per se.

Like individual scientists, research units pursue theories by testing those theories, refining hypotheses and models, applying for grants, allocating resources towards the development of those theories, etc. Darden and Maull's criteria of a field help to distinguish between a group of scientists who make up a single research unit and those who do not: members of a research unit have common problems, explanatory goals, techniques, etc. Additionally, certain social explanations will apply to research units, but not to groups of scientists that are not part of a common research unit. For instance, members of a research unit can function as cooperative parts of a whole; scientists who do not make up a research unit do not (though they might still engage in productive collaboration with other research units).

There are multiple ways of delimiting fields, e.g. by having finer- or coarser-grained ways of describing problems, explanatory goals, and techniques. Consequently, research units assume many different institutional shapes in scientific practice. They might consist of a laboratory team that distributes research labour among its members; a whole scientific field (e.g. population geneticists); or a professional organization (e.g. the American Physical Society's Division of Condensed Matter Physics). At the smallest extreme, a research unit might consist of a single scientist. Membership in research units is fluid, and scientists can be members of more than one research unit at a time. Moreover, a certain amount of imprecision is to be expected. How many research units are there, for example, in a large research organization like CERN? While there may be no precise answer to this, it seems clear that it is a far smaller number than the total number of individual scientists affiliated with the organization. While we will not conjecture any further about how research units are individuated, we assume that they track with however scientific communities divide themselves in the context of theory-pursuit, save for the simplifying assumption that an individual scientist is a member of exactly one research unit.

In summary, we will be focusing on research units' pursuit of unified theories. We will ultimately argue that, unless more social-epistemological work is done, it can be rational for research units to forgo this pursuit, even when unified theories provide the greatest cognitive benefit.

2.3. Unification

In order to argue that unification is sometimes rational to forgo, we must also clarify what we mean by unification. We aim for as *general* an account of unification as possible, so that our social-epistemological challenge applies to most (if not all) accounts of unification currently in the philosophical marketplace of ideas.

Let a theory *h* fit a phenomenon or item of evidence *e* only if *h* stands in a scientifically significant and accepted relationship to *e*. Candidates for fitting relationships include confirmation, deduction, explanation, prediction, etc. Then the transition from one stage of scientific inquiry t_1 to another such stage t_2 marks an increase in unification if and only if:

- (U1) Fewer theories at t_2 are required to fit the same phenomena than those required at t_1 ; or
- (U2) The same number of theories fit more phenomena at t_2 than at t_1 ; or
- (U3) Fewer theories fit more phenomena at t_2 than at t_1 .

For simplicity's sake, we restrict our attention to (U1) and we assume that the set of phenomena to be fit is finite and fixed over time. We also assume that a unique theory fits any subset of phenomena. This omits the possibility of scientists pursuing different theories that fit the same collection of phenomena, but the variety of theories that can be pursued remains sufficiently rich to capture interesting ideas about unification.

As simplistic as this picture appears, it suffices for our purposes. In motivating the need for a social epistemology that is relatively invariant to the account of unification being endorsed, we thereby allow the reader to fill in the "fitting" relation and to count and define "theories" and "evidence" in whatever manner furnishes her favourite account of unification (Table 1).

Instead of fighting turf wars over the nature of unification, our primary focus concerns aspects of unification most relevant to scientists' decisions about theory pursuit. To that end, imagine a scientific community in which each research unit deliberates about pursuing one of many theories that differ in both the degree and kind of unification they potentially offer. First, *local* theories fit only phenomena within the research unit's field. Second, *cosmopolitan* theories fit phenomena both within the research unit's field and in other fields. This is a sliding scale: maximally cosmopolitan theories fit at least one phenomenon from every field at a given time.

Unification has another dimension. Theories might be *grand*, and fit *all* of the phenomena within their purviews; or they might be *humble*, fitting only a subset of these phenomena. As with cosmopolitanism, grandeur comes in degrees. At one extreme, maximally humble or *parochial* theories fit only one phenomenon. At the other extreme, maximally grand or *global theories* fit all of the phenomena. Localism and cosmopolitanism are thus defined relative to the fields of scientists, while grandeur and humility are defined relative to the total set of phenomena.

Table 1. A sketch of how the leading accounts of unification would be subsumed within our framework

Account of Unification	Theory	Phenomena	Fitting Relationship
Bartelborth 2002	Theoretical Models	Pre-Theoretical Models	Embedding (Explanation)
Kitcher 1989	Argument Patterns	Explananda	Deduction (Explanation) ²
Morrison 1999 ³	Theoretical Parameters	Phenomena	Necessity, Mathematical Structure
Myrvold 2003 ⁴	Hypotheses	Evidence	Conditional Probability, Informational Relevance
Schurz and Lambert 1994; Schurz 1999	Premises	Explananda	Inference
Thagard 1978	Hypotheses	Classes of Facts	Explanation

To illustrate these ideas, consider a simplified model in which there are only three phenomena e_1 , e_2 , and e_3 , and two research units A and B . A 's field consists of e_1 and e_2 ; B 's, of e_3 .⁵ Then the possible theories would be: h_1 , h_2 , h_3 , h_{12} , h_{13} , h_{23} , h_{123} . Here, the subscripts refer to the phenomena the theory fits, e.g. h_{23} fits e_2 and e_3 . Using our earlier terminology, h_1 , h_2 , and h_3 are parochial theories; and h_{123} is a global theory. Given A and B 's fields, h_1 , h_2 , h_{12} , and h_3 are local theories, while h_{13} , h_{23} , and h_{123} are cosmopolitan theories.

We do not assume cosmopolitan theories *inherently* increase unification. However, so long as there is more than one field in a scientific community, at least one cosmopolitan theory is always grander than any local alternative. In this simple model, h_{123} is the only cosmopolitan theory grander than any local alternative. Consequently, so long as there are two distinct research units, the challenge that we raise to unification can—at least in principle—arise.

In what follows, we will be focusing primarily on the conditions under which it is more rational to pursue local theories rather than cosmopolitan theories. We return to these de-

² Both Bartelborth and Kitcher provide unificationist theories of explanation. Consequently, if an explanandum can be embedded/deduced from a theoretical model/argument pattern that is a member of a unifying theory or “explanatory store,” it is explained by an application of that model/pattern.

³ While Morrison states that “there is no ‘unified’ account of unity—a trait that makes it immune from general analysis,” this entry captures one kind of unification that might not be subsumed under other entries, and springs from her discussion that “true cases of unification... have a mechanism or parameter represented in the theory that fulfills the role of a necessary condition required for seeing the connection among the phenomena” Morrison, Margaret. 1999. *Unifying Scientific Theories: Physical Concepts and Mathematical Structures*. New York: Cambridge University Press, page 32.

⁴ See also McGrew, Timothy. 2003. “Confirmation, Heuristics, and Explanatory Reasoning.” *British Journal for the Philosophy of Science* 54: 553-567; Schubach, Jonah N. 2005. “On a Bayesian Analysis of the Virtue of Unification.” *Philosophy of Science* 72: 594-607.

⁵ We use this notation only to illustrate how we classify different degrees and kinds of unification; it does not figure in what follows.

tails below. For now, let us briefly discuss how we will be treating the other dimension of unification—the degrees of grandeur. We will be assuming that, *ceteris paribus*, a grander theory is preferable to a humbler theory. Presumably, unificationists would agree on this point, since, according to (U1)-(U3) above, a grander theory is more unifying. The question remains whether it is also more rational to *pursue* cosmopolitan theories than local theories; our answer will be that there are many cases in which it is not. Moreover, we will be focusing exclusively on the decision between pursuing between a grander cosmopolitan theory and a humbler local theory.

Because we are focusing on cosmopolitan unification, we assume that unification always involves two or more research units. Undoubtedly, this assumption is a simplification; there have been several “lone unifiers” throughout the history of science (e.g., Newton, Maxwell, Einstein)⁶. However, our claim is only existential, not universal: it is *sometimes* rational to forgo unification, even when it promises the greatest benefit. Correspondingly, our claim need not be too strong: *sometimes* unification requires multiple research units. On this front, there are many episodes in the history of science that required substantial coordination, e.g. the modern synthesis of evolutionary theory and genetics, the Standard Model in particle physics, the tectonic theory of continental drift, etc.

Additionally, it is not altogether clear that unificationists are well served by resisting our argument via appeal to lone unifiers. After all, it would then follow that unification is only worth pursuing independently of social conditions when a scientist on the order of Newton, Maxwell, or Einstein dedicates his or her research efforts towards unification. Since few scientists can be presumed to match these figures’ talents, it would then follow that *most* scientists ought not to pursue unified theories. Consequently, we take the assumption that some cases of unification require multiple research units to be defensible.

2.4. *Interfield Hooks and Latches*

On the assumption that unification involves multiple fields, it then seems quite reasonable that some research units must devote special efforts to effecting the desired unification. Let an *interfield hook* be the desired product of these efforts. Each hook can be accepted or rejected by other research units. If one research unit accepts another’s hook, we will say that the former has *latched onto* the hook. We will first define interfield hooks in the abstract, and then provide several examples from the history of science. We will then do the same with interfield latches.

Above, we defined a field as consisting of shared: (i) problems⁷, (ii) domains, (iii) explanatory patterns, (iv) methods, (v) concepts, and (vi) vocabulary. Call items (i)-(vi) the *elements* of the field. Then, research unit *A* builds an interfield hook to research unit *B*’s field if and only if:

- (H1) *A* either uses some of the elements in its field to solve some of the problems in *B*’s field, or uses some of the elements in *B*’s field to solve some of the problems in its field;

⁶ We thank an anonymous referee for raising this issue.

⁷ We are including the fulfillment of ‘explanatory goals’ as a kind of problem to be solved.

- (H2) The solution to these problems increases unification in one of the ways prescribed by (U1)-(U3) above⁸.

We will describe the problems that a hook is intended to solve as its *characteristic problems*. Darden and Maull (1977, 49) provide several examples of interfield hooks. To quote them at length:

- (1) A field may provide a *specification of the physical location* of an entity or process postulated in another field. For example, in its earliest formulation, the chromosome theory of Mendelian heredity postulated that the Mendelian genes were in or on the chromosomes; cytology provided the physical location of the genes...
- (2) A field may provide the *physical nature* of an entity or process postulated in another field. Thus, for example, biochemistry provided the physical nature of the repressor, an entity postulated in the operon theory.
- (3) A field may investigate the *structure* of entities or processes, the *function* of which is investigated in another field. Physical chemistry provides the structure of molecules whose function is described biochemically.
- (4) Fields may be linked *causally*, the entities postulated in one field providing the causes of effects investigated in the other. For example, the theory of allosteric regulation provides a causal explanation of the interaction between the physicochemical structure of certain enzymes and a characteristic biochemical pattern of their activity.

These examples focus quite a bit on building hooks between domains (element ii, above), explanatory patterns (iii), and concepts (v). Consequently, they emphasize an important class of interfield hooks in which one research unit *constructs* a theory that purports to fit evidence in another field. There are many other kinds and examples of interfield hooks. For instance, in the modern synthesis of genetics and natural selection, Dobzhansky's (1937) *Genetics and the Origin of Species* aimed to make the ideas of population genetics—often expressed in complex statistical notation and biological theory—accessible to field naturalists. In this way, the hook is primarily one of presenting a vocabulary that makes population genetics relevant to field naturalists (element vi above).

Our argument in Section 4 assumes that the evidence in *A*'s field is unified with the evidence in *B*'s field only if *A* builds a hook and *B* "latches onto" that hook (or vice versa). This latching is basically *B*'s positive reception of *A*'s hook. More precisely, if *A* builds an interfield hook to *B*, then *B* latches onto that hook if and only if:

- (L1) *B* accepts the way that *A*'s hook solves its characteristic problems; or
- (L2) *B* modifies (elements in) its field so that *A*'s hook solves its characteristic problems in a manner that *B* accepts.

We will take the paradigmatic kind of latch to be one of *testing* whether a cosmopolitan theory fits the evidence in the latching unit's field⁹. Continuing with an earlier ex-

⁸ There is almost certainly a class of interfield relations that satisfy H1, but not H2. These are worthy of further exploration, but are incidental to the present paper's concerns.

⁹ Returning to the previous section, this provides yet another reason to think that "lone unifiers" still require other scientists to effect unification. For instance, Einstein needed Eddington to test his prediction.

ample, Dobzhansky's work simply presented the findings of population genetics in a vocabulary accessible to field naturalists. Field naturalists could have read and understood Dobzhansky's hook, but then rejected the ideas of population genetics. Thus, latching was required to unify population genetics with field studies. For instance, by applying Dobzhansky's ideas to his fieldwork on the birds of New Guinea, Ernst Mayr (1942) went on to write *Systematics and the Origin of Species*, one of the most influential texts of the modern synthesis of natural selection and genetics.

Latching may assume other forms. In some cases, it may involve something other than (perhaps in addition to) testing, e.g. developing methods and models that complement *A*'s hook, interdisciplinary collaboration, or extensive research in *A*'s field. In some cases, it may involve something substantially less labour-intensive than testing, e.g. simply accepting *A*'s hook. This more 'passive' form of latching may seem to make *B*'s role in unification trivial. However, research units tend to have expertise and authority over their own fields. Consequently, if *B* does not accept *A*'s hook, then doubts or uncertainty about whether *A*'s cosmopolitan theory actually fits the evidence in *B*'s field are warranted. Thus, we will say that *A*'s theory stands in a *questionable fit* to *B*'s evidence when *B* does not latch onto *A*'s hook.

Finally, we note that a research unit builds or latches onto a hook only if it pursues a cosmopolitan theory. Lest this seem too strong, recall that pursuit of a theory means allocating resources towards its further development. Building a hook to another field clearly fits this description, since it involves marshalling resources from multiple fields. Latching is also a cosmopolitan pursuit, for its paradigmatic form involves testing whether elements from another field solve problems in one's own field. Since testing of this sort requires resources, latching also entails pursuit of the theory.

To summarize, we are interested in whether a research unit is always rational in pursuing a unifying theory. Given the way that we have defined unification, this question amounts to whether a research unit is always more rational in pursuing a grander cosmopolitan theory instead of a humbler local theory. Moreover, we assume that producing a cosmopolitan theory requires one research unit building an interfield hook, i.e. determining which elements in one field are relevant to another, and another research accepting or latching onto the central features of that hook.

3. *A Decision-Theoretic Model of Theory Choice*

Recall that our main thesis is that it is sometimes rational not to pursue unification, even when unified theories provide the greatest benefits. What is meant by "rational" and "benefit"? According to the basic tenets of "economic approaches to science" (Strevens 2011), it is rational to pursue a theory *b* if and only if pursuing *b* maximizes expected payoff. In turn, the expected payoff of pursuing *b* is a function of the costs of pursuing *b*, the benefits provided if the pursuit is successful, and the probability that pursuing *b* will succeed¹⁰.

¹⁰ Included in consideration of costs are factors such as the amount of resources it might take to successfully produce the theory, the amount of time it might take for the theory to be produced, and so forth.

Here and throughout, “rational” means *subjectively* rational (likewise for “cost,” “benefit,” and “probability”). In other words, our account only appeals to scientists’ mental states (even if these states are objectively irrational or unjustified). Note further that the theory h is defined as a theory that fits some set of phenomena E ; pursuing h means attempting to produce a theory that fits E . That pursuit is *successful* if and only if a theory is produced that actually fits E .

We can model the decision problem faced by some research unit A as follows: A has to choose one theory out of a set H of possible theories to pursue. A ’s choice is rational if the expected value of pursuing that theory is at least as great as that of pursuing any other theory in H .

Suppose performing an action h can produce one of a set of outcomes O , and that a probability function P assigns to each outcome $o \in O$ a probability $P(o)$, such that $0 \leq P(o) \leq 1$ and $\sum P(o) = 1$. Suppose also that the cost of performing an action is given by a cost function $Z(h)$ (where the cost is the same regardless of which outcome is produced), and that the value of each outcome can be given by a value function $V(o)$. Then we can calculate the expected value $U(h)$ of performing h as $U(h) = -Z(h) + \sum_{o \in O} P(o) \cdot V(o)$.

For each research unit, the set of actions corresponds to the set of theories H ; $U(h)$ then represents the expected value (to the research unit) if h is pursued, and each research unit thus aims to maximize the value of $U(h)$. As mentioned earlier, we assume that scientists value evidence (among other things). Consequently, the amount of evidence fit by the theory h influences, but does not determine, the function $U(h)$.

The set H can be partitioned into two subsets, H_L and H_C . H_L represents the *local* theories that research unit A can pursue, such that for every theory $h_L \in H_L$ the phenomena fit by h_L are all in A ’s field, while H_C represents the *cosmopolitan* theories that A can pursue, such that for each theory $h_C \in H_C$ at least one phenomenon fit by h_C is not in A ’s field. Ultimately, we will compare the *optimal local strategy* h_L^* to the *optimal cosmopolitan strategy* h_C^* , where $U(h_L^*) \geq U(h_L)$ for all $h_L \in H_L$ and $U(h_C^*) \geq U(h_C)$ for all $h_C \in H_C$. Given the discussion in Section 2.3, optimal cosmopolitan strategies aim to produce grand cosmopolitan theories; the optimal local strategies, humble local theories. If circumstances exist in which it is more rational for a research unit to pursue an optimal local strategy over an optimal cosmopolitan strategy, then our challenge to unificationism succeeds.

Consider first the subset H_L of local theories. For each theory $h_L \in H_L$, there are two possible outcomes if a scientist pursues it:

- o_1 . *Local Failure*: the pursuit of h_L fails to produce a theory that fits the relevant evidence in the research unit’s field; or
- o_2 . *Local Success*: the pursuit of h_L succeeds in producing a local theory that fits the relevant evidence in the research unit’s field.

In general, $U(h_L)$ depends on the values associated with o_1 and o_2 , on the respective probabilities of these outcomes, and on the costs associated with pursuing a local theory. Label the values associated with these two outcomes of h_L^* as $V(o_1) = 0$ and $V(o_2) = l^{11}$. Further,

¹¹ We set the value of local failure as 0 for convenience. Our conclusion does not change if we set the value of local failure as another value, so long as it is less than the value of local success.

let p be the probability that the pursuit of h_L^* succeeds (i.e. that o_2 is brought about), and $1 - p$ be the probability that the pursuit fails (i.e. that o_1 is brought about). Finally, label the cost associated with pursuing h_L^* as $Z(h_L^*) = z_L$. We can now calculate the expected utility of pursuing h_L^* as

$$U(h_L^*) = (p \cdot l) - z_L.$$

Next, consider the subset H_C of cosmopolitan theories. For each theory $h_C \in H_C$, consider three possible outcomes:

- o_3 . *Complete Failure*: the pursuit of h_C fails to produce a theory that fits the relevant evidence within or outside of the research unit’s field, or
- o_4 . *Questionable Success*: the pursuit of h_C produces a theory that fits the relevant phenomena in the research unit’s field, but that questionably fits some of the relevant phenomena in other research units’ fields; or
- o_5 . *Unquestionable Success*: the pursuit of h_C produces a theory that fits the relevant phenomena in the research unit’s field, and that bears no questionable fits to the relevant phenomena in other research units’ fields.

As discussed in Section 2.4, questionable *fits* arise when affected research units do not accept how the cosmopolitan theory bears on their field (i.e. they don’t “latch onto the hook”). A questionable *success* is when a research unit pursues a theory that results in questionable fits. By contrast, *unquestionable success* is achieved when all affected units accept how the cosmopolitan theory bears on their field. Unquestionable successes are successful cosmopolitan theories, and are grander than either questionably or locally successful theories. In short, unquestionable successes are successful unifications.

Turning our attention to the optimal cosmopolitan theory h_C^* , let $1 - r$ be the probability of o_3 . Then r is the probability that the pursuit succeeds in some way—that is, r is the probability of either o_4 or o_5 . Call r_q the probability of o_4 and r_u the probability of o_5 , such that $r_q + r_u = r$. Denote the value of failure as $V(o_3) = 0$, the value of questionable success as $V(o_4) = h_C^q$, and the value of unquestionable success as $V(o_5) = h_C^u$. Again, we assume for simplicity that the values of complete and local failure are the same, namely 0. Thus, we assume that the value of a theory that fails to fit any relevant evidence is largely unaffected by its being local or cosmopolitan. Finally, we denote the cost of pursuing h_C^* as $Z(h_C^*) = z_C$. Then the expected value of pursuing h_C^* is $U(h_C^*) = (r_q \cdot h_C^q + r_u \cdot h_C^u) - z_C$.

Since it figures prominently below, consider one important factor in determining $U(h_C)$: the relative values of r_q and r_u . Holding the probability of failing to produce a cosmopolitan theory fixed at $1 - r$, the probability of successfully producing a cosmopolitan theory of some kind is fixed at r . The value of that success depends on the ratio between r_q and r_u : the higher r_q , the more the value of $U(h_C)$ will approach h_C^q , while the lower r_q is, the more the value of $U(h_C)$ will approach h_C^u . Let us call the value $k = r_u/r$ the *expectation of unification*. The value of k will range from 0 to 1, and represents A ’s judgment about the likelihood of o_5 given that o_3 does *not* occur—or, equivalently, the likelihood of o_5 given either o_4 or o_5 .

We will argue below that social factors significantly influence k ’s value. If this is right, and if there are values of k that favour pursuing local strategies to cosmopolitan strategies, then we can conclude that it is sometimes rational to forgo unification. In other words,

when faced with the choice between pursuing an optimal local strategy and pursuing an optimal cosmopolitan strategy, there are social conditions wherein the expected value of the former is greater than the expected value of the latter.

4. *Social-Epistemological Challenges for Unificationists*

With these clarifications and assumptions in hand, let us state the precise sense in which it is sometimes rational to forgo the pursuit of unified theories, even when they provide the greatest cognitive benefit:

The Limits of Unification Thesis (LU): There exist research units with sets H and P such that $U(h_L^*) > U(h_C^*)$ and $l < h_C^{u12}$.

In other words, the utility of an optimal local strategy is sometimes greater than that of an optimal cosmopolitan strategy, even when the value of a local success is lower than that of an unquestionable, and hence cosmopolitan, success.

Specifically, we will argue that any plausible response to this challenge will invoke social-epistemological considerations that have, to date, not been forthcoming. Thus, we are challenging the claim that the value of pursuing unified theories is independent of any social conditions. So far as we know, this fairly characterizes every proponent of unificationism (e.g. the authors listed in Table 1). Note that this modest position is compatible with the claim that unification's value is not ubiquitous. Of course, if our arguments undermine even this modest position, they apply *a fortiori* to the stronger thesis that unification's value is not only independent of social considerations, but is also a desideratum of all theories. Should either the moderate or the strong unificationist grant that the value in pursuing a unified theory depends on social conditions, she will have conceded precisely the point we are underscoring.

Our argument will proceed as follows. First, in Section 4.1, we will argue that the conditions that justify LU depend on research units' expectation of unification —i.e., the parameter k , as described in the previous section. Then, we will argue that the expectation of unification is deeply tied to social-epistemological considerations in Section 4.2. It thereby follows that whether it is rational to pursue unification is a social-epistemological question.

4.1. *An Argument for LU*

There are cheap ways to establish LU: one could simply jerry-rig the probabilities, costs, and benefits in an ad hoc manner to get the result. This should convince no one. To that end, we will make assumptions that either grant the unificationist a favourable dialectical position, or are independently plausible. Thus, when we show that it is rational to forgo the pursuit of unified theories, we cannot be charged with tilting at windmills or burning down straw men.

¹² Strictly speaking, it only needs to be the case that the pursuit of a local theory is *at least as rational* as the pursuit of a cosmopolitan theory; that said, we are happy to grant to the unificationist the case in which the values are exactly equal.

Having said that, we wish to underscore a few things. First, we are making an existential claim: it is *sometimes* rational to forgo unification, even when it promises greater benefit. Hence, counterexamples to these assumptions — especially the ones that are unfavourable to the unificationist— must show that our assumptions *never* obtain in the history of science. We suspect that these kinds of rebuttals are unlikely to succeed. Indeed, we suspect that the assumptions that are not merely made for the sake of the unificationist are routinely true. Thus, in talking about the relative value, cost, and probabilities of different outcomes in this section, we assume implicit *ceteris paribus* clauses throughout.

Let us begin with one of LU's two conjuncts:

$$(LU1) \quad l < b_C^u$$

LU1 states that the value or benefits of a local success are less than those of an unquestionable success. This claim is both “unificationist-friendly” and independently plausible. Clearly, any unificationist should accept this claim, since a local success is less unifying than an unquestionable success. Moreover, there are independent reasons for accepting LU1. An unquestionable success fits with more phenomena than a local success, and we have defined a fit as any scientifically significant and accepted relationship between a theory and a phenomenon. More fits clearly are valuable—scientists clearly want to explain, predict, etc. more phenomena than fewer. From this, it follows that unquestionably successful theories stand in more valuable relationships to the phenomena than local successes. Indeed, this reasoning suggests that unquestionable success is greater than *any* of the four alternatives (local failure, local success, complete failure, questionable success) discussed in Section 3.

With this, let us turn to the second conjunct of LU, which states there are conditions where the optimal local strategy is of greater utility than the optimal cosmopolitan strategy:

$$(LU2) \quad U(b_L^*) > U(b_C^*)$$

Now, given the discussion in Section 3, this is equivalent to:

$$(p \cdot l) - z_L > r \cdot (b_C^g + b_C^u) - z_C$$

Given LU1, we know that the benefits or values of pursuing a grander cosmopolitan theory must be greater than those of pursuing a humbler local theory. This is good news for the unificationist, since the remaining elements of the decision problem —the costs and probabilities— are most plausibly construed as favouring humbler local theories.

First, the costs of pursuing a grander theory are greater than those of pursuing a humbler theory. For each phenomenon e , ascertaining whether a theory fits e requires an expenditure of resources (time, effort, thought, equipment, etc.). Since grand theories require more of these expenditures, they cost more to pursue than humble theories.

Second, the costs of pursuing a cosmopolitan theory appear greater than those of pursuing a local theory. In particular, pursuing cosmopolitan theories requires the development of hooks and latches, which consume resources. By contrast, only pursuing a local theory requires neither hooks nor latches. For instance, when pursuing local theories, scientists are familiar with the relevant phenomena, leading hypotheses, etc. but this is not the

case with cosmopolitan theories. Consequently, the scientist must allocate resources to familiarize herself in the adjacent field. Costs, then, increase monotonically along both the humble-grand axis and the local-cosmopolitan axis. Thus, the cost of pursuing a grander cosmopolitan theory is greater than the cost of pursuing a humbler local theory:

$$z_C > z_L$$

However, even here we will be charitable to unificationists, and assume that these costs are vastly outstripped by the benefits of producing a grand cosmopolitan theory. To do otherwise is essentially to magnify the disincentives for pursuing unified theories. Since our goal is to show that one can rationally forgo pursuit of a unified theory even when its cognitive benefits are substantially greater than the alternatives, we readily downplay the greater costs of pursuing a grand cosmopolitan theory. Given the preceding, this entails that the net gain of an unquestionable success is greater than that of a local success, and the net gain of a local failure is greater than that of a complete failure.

With the relative values and costs of the different outcomes spelled out, all that remains is the probability that a research unit assigns to these outcomes. The probability r of producing either a questionable or an unquestionable success is smaller than that of a local success. After all, fitting a theory to more evidence is harder than fitting it to less. Nevertheless, we will assume that this difference is slight —once again in the name of a being dialectically kind. Thus:

$$p > r$$

If things were to end here, LU2, and hence LU, would already seem plausible: it is a wash whether the greater value of an unquestionable success offsets its greater costs and its lower probability of success. Presumably, *sometimes* it does not, which renders the pursuit of a local, non-unifying theory more rational. If that is correct, then LU2 is proven.

However, we are not content to let things rest there. In particular, we will now argue that k , the expectation of unification, is the crucial fulcrum. To see this, consider a limiting case, namely when there is no expectation of unification. In such a case, a research unit expects that no other research unit will accept the way it reconfigures various interfield elements, i.e. it expects that no other research unit will latch onto its hook. In this case, $k = 0$. As a result, $r = r_q$. This in turn, yields the following:

$$U(h_C^*) = r \cdot h_C^q - z_C$$

We have already argued that $r < p$ and that $z_C > z_L$. Consequently, since we are (provisionally) assuming that $k = 0$, if $h_C^q \leq 1$, then $U(h_C^*) < U(h_L^*)$. In other words, we must show that questionable cosmopolitan successes are no more valuable than local successes. This is plausible, and may well be unificationist-friendly, since there is a good case to be made that questionable cosmopolitan successes are strictly less valuable than local successes, i.e. $h_C^q < 1$. By definition, questionably successful theories have questionable fits, which presumably diminish their value. Indeed, because questionably successful theories fail to apply to the phenomena to which they purport to fit, something is frequently wrong with them, e.g. they make false predictions. This, in turn, raises doubts as to whether the fits they actually mus-

ter are bona fide or merely fortuitous. By contrast, a local success fits everything that it purports to fit, so these doubts need not arise.

Indeed, we are hard pressed to find historical examples of questionable successes that were still held in high regard by the scientific community. Relatively close cousins might be the “selective successes” used in some prominent defenses of scientific realism (e.g., Kitcher 2001; Psillos 1999; Worrall 1989)¹³. Perhaps these examples can be repurposed to motivate questionable successes. For instance, while the ether theory failed to unify optical, electromagnetic, and gravitational phenomena, certain models of the ether theory fit their phenomena, e.g. Fresnel’s theory of diffraction.

However, even here we have some reservations, for these arguments take a relatively “long view” of the history of science, and it’s unclear whether scientists take a similarly long view when pursuing theories. Presumably, theories that purport to unify but that only partly fit their evidence are frequently discarded quite abruptly, and hence are unlikely to be salient when one takes a long view. Nevertheless, for the scientist who pursues such a theory, they are very salient —perhaps even more so when they are so quickly discarded. Indeed, when Fresnel pursued his theory, it seemed to be something much closer to an *un*-questionable success —problems with the ether theory would not reach fever pitch for decades.

Despite these worries, we are bracketing these concerns so as to give the unificationist as soft of a landing as possible. Thus, we hold that questionable cosmopolitan successes can be as valuable as local successes. With this, we have shown that it can be more rational to pursue a humbler, local (i.e. non-unifying theory) rather than its grander cosmopolitan counterpart. The crucial condition is that $k = 0$, i.e. the research unit does not expect others to accept its reconfiguration of interfield elements, i.e. to latch onto its hook. Thus, we have established LU.

However, if LU could only be true under these conditions, then it would be of limited interest. After all, it seems implausible that one segment of the scientific community would expect *no* other segment to accept a theory that was relevant to the latter’s field. As we shall eventually argue, higher values of k can also underwrite LU. Indeed, it is somewhat tempting to suggest that even when there are extremely high expectations of unification, it is still not altogether obvious that it is rational to pursue a grand cosmopolitan theory. At its maximum, $k = 1$. This would mean that a research unit expects all of the other research units to latch onto its hook, i.e. to find it an acceptable way to reconfigure elements in their field. In this case, $r = r_u$. This, in turn, entails the following:

$$U(b_C^*) = r \cdot b_C^u - z_C$$

We have argued that $r < p$, $b_C^u > l$ and $z_C > z_L$. It does not thereby follow that $U(b_C^*) > U(b_L^*)$ when $k = 1$, for even here, a cosmopolitan theory’s greater costs and lower probability of success might eclipse its greater benefits. However, we think that this would be unfair to the unificationist. Presumably, when $k = 1$, the conditions are ripe for pursuing a unified theory. Consequently, we assume that the benefits of an unquestionable success

¹³ We stress that the paper is non-committal about the scientific realism debate; we are only using these examples to make a point about unification.

are so much greater than those of a local success, that it outshines the former's greater costs and lower probability of success. Thus, we grant that when $k = 1$, it is more rational to pursue a grander cosmopolitan theory than a humbler local theory.

Despite the various assumptions we have made on behalf of the unificationist, there will still be values of k greater than 0 that make it rational to forgo unification, i.e. that will justify LU. Specifically, there is some value of k between 0 and 1 such that the expected value of pursuing a local theory is equal to the expected value of pursuing a cosmopolitan theory. If the actual expectation of unification is below that level, then the expected payoff of local pursuit is higher, so the rational agent should act accordingly. In all such cases, LU holds. By contrast, if the expectation of unification is higher, then it is rational to pursue a cosmopolitan, unifying theory.

At this point, one may object that unification can be achieved without being pursued. If this objection were sound, then the expectation of unification is irrelevant. For example, it might be thought that scientists might only pursue local theories that serve as inputs or fodder for subsequent unification. However, this merely postpones the problem, for even if earlier scientists don't face the decision problem we have just described, subsequent scientists who unify these local theories do. Specifically, later unifiers will have to build hooks that latch onto other fields, and the probability of this succeeding is precisely what the expectation of unification represents.

4.2. *Social-Epistemological Considerations*

Thus, we have seen that whether or not it is rational to pursue a unifying theory depends primarily on what we have called the expectation of unification. In other words, it depends on a research unit's expectation that it will build a hook and other research units will latch onto that hook. Stripping this of our handy labels, this means that the rationality of pursuing a unifying theory depends on scientists' expectations that they can reconfigure the elements in their fields and adjacent fields in a manner that other scientists will find acceptable.

To complete our story, we point out that this is a fundamentally social-epistemological question. It fundamentally concerns the kind of trust that scientists have towards each other; how the scientific community divides itself and its labour into fields and research units; the scientific incentives for interdisciplinary work; whether the scientific community produced a sufficient amount of background knowledge to support cosmopolitan theories; the compatibility of elements between different fields; the mechanisms and practices of accepting and testing claims from other fields; and so on. There might be other factors affecting k . The crucial point for us is that *however k is determined, it appears to be inextricably connected to social-epistemological factors that unificationists have not adequately developed.*

5. *Conclusion*

We have argued that even if a cosmopolitan theory fits more evidence (i.e. it is more unifying), it may be more rational to pursue a local (i.e. less unifying) theory. This occurs when scientists fail to know whether other scientists will cooperate in the pursuit of the grand

cosmopolitan theory. Thus, since grand cosmopolitan theories are the most unifying, it appears that other, less unified theories can be equally rational to pursue.

Moreover, some of our assumptions about the structure of scientific communities may well understate the tractability of the problem. In particular, both cosmopolitanism (localism) and grandeur (humility) admit of degrees, so a fully general model would give agents far more theories to pursue than our basic model. So, it is also possible that a less idealized account of the scientific community makes it even more puzzling that scientists should always pursue unified theories (and, in particular, global theories).

Nevertheless, we do not claim that our challenge to the value of unification is insurmountable, but we do take it to be instructive, for it appears that any rebuttal to our challenge will involve revising our assumptions about the social structure of scientific communities, or our assumptions about unification. We end by suggesting future lines of research in which either of these assumptions might be challenged, and the ramifications of these challenges to the epistemic value of unification.

First, suppose that our assumptions about the social structure of scientific communities are false, or that some other decision- or game-theoretic structure would adequately describe scientists' pursuit of unified theories. Then, this still concedes the point that we have underscored but that has hitherto been ignored: unification depends on the social structure of the scientific community. Consequently, we welcome this kind of critique.

To make this point more precise: one might think that there are two (possibly competing) sets of interests involved when research units deliberate about which theory to pursue: there are the interests of the scientific community *as a whole*, and there are the interests of each research unit specifically. Given that our paper has analysed deliberation from the standpoint of research units, our conclusion has been that it is sometimes rational for research units not to pursue unifying theories. We might think, however, that an important related problem is that of finding a way to get research units to deliberate from the point of view of the scientific community *as a whole*, which would solve the decision problem by making only the pursuit of grand cosmopolitan views rational¹⁴. While we do not address this issue in this paper, it is noteworthy that most solutions to this problem involve changing the agential framework of the research units, which itself involves a type of intervention on the social structure of the scientific community.

A more damning challenge would target our assumptions about unification, for this may well entail that the social-epistemological dimensions are not relevant to unification. However, even here, the worst possible outcome is not so bad. As we have seen, our general constraints in §2.3 apply to a majority of the extant philosophical accounts. So, if only a few accounts of unification are to be pursued regardless of the scientific community's structure, then these accounts will have a privileged status. Moreover, if these privileged accounts only involve "heterodox" views of unification (e.g., Cartwright 1999; Darden and Maull 1977; Mitchell 2002), then this would require radical revision as to what ought be part of the conceptual or normative core of unification. Consequently, we take our challenge as a call to arms to re-envision the nature and purpose of unification.

¹⁴ Thanks to an anonymous referee for raising this concern.

REFERENCES

- Bartelborth, Thomas. 2002. Explanatory Unification. *Synthese* 130: 91-107.
- Cartwright, Nancy. 1999. *The Dappled World : A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.
- Darden, Lindley, and Nancy Maull. 1977. Interfield Theories. *Philosophy of Science* 44: 43-64.
- Dobzhansky, Theodosius. 1937. *Genetics and the Origin of Species* New York: Columbia University Press.
- Kellert, Stephen H., Helen E. Longino, and C. Kenneth Waters. 2006. *Scientific Pluralism*. Minneapolis, MN: University of Minnesota Press.
- Kitcher, Philip. 1989. Explanatory Unification and the Causal Structure of the World. In *Scientific Explanation*, edited by P. Kitcher and W. C. Salmon, vol XIII, 410-506. Minneapolis: University of Minnesota Press.
- . 2001. Real Realism: The Galilean Strategy. *Philosophical Review* 110: 151-197.
- Mayr, Ernst. 1942. *Systematics and the Origin of Species from the Viewpoint of a Zoologist*. Cambridge, MA: Harvard University Press.
- McGrew, Timothy. 2003. Confirmation, Heuristics, and Explanatory Reasoning. *British Journal for the Philosophy of Science* 54: 553-567.
- Mitchell, Sandra D. 2002. Integrative Pluralism. *Biology and Philosophy* 17: 55-70.
- Morrison, Margaret. 1999. *Unifying Scientific Theories : Physical Concepts and Mathematical Structures*. New York: Cambridge University Press.
- Myrvold, Wayne C. 2003. A Bayesian Account of the Virtue of Unification. *Philosophy of Science* 70: 399-423.
- Psillos, Stathis. 1999. *Scientific Realism: How Science Tracks Truth*. London: Routledge.
- Schupbach, Jonah N. 2005. On a Bayesian Analysis of the Virtue of Unification. *Philosophy of Science* 72: 594-607.
- Schurz, Gerhard. 1999. Explanation as Unification. *Synthese* 120: 95-114.
- Schurz, Gerhard, and Karel Lambert. 1994. Outline of a Theory of Scientific Understanding. *Synthese* 101: 65-120.
- Strevens, Michael. 2011. Economic Approaches to Understanding Scientific Norms. *Episteme* 8: 184-200.
- Thagard, Paul. 1978. The Best Explanation: Criteria for Theory Choice. *Journal of Philosophy* 75: 76-92.
- Worrall, John. 1989. Structural Realism: The Best of Both Worlds? *Dialectica* 43: 99-124.

RANDALL HARP is assistant professor of philosophy at the University of Vermont. His research interests include the philosophy of action and collective action, decision-theoretic models of agency, and social ontology.

ADDRESS: 70 South Williams Street, University of Vermont, Burlington, VT, USA 05401. Email: rharp@uvm.edu

KAREEM KHALIFA is associate professor of philosophy at Middlebury College. His research interests include general theories of scientific explanation, inference to the best explanation, the concept of understanding, and social epistemology.

ADDRESS: 303 Twilight Hall, Middlebury College, Middlebury, VT, USA 05753. Email: kkhalfi@middlebury.edu