Discussion

Philosophy of the Social Sciences
2023, Vol. 53(5) 424–438

DOI: 10.1177/00483931231188011
journals.sagepub.com/home/pos
Sage

Sets, Net Effects, Causal Mechanisms, Subpopulations, and Understanding: A Comment on Mahoney

Stephen Turner[1]

Abstract

This comment discusses the suggestions made in Mahoney's "Constructivist Set-Theoretic Analysis: An Alternative to Essentialist Social Science" (2023). Mahoney presents an approach to cases of intersectionality or confounding which produce causal results unlike those that result from traditional net effects causal modeling. He presents it as an alternative to "essentialism," which he describes as a cognitive error. These alternatives have the same problems as those he attributes to net effects analysis, with one exception: the method does allow for the identification of subpopulations with different causal mechanisms producing different outcomes for this subpopulation.

Keywords

Simpson's paradox, statistical causality, neo-Kantianism, policy analysis, understanding

Received 27 June 2023

[1]Department of Philosophy, University of South Florida, Tampa, FL, USA

Corresponding Author:
Stephen Turner, Department of Philosophy, University of South Florida, FAO 226, 4202 E. Fowler Avenue, Tampa, FL 33620, USA.
Email: turner@usf.edu

The dead hand of neo-Kantianism lies heavily over the social sciences, for multiple reasons. It revivifies in different forms, such as constructivism, post-modernism, relativism, and doctrines about the positionality of knowledge. From Kant we inherit the idea that some conceptual categories are conditions of and bound into our experience of the world (Kant 1781 [1958], 93), but that others—teleology, for example, not bound into our experience but which our minds have a propensity to impose (Kant 1790 [1968], §67, 224–25). Neo-Kantianism added the idea of essential conceptual features of a domain that were presupposed but discernable only through philosophical analysis. At the same time the neo-Kantians insisted on the distinction between psychology and "logic," meaning for them the intrinsic conceptual relations within a domain. But they also embraced the independent idea that there was collective psychology and that this psychology imposed itself on individual, providing them with forms of thought which were not understood by the individual to be personal, and were not the subject of self-consciousness, but were taken for granted. Mahoney's introductory discussion of the problem of essentialism frames it in a manner that is analogous to Kant on teleology: we have a propensity to see the world in terms of essences, which are not there, but we see to be there and take to be real. Kant's version is not so much an error about the world as a mistake about the status of teleological thinking, as well as, implicitly, a theory of the mental processes that produce the result. Teleology makes something real into something intelligible. But the intelligibility is not, so to speak, in the thing itself. This ambiguity is what makes the notion of essences, like the notion of teleology, so cognitively sticky. We cannot see our contribution to what we are apparently finding in the world. Everything is consistent with there being an essence; nothing warns us that we are imposing the intelligibility we think we are discovering, other than, perhaps, the fact that it is sometimes hard to find.

Mahoney comes to the problem from a different direction: applying cognitive science to methodology. He argues that essentialism is a cognitive trick the mind plays on us which needs to be overcome to see other empirical patterns that are in some sense better. And he gives an account of what kinds of patterns can be revealed if we set our essentialist cognitive predispositions aside. This is an intriguing and important line of argument that has many analogues. We make typical cognitive errors about probability, for example, for similar reasons: we confuse typical cases with actual probabilities (Gigerenzer 2002). Base rate errors are commonplace and seem to be grounded in deep cognitive biases. Stereotyping produces the same kinds of effects: stereotypes arise through normal cognitive processes and are economical for thought, but do not match the empirical probabilities. What is good to think, to quote Levi-Strauss's phrase, may not be as good for prediction or summarizing data. And we can invert notions like natural kinds by

reinterpreting them in cognitive science terms, as packages of similar affordances that lead us unconsciously to group objects into a kind.

In most of these cases, and the ones made famous by Kahneman (Kahneman 2011) the distinction is between an explicit formal model of correct reasoning and the cognitive shortcuts that are the product of evolution. Much of this research, for example on priming, has turned out to be thin and problematic, and caught up in the replication crisis (Hartmann and Meijs 2012).[1] There is a general problem here, which is difficult to be precise about, but which can be seen in a few examples, some very familiar from an older train wreck: neo-Kantianism. The demise of neo-Kantianism came over the multiplicity of issues, but one central one concerned the contrast between "life," which was dynamic and changing, and the constitutive concepts that neo-Kantianism understood to organize experience, which were rigid and fixed. The problem extends to formalisms generally. There is another problem: the informal, tacit, "cognitive error" processes are interwoven with thinking generally: we are not rule-driven automatons. "Essentialism" itself is indistinguishable from extracting "meaning" from a situation. It is the absolutization of such extractions that leads to philosophical problems.

The general issue is this: we have formalisms, such as the predicate calculus, correlation, directed arrow graphs, decision theory, and, Mahoney's preference, set theory, which we can use to represent patterns. And we can just put things in categories based on formal and explicit criteria of some kind and count them. But the formalisms never quite match up to the things we want to say with them. It is notorious that correlation is not causation. But we have no formalism for causation. It is not a formal concept. We can try to capture the notion of scientific law, another case where there is a mismatch, by specifying its logical features, as Hempel and Oppenheim did (1948), but adding a non-formal rider that distinguishes mere deduction from a generalization from a law-like explanation. And, despite claims suggesting otherwise, we cannot "prove cause mathematically," as Clark Glymour once claimed in a conference (cf. McKim and Turner 1997). We need additional assumptions that add the element that makes the directed arrow graph representing a causal model or correlation "causal."

We have the same problem with sequences, as I ruefully discovered myself many years ago (Turner 1980) trying (not very successfully!) to provide a formalization that could be given a statistical interpretation. One could formalize sequences set-theoretically in terms of order pairs or ordered n-tuples,

[1] Cf. https://replicationindex.com/2017/02/02/reconstruction-of-a-train-wreck-how-priming-research-went-of-the-rails/comment-page-1/. https://replicationindex.com/2017/02/02/reconstruction-of-a-train-wreck-how-priming-research-went-of-the-rails/comment-page-1/#comment-1454. There are also issues of interpretation (cf. Hartmann and Meijs 2012).

but the "sequence" part, if it implies anything more that ordered, needed to be added. Temporal succession is not a formal concept. This is relevant to what follows, because Mahoney's positive suggestion also involves a sequence concept, causal pathways, which has a similar problem: the set machinery, such as the order in ordered pairs and n-tuples is not itself sequential: it is merely ordered. Nor is it causal. That part needs to be added.

There is another side to this problem, exemplified by the discovery—using AI and powerful computers—of chess strategies that were successful but had never been discovered by human chess players in the long and carefully recorded history of chess matches. This suggests something more general: that there are patterns that we are cognitively unable to see unless they are found for us. Addition is "natural," at least in the sense of being part of our ordinary experience with objects, like apples. Quaddition is something we need to be instructed in on the basis of a theory, even though, mathematically, it is just another pattern which we can formalize. And there are patterns that can be identified for us that we cannot make sense of, even if we can apply them formulaically. Algorithms derived from connectionist machine learning, indeed, are such patterns which we cannot understand meaning cannot restate it an idiom we understand. This suggests, minimally, that there is a gap between mere empirical patterns and the normal contents of science, including sociology, and indeed between what is humanly understandable in the patterned world. But perhaps some basic cognitive science considerations can account for this gap without reviving essentialism. Nor is this a hypothetical: "big data" now routinely generates these relations, mindlessly, without theory, solely by searching for patterns. So we are faced here with a real problem that is worth addressing, but one that is bigger and more complex than simply a choice between essentialism and some alternative. It goes to the heart of what social science understanding amounts to.

1. Mahoney's Solution

The first part of Mahoney's argument is this. Getting cognitively tricked into essentialist thinking is bad. Stereotypes, as one kind of essentialist thinking, may lead us to misrepresent the kinds of relations we might discover by attending to the actual data, both by not corresponding to the data but by leading us away from better explanations. He wishes "to escape this insidious essentialism," and to do this "we need an approach that allows us to analyze the social world as composed of something other than entities in possession of internally or externally-derived properties that do not depend ontologically on human minds" (Mahoney 2023, 339). As we will see later, his example of such an internally derived property is intelligence. He assumes that we normally do experience entities in this insidious essentialist way. This is a complex and puzzling thought, but the punchline is this: "We need an approach that views

human reality in a way that does not correspond to how we ordinarily experience it" (Mahoney 2023, 339).

A banal and deflated construction of this thought might be this. We can categorize things in a better and more objective or explicit way and come up with different results. But as Mahoney's later discussion will make clear, he is not arguing that the goal is identifying mere empirical patterns. The point is contained in the term "ordinarily." The way in which we are called on to experience human reality is still for it to be intelligible, but in a non-ordinary, meaning a non-essentialist, way. If the new categories are themselves intelligible or intelligibility producing, we have a potential improvement over our essentialist ones. He is not promoting mindless big data, but something else: new approaches that do "depend ontologically on human minds" (Mahoney 2023, 339). There is another bit of Kant that applies here. We only understand what we ourselves have created (Kant 1781 [1958], 93).

But this is where things get tricky. The problem is familiar from Max Weber, who faced it in a related but different way. He wanted explanations of action to be both objectively valid and subjectively meaningful: he knew that social action could not be reduced to causal laws, meaning true generalizations, because there were none in social science. The patterns were probabilistic. But there were probabilistic patterns everywhere, some very weak, some very strong. Weber borrowed his colleague Johannes von Kries's account of probabilistic causation as it had been applied to questions of legal liability (Turner and Factor 1981). The method was this: if one defined equivalence classes of conditions or events, and calculated the probability of a given outcome on the probability of a set of conditions or events defined in terms of these equivalence classes, and then subtracted the probability of the outcome of the set minus a potential cause, one could get a number. The method worked with, in principle, objective categories. The number was an objective result. The difference, if it reached a certain threshold, could be regarded as "adequate" and therefore for legal purposes a cause rather than a coincidence or an irrelevance. But this threshold was conventional.

This reasoning was in practice mostly hypothetical or heuristic because the relevant probabilities only existed in principle (Turner 2018). But it provided a usable account of causality for legal and also sociological purposes: an event or condition was not a cause unless it raised the probability of an outcome beyond a negligible level, and made sense in terms of intelligible actions. Just as Karl Pearson knew that everything was correlated, however slightly, with everything else, Weber knew that there were categories with probabilities—or in our more familiar terms, correlations—that existed but did not make sense, or at least were not causal relations. These would include the familiar cases of accidental correlation, confounding, spurious relations, and so forth, but also meaningless correlations or probabilities. The requirement of adequacy on the level of meaning was a way of eliminating at least some of these cases, for

action at least, because action was already defined by him as subjectively meaningful. But underdetermination was still possible.

Mahoney does something parallel to this, but with the concept of sets, taken from a methodology pioneered by Charles Ragin. Sets, like dependent probabilities and correlations, are everywhere, and do not have any intrinsic meaning, causal or otherwise. Just using the concept of sets does not help: there is an infinity of sets. The problem is to find the sets that tell us something about the thing we want explained. But he does not want to find these by succumbing to the cognitive trick of essentialization. He needs a way to identify the sets that he wants to use as an alternative to essences. But sets are merely mathematical objects with members. What he is really after is kinds. And we need some grounds for putting thing into kinds. Social things, as Mahoney is quick to point out in criticizing others, present special problems. They are not "natural kinds," but some other kind of kind: social kinds.

What defines a social kind, without essentialism? His solution to the problem oddly parallels Weber. But rather than using the language of neo-Kantianism, or notions like "constitutivity," he appeals to cognitive science, now in a positive way. He appeals to the concept of conceptual spaces.[2] The idea is that this concept will help us with non-natural, social kinds, which allow us to think about mind-dependent social entities, of the kind we are familiar with. Conceptual spaces are, for him real parts of the cognitive machinery that perform a certain task: "Through conceptual spaces, human minds transform mostly incomprehensible natural substances and properties into the coherent social entities of human reality" (Mahoney 2023, 331). In short, cognitive spaces create what we experience as entities.

Neo-Kantianism was anti-psychological in principle and located presuppositions in the "logical" side of the logic-psychology divide. Presuppositions were the logical condition of the conceptual domains it analyzed. Mahoney moves them to the psychological or cognitive side. Thus, instead of, for example, using Weber's term, "ideal-types," which applied to constructions that were meant to replace "essentialist" terms with meaningful but less ideologically loaded terms, Mahoney uses the term used in cognitive space theory, prototypes, for an analogous purpose, of escaping from faulty stereotypes, or essentialisms (Mahoney 2023, 341).

[2]This is not as far from neo-Kantianism as it might appear. Cognitive science thinking about thinking diverges into two groups: one, like this one, simply takes over neo-Kantian ideas in the form of cognitive frames, which can be "shared"; the other rejects this in favor of connectionist learning and predispositions which reflect the fact that brain signatures are individualized, and that there is no mechanism for sharing frames.

Comparing slow thinking to decision theory works by assuming decision theory as a model for the rationality in question. Mahoney does not have a "correct" alternative. But he needs to ground his alternative somehow. He goes directly to the cognitive for an approach to this. He thinks objects are socially constituted. So he needs an account of how cognitive spaces are shared.

> The conceptual space model offers a way of representing what it means for two or more individuals to *share an understanding* of a social category. When individuals share an understanding of a social category, a conceptual space for the category exists in each of their minds. These conceptual spaces stand in similar distance relations to other conceptual spaces for related categories and background understandings of society. For example, if two or more individuals share an understanding of the category *marriage*, similar distance relations exist between their prototypes of *marriage* and other related categories, such as *two-person relationship, legal union, romantic relationship, and financial relationship*. (Mahoney 2023, 341–42; italics in original)

Sharing thus plays a large role in his account. And social kinds, as he is at pains to stress, are not like natural kinds. So "sharing" is itself constitutive of social kinds: "The one thing that all instances of a given social category have in common is their activation of a conceptual space in the minds of those who share an understanding of the category" (Mahoney 2023, 343). Thus, "The members of social categories (e.g., all scientists, all marriages, all crimes) are similar by virtue of their membership in the same conceptual spaces, not by virtue of any mind-independent properties—whether internally or externally derived—that they possess" (2023, 346). So social categories are mind dependent, but real.

The reasoning here is roughly this: conceptual spaces are real or ontological, and they are shared; the categories or sets they define are therefore valid or real in a special sense derived from this real condition; therefore, the entities one can define using these sets or categories, using such set theoretic operations as union and intersection, or in diagrams, are also real or valid. Why is this an improvement on mere categorization according to explicit criteria for set membership? "Constructivism enriches set-theoretic analysis by ridding it of essentialist foundations, thereby allowing for valid inference with social categories" (Mahoney 2023, 346). Spatial notions are thus better foundations. But they are relative foundations: "Constructivist set-theoretic analysts seek to define, use, and code their categories in ways that are well understood by and meaningful to their readers. These readers are usually the principle semantic community of interest" (Mahoney 2023, 347).

> Constructivist set-theoretic researchers do not arbitrarily stipulate category definitions and meanings. … researchers explicitly follow existing category usage in relevant communities. These existing communities may be specialized academics who already agree about the meaning of scholarly terms. They may also include non-academic groups who employ categories in clear and consistent ways within their communities. (Mahoney 2023, 347)

Why does this matter? Perhaps for this reason: the categories are in some sense pre-validated as intelligible. So, we are not faced with the problem of meaningless statistical relationships. But this takes us only so far. We are still faced with the rest of the problems of interpreting statistical relationships involving these categories.[3] This is an important point to which I will return. It will play a large role in interpreting his example.

2. The Example

Mahoney wants to show the general superiority of his preferred set-theoretic approach to the dominant variables model and the net effects view of causality.[4] The example he chooses is a particularly vexed one. He, like the authors he discusses, Charles Ragin and Peer Fiss, refers to the earlier dispute over Richard Herrnstein's claims (Herrnstein and Murray 1994) about the effects of intelligence (Mahoney 2023, 353). Ragin and Fiss are critical of the over-simple model applied by Herrnstein, and of its quietistic policy implications. This argument raises many questions that cannot be gone into here. Suffice it to say that all of the usual problems of causal analysis arise in this case, and in the vast literature that has developed on racial differences in income and wealth. Confounding, spuriousness, selection bias, causal arrow ambiguity, homogeneity assumptions, circularity, redundant casues, assumptions about populations, the problem of whether a fixed characteristic is a cause: they are all there in profusion. Different measures and different outcomes, which do not differ except for apparently minor details of measurement, can produce major

[3]The relationships are statistical, as is the causality, a point obscured by Mahoney's appeal to a regularity view of causality. The relations are not regularities. Nor is the requirement of continuity as usually understood met in these cases. The issues were discussed long ago in relation to Salmon's statistical relevance theory, which was undermined, as Nancy Cartwright showed in a counterexample, by a kind of confounding intervention in the causal process (Cartwright 2002 [2006]; cf. Salmon 1984, 1994). The irrelevance of continuity in this sense to human action was discussed earlier (Turner 1982).

[4]For this term see Turner (1997).

differences. And outcomes that should be similar turn out to be radically different. Untangling these issues is exceptionally difficult, and studies are routinely reinterpreted with different assumptions and statistical measures. I do not propose to do that here. But I will provide an alternative interpretation of the case, and of the methods, in terms of "net effects" causality.

The core issue here is confounding. The secondary one is about the categories themselves, which can be explained first. Mahoney's reasoning is that the properties model of explanation assumes coherent internal properties possessed by individuals, which are taken to be like natural kinds. He is particularly concerned with one: intelligence. His alternative is to reinterpret categories to make them community relative. "A constructivist set-theoretic analyst could study *intelligent person* and *dumb person*, but these categories would be understood as beliefs and understandings within a specific community or society, as opposed to representing any coherent and shared features possessed by individuals, whether internally or relationally derived" (Mahoney 2023, 358; italics in original).

The distinction has implications for the causal stories we tell, but also in the questions they answer. The difference is in cognitive aims. Mahoney explains it in this way: "These set-theoretic findings are not intended to estimate the average effect of a change on a variable for poverty outcomes. Instead, they are intended to help readers understand regularities in the social world" (Mahoney 2023, 257). The understandings in question are about what category membership properties mean for different groups.

Saying some proportion of some difference results from intelligence runs into a problem, for Blacks, of confounding, which makes net effects reasoning problematic. As Mahoney explains,

> … if an individual is in the category *Black person*, it is hard to separate the individual's membership in *not-high-AFQT-score* from their membership in *not-high-income-parents*. Ragin and Fiss show that these overlapping sets "should be considered jointly and not treated as separate or 'independent' ... it is hazardous to try to separate the effects of test scores and parental income on poverty" [Ragin and Fiss 2017, 98–99]. (Mahoney 2023, 355; italics in original)

For White people, this is not a problem: the two categories do separate. For Black people, they are often the same.

Although Mahoney does not say this, there is a known problem here that undermines causal analysis of racial differences generally. The assumptions about the population needed to do the kind of causal modeling done here include the assumption that we are looking at one population; one in the sense

that the causal processes are the same for everyone in the population, rather than different for different subpopulations. The problem here can be briefly explained, and is highly relevant to what follows. The basic idea is this: if we had an actual law governing something, such as the velocity of falling objects, which we display empirically as a bunch of data points around a line involving random error, such as measurement error, and we selected out a subset, such as "red objects," and it turned out that they behaved differently, perhaps being defined by a line with a slope in the other direction, we would have to reject the "law" and reject the idea that these are the "same" populations. If there are distinct "populations" or subsets, we are faced with the possibility of Simpson's paradox type outcomes: the observed correlation might disappear or be reversed if we break the population into subpopulations.[5]

This becomes crucial for net effects reasoning. If we use the general correlation between AFQT scores and a given outcome to correct other correlations in order to identify causal contributions using standard multiple regression methods, we are assuming that the causal mechanisms in both groups are the same. What Ragin and Fiss show, in effect, is that this assumption is problematic. For one population the two variables almost completely coincide; for that population the effects of the two variables cannot be separated. If we understand this fact of coincidence in causal terms, we have to face the possibility that the coinciding subgroup is causally deviant, that different causal relations hold for it, and that it is thus non-comparable using the usual correlational methods.

This seems like a dead end. But the virtue of the Ragin approach is that it suggests a way of turning this difference into a comparative question, and even, with suitable assumptions, of quantifying it. If we turn degree of coincidence into a variable, we can ask some novel empirical questions. Ragin provides an example in his analysis of persistent poverty, an

[5]Cartwright gives a clear account of the paradox and its relevance to these kinds of cases: "Take any fact about the conditional probability of one factor, say A, on another, B. B may increase the probability of A or decrease it or leave it unchanged. Consider a third factor, C, which is probabilistically dependent on both. Then, depending on how the numbers work out, if we stratify on ± C, the original relations between A and B can be shifted in any way at all: B may increase the probability of A in both subpopulations, or decrease it, or leave it unchanged. A standard example is the Berkeley graduate school, which appeared prima facie to be discriminating against women: the probability of admission given one was a woman was lower than the probability of admission if one was a man. But, department by department, this turned out not to be the case. What was happening was that women were applying to the departments that were more difficult to get into. The 'true' relation between admission and sex is revealed by stratifying on departments" (Cartwright 2002 [2006], 6).

important and controversial policy problem which net effects analysis has not resolved, largely because of the coincidence/confounding problem.[6] The confounders Mahoney considers are test scores and parental poverty. The analysis assumes that the differences in outcome are the result of differences in causal paths, which involve advantages and disadvantages. The novel finding is that there is "a strong connection between advantages and avoiding poverty for whites and a strong connection between disadvantages and experiencing poverty for blacks" (Ragin and Fiss 2017, 147). In short, the causal paths for Blacks drive them toward the category of poverty, while those of Whites do not.

> We show that whites' high coincidence of advantages is coupled with a low coincidence of disadvantages, while blacks' lower degree of coinciding advantages is coupled with a higher degree of coinciding disadvantages. The six pairwise correlation coefficients, by contrast, show virtually no racial differences, indicating that race-linked coincidence patterns are neutralized in correlational analyses. (Ragin and Fiss 2017, 148)

What does this mean in terms of explanation? The neutralization occurs on the population level. The correlational analysis in terms of the total population does not tell us anything about racial differences. They vanish. The intersectional one, i.e., one dealing with a subpopulation, does.

> As we have argued, the customary focus on net effects of independent variables and the use of correlation-based methods tend to hide from view the intersectional nature of social inequality. Instead, the set-analytic approach we offer shifts the focus from the separate effects of independent variables to "causal recipes," asking "What combinations of causally relevant conditions are linked to the avoidance versus the experience of poverty?" Our approach reflects the

[6]There are good reasons for doing this. Paul Holland has long argued against taking unchangeable properties as causes: "I take the position that causes are only those things that could, in principle, be treatments in experiments" (Holland 1986, 954). And a focus on policy, as in Ragin and Fiss, as well as in the interventionist rather than the net effects interpretation of causal analysis all point in the same direction. There is also the difference between social and natural explanations, with the latter bottoming out in understandable actions. This bears on Mahoney's rejection of intrinsic properties as variables. But banning intrinsic qualities or features that cannot be the subject of interventions or for some other reason, such as involving classifications into kinds that are not part of the local culture, merely turns them into subpopulation properties which make the populations different and non-comparable using these standard methods. The causal issues do not go away, they simply arise in this new form.

> fundamentally intersectional nature of social inequality—a quality that is relevant to the analysis of many social and "wicked" problems, beyond poverty. It is this quality that makes the application of a set-theoretic approach so powerful [Blackman, Wistow, and Byrne 2013]. (Ragin and Fiss 2017, 155; italics in original)

The alternative, however, is also causal. It asks, in effect, what coincidental/confounding combinations of conditions are "linked" to the experience or avoidance of poverty. The "link" from the formalism of "set-coincidence," however, it does not tell us anything about causality. That needs to be added. The "causal recipes" amount to an assumption about causality. Similarly for the quasi-explanatory notions of advantage and disadvantage. They are, formally, merely categorical concepts which, combined, can be associated with outcomes. But they do produce a striking result:

> whites have a very high coincidence of advantages—not-low-test-scores coincides strongly with not-low-income-parents. Blacks, by contrast, have a strong coincidence of disadvantages—not-high-test-scores coincides strongly with not-high-income-parents. (Ragin and Fiss 2017, 147; italics in original)

The findings also indicate a difference in causal mechanisms, at least in terms of their effect.

> It reveals that specific combinations of disadvantages—low-test-scores, unfavorable-family-background, and unfavorable-domestic-situation—are more consequential for blacks than for whites when it comes to avoiding poverty. (Ragin and Fiss 2017, 152; italics in original)

"More consequential" is the key result. The overlapping character of the specific disadvantages in the specific subpopulation produces a result indicating a difference in causal mechanisms: the same variables have a greater effect in this subpopulation. From the net effects perspective, what this shows is the systematic deviation of a particular population from the general trend with respect to causal mechanisms—this is what "more consequential" means here. The combination itself, not just its additive effects, has an independent effect.

The kinds of dramatic subpopulation differences one can visualize on a scattergram to illustrate Simpson's paradox do not often occur with real data. The concepts of subpopulation and the causal mechanisms that operate within them are probabilistic and messy. Subpopulations are not cleanly defined. The extent to which a subpopulation's causal mechanisms deviate from those that appear in the total population or from other subpopulations are matters of degree. But the effects, however hard to see, are there. Ragin and Fiss find one

when they show that the payoffs of "advantages" are not as great for Blacks with multiple disadvantages. As a policy matter, it is an important finding: it shows us why this subpopulation is not as greatly benefited by this intervention.

As a methodological innovation, it does even more. It provides us with a way of thinking about subpopulations generally: not only are they groups within which there are different causal mechanisms,[7] we can use the fact of confounding/coincidence to describe the differences between the populations in terms of the particular coincidences that appear in them. Ragin and Fiss do this by identifying the features, which they call advantages and disadvantages, which differentiate them from one another. They understand these differences as producing different possible paths. We can adopt this basic idea, and apply it to all the subpopulations that make up a correlation visible at the population level, and identify the different sequences or paths that are characteristic of the these subpopulations. Just as certain combinations of disadvantages have a larger than predicted effect, so would other combinations.

We can go further by discarding the language of advantage and disadvantage. We can make it an empirical question as to what set membership characteristics are linked to particular outcomes—such as probability of becoming a nuclear physicist, or a professional athlete, or whatever outcome interests us—in particular subpopulations. And with this we could define empirically what are "advantages" and "disadvantages" in relation to particular outcomes within particular subpopulations. This would give us a mosaic of the subpopulations that make up the larger population, each of which may, as we can expect from probabilistic and fuzzy relations, have somewhat different causal relations, different confoundings, in different degrees, different combinations of conditions for outcomes, and therefore different "paths," if we chose to conceptualize the relations in this way.

Ironically, however, by this kind of conceptualization into causal paths, narrativizing them to make the sequences intelligible, making them into "paths," is just the kind of sense-making essentializing the paper begins by attacking. As Weber saw when he discussed the need for one-sided ideal-types, some kind of "essentializing" simplification is inseparable from understanding when it is applied to groups in this way. In short, we have not escaped either the problems of net effects reasoning or the problem of essentializing. But we have added a potentially useful way of approaching a basic problem, by replacing an assumption about the undifferentiated causal

[7]It also suggests a different view of populations: as consisting not of causally homogenous bell-curve like populations, but of different subpopulations with different causal mechanisms, indeed of different individuals whose internal causal mechanisms differ, but whose differences largely balance out at the population level. Identifying the differences is relevant both to explanation and policy, as Ragin and Fiss suggest.

character of the larger population with an empirically grounded means of representing differentiation into subpopulations.

We are still faced with the problem of what makes something causal, and what makes something intelligible. If we want to draw causal conclusions, we cannot avoid the task of making sense of the results by constructing a causal narrative that fits the subpopulation and the specific differences in causal outcomes that our analysis has identified. Ragin and Fiss themselves acknowledge this when they defer to the qualitative inquiries that would flesh out their results: inquiries in the category of understanding. What makes sense in one causal context may not in another. If this is what Mahoney is trying to capture with the notion of symbolic space we can leave it there, without going into the wilds of ontology. These are the familiar conditions of social science generally.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iD

Stephen Turner https://orcid.org/0000-0002-7538-0533

References

Blackman, Tim, Jonathan Wistow, and Dave Byrne. 2013. "Using Qualitative Comparative Analysis to Understand Complex Policy Problems." *Evaluation* 19 (2): 126-40. doi:10.1177/1356389013484203

Cartwright, Nancy. 2022 (2006). "From Causation to Explanation and Back." In *The Future for Philosophy*, edited by Brian Leiter, 230–45, Oxford: Clarendon Press. https://core.ac.uk/download/pdf/73975.pdf.

Gigerenzer, Gerd. 2002. *Calculated Risks: How to Know When Numbers Deceive You*. New York: Simon & Schuster.

Hartmann, Stephan, and Wouter Meijs. 2012. "Walter the Banker: The Conjunction Fallacy Reconsidered." *Synthese* 184 (1): 73-87. doi:10.1007/s11229-009-9694-6

Hempel, Carl G., and Paul Oppenheim. 1948. "Studies in the Logic of Explanation." *Philosophy of Science* 15 (2): 135-75. doi:10.1086/286983

Herrnstein, Richard J., and Charles Murray. 1994. *The Bell Curve: Intelligence and Class Structure in American Life*. New York: Free Press.

Holland, Paul W. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association* 81 (396): 945-60. doi:10.1080/01621459.1986.10478354

Kahneman, Daniel. 2011. *Thinking: Fast and Slow*. New York: Farrar, Strauss and Giroux.

Kant, Immanuel. 1781 (1958). *Critique of Pure Reason*, translated by Norman Kemp Smith. London: Macmillan.

Kant, Immanuel. 1790 (1968). *Critique of Judgment*, translated by J. H. Bernard. New York: Hafner Publishing Co.

McKim, Vaughn, and Stephen Turner. 1997. *Causality in Crisis? Statistical Methods and the Search for Causal Knowledge in the Social Sciences*. Notre Dame, IN: University of Notre Dame Press.

Mahoney, James. 2023. "Constructivist Set-Theoretic Analysis: An Alternative to Essentialist Social Science." *Philosophy of the Social Sciences* 53 (4): 327-66. doi:10.1177/00483931231178100

Ragin, Charles C., and Peer C. Fiss. 2017. *Intersectional Inequality: Race, Class, Test Scores, and Poverty*. Chicago: University of Chicago Press.

Salmon, Wesley C. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press.

Salmon, Wesley C. 1994. "Causality Without Counterfactuals." *Philosophy of Science* 61 (2): 297-312. doi:10.1086/289801

Turner, Stephen. 1980. "Modeling and Evaluating Theories Involving Sequences: Description of a Formal Method." *Quality and Quantity* 14 (4): 511-18. doi:10.1007/BF00184717

Turner, Stephen. 1982. "On the Relevance of Statistical Relevance Theory." *Theory and Decision* 14 (2): 195-205. doi:10.1007/BF00133977

Turner, Stephen. 1997. "Net Effects: A Short History." In *Causality in Crisis? Statistical Methods and the Search for Causal Knowledge in the Social Sciences*, edited by Vaughn McKim, and Stephen Turner, 23-45. Notre Dame, IN: University of Notre Dame Press.

Turner, Stephen. 2018. "What Do Narratives Explain? Roth, Mink and Weber." In *Towards a Revival of Analytic Philosophy of History: Around Paul Roth's Vision of Historical Sciences*, edited by Krzysztof Brzechczyn, 130-47. Leiden, the Netherlands: Brill.

Turner, Stephen, and Regis Factor. 1981. "Objective Possibility and Adequate Causation in Weber's Methodological Writings." *Sociological Review* 29 (1): 5-29. doi:10.1111/j.1467-954X.1981.tb03019

Author Biography

Stephen Turner is Distinguished University Professor of Philosophy in the University of South Florida. His *Making Democratic Theory Democratic: Democracy, Law, and Administration after Weber and Kelsen,* with George Mazur, appeared in March 2023. A discussion of topics related to this paper, "Progress in Sociology?" was recently published in *New Philosophical Perspectives on Scientific Progress*, edited by Yafeng Shan.