



Mapping higher-level causal efficacy

Horia Tarnovanu¹ 

Received: 10 April 2020 / Accepted: 23 April 2021 / Published online: 4 May 2021
© The Author(s) 2021

Abstract

A central argument for non-reductive accounts of group agency is that complex social entities (states, companies, churches, political parties) are capable of exerting causal influence independently of and superseding the causal efficacy of the individuals constituting them. A prominent counter is that non-reductionists run into an insuperable dilemma between identity and redundancy – with identity undermining independent higher-level efficacy and redundancy leading to overdetermination or exclusion. This paper argues that critics of non-reductionism can manage with a simpler and more persuasive reductio strategy called *mapping*: allow that group agents are causally efficacious in their own right and chart (a) how their causal efficacy is carried out; (b) how it relates to the causal efficacy of individual determiners; (c) how it connects to the causal relevance of background structural factors. The focus exclusively on whether groups *are* or *are not* causally efficacious black-boxes implementation, while close attention to how causation is wired increases the visibility of individualist arguments and countenances structure-oriented explanations.

Keywords Causation · Group agency · Individualism · Mapping · Non-reductionism

1 Introduction

A central argument for non-reductive accounts of group agency¹ is that complex social entities (states, companies, churches, political parties) are capable of exerting causal influence independently of and superseding the causal efficacy of the

¹ Non-reductionists define groups as real and robust entities, irreducible to the networked individuals constituting them (Elder-Vass, 2010; Epstein, 2015; Hindriks, 2008; List and Pettit, 2011). Although reductionist claims standardly involve identity relations and entail eliminativist or conservative views, I take ‘irreducibility’ and ‘non-reductionism’ to describe the distinct, detached, high-level profile of group agents, irrespective of the preferred building relation (composition, constitution, grounding, etc.).

✉ Horia Tarnovanu
tarnovanu@gmail.com

¹ SPAIS, University of Bristol, Bristol, UK

individuals constituting them (Elder-Vass, 2010; Menzies & List, 2010; Sawyer, 2003). A prominent counter is that non-reductionists run into a dilemma between identity and redundancy—with identity undermining higher-level efficacy and redundancy leading to overdetermination or exclusion. In this paper, I argue that critics of non-reductionism may appeal to a simpler and more persuasive reduction strategy called *mapping*: allow that group agents are causally efficacious and chart (a) how their causal efficacy is carried out; (b) how it relates to the causal efficacy of individual determiners; (c) how it connects to the causal relevance of background structural factors. Mapping has several advantages: (i) it is straightforward; (ii) it exposes the way efficacy is implemented independently of standard overdetermination- and Eleatic Principle-related arguments; (iii) it shows that non-reductive accounts lack a well-defined causal story about how group agents impact the world; (iv) it shifts attention from the agential capacities of social entities (and agency-building relations) to the way collectives actually act. The focus exclusively on whether groups *are* or *are not* causally efficacious black-boxes implementation, while close attention to how causation is wired increases the visibility of individualist arguments and countenances structure-oriented explanations. While mapping motivates a methodological realignment towards efficacy and offers reasons for favouring reductionism, it is not intended as an intransigent, decisive move against non-reductionist intuitions, but as a stimulant for further discussion and refinement. The possibility of further anti-reductionist moves with respect to mapping remains open.

Plan: Section 2 spells out the causal constraint for higher-level agency and the limitations of the identity-redundancy dilemma. Section 3 introduces mapping. It argues that charting the way group-level causation carries out and relates to individual determiners and auxiliary structural factors tones down the attractiveness of higher-level efficacy. Section 4 considers the non-reductionist appeal to realisation-insensitivity of higher-level causal claims as a study case for the application of mapping principles. Section 5 is reserved to conclusions and methodological observations.

2 A causal clause and a dilemma

Several theorists of the nature and properties of social entities have argued that groups are causally efficacious agents in their own right. Examples include arguments that higher-level causation *supervenes* on, relationally *emerges* from, or is *enacted* by lower-level causation (Elder-Vass, 2010; Hindriks, 2013; Sawyer, 2003), defences of multiple realisability-*cum*-realisation insensitivity of higher-level causal properties (Kincaid, 2009; List & Spiekermann, 2013; Menzies & List, 2010); or implicit admissions that efficacy is a prerequisite feature of agency (Gilbert, 1989; List & Pettit, 2011). In a broad sense, the test of efficacy for agency [TEA] could be articulated as a causal clause as follows:

[TEA] An entity A is a group agent only if.

- a) A performs an act (action, omission) that causes an outcome Ω (event or state of affairs in the world) or one of Ω 's causes.
- b) A's causal profile is irreducible² to the aggregated³ individual causal capacities c_1, c_2, \dots, c_n of its members.

[TEA] is an essential proviso for a number of reasons. Higher-level agency requires efficacy since groups venture into spheres of action that no member individually follows, understands, or sanctions. Agents are what agents do, i.e., an entity and the laws governing its behaviour are synchronised correlates. Second, group agents are said to exist as real entities, and for something to exist is for it to be causally efficacious (*as per* the Eleatic Principle⁴). Although the Eleatic Principle might be subject to controversy (Colyvan, 1998), no non-reductive theorist defends epiphenomenalism or takes higher-level agents to be causally inert like fictional or abstract objects (unicorns, ghosts, or sets). Third, non-reductionists often take groups to be morally responsible for concrete results (events or states of affairs), and the only way to impact the world is via causal capacities.

If not all theories of group agency explicitly endorse [TEA], the most prominent commit their adherents to it. Consider for instance social integrates views (crediting groups with representational and motivational states [List & Pettit, 2011]), status views (holding that a group is an agent in virtue of having a certain status projected on it [Hindriks, 2008; Searle, 1995]), or Epstein's (2015) 'grounding-anchoring-framing' view (holding that group agency requires individual members *plus* infrastructures, external constraints, membership requirements, existence conditions, assignment of positions of power, etc.). Strategy-wise, such theories extend key features of agency to groups⁵ and rely on the difficulty of deriving group-level facts from individual- or structural-level facts, on the concept of constitution⁶ to deliver cohesiveness without identity, and on the notions of supervenience or grounding⁷

² 'Irreducible' in the sense that A has distinct and novel causal significance, non-identical to and unaccountable for in lower-level terms, in a moderate rather than strong emergentist sense (as evidenced by arguments for social-level causation enacted by individual-level causation in a 'unity without identity' constitution relation (Hindriks, 2013); relational-emergent causation (Elder-Vass, 2010), realisation-insensitive high-level causal capacities (Kincaid, 2009; Menzies and List 2010); or holistic causal explanations (List and Spiekermann, 2013).

³ Although 'aggregation' is typically synonymous with 'mereological summation' (Wahlberg, 2013: p. 78), clause b) aims to convey distinctness in A's causal profile irrespective of whether A's members are pluralities, fusions, aggregates or sets (e.g., I presume that one may aggregate causal capacities in structured pluralities).

⁴ According to the Eleatic Principle, necessarily, if x is a concrete object, then x is causally efficacious (an entity x is real iff it participates in a causal process or brings about changes in the world).

⁵ The idea that systems exhibiting key features of agency count as agents is recurrent—groups are credited with psychological autonomy and rational unity (Pettit, 2003), normative autonomy Copp (2006), intentional states and second order reflection (Tuomela, 2005), rights and obligations (Hindriks 2008), reactive attitudes and sensitivity to moral reasons (Toleffsen, 2003), or action control (Gilbert, 2013).

⁶ Individuals (plus further constraints) *constitute* but they are not *identical* with the group agent. Group agency is *grounded* in the agency of individuals (plus further constraints), but not identical with them.

⁷ For two sets of properties A and B, A *supervenes* on B if there can be no difference in A without a difference in B. B *grounds* A if A obtains in virtue of (the more fundamental) B.

to deliver dependence without reduction. But endorsing a distinct intentional profile and independent reasoning appears to commit them to distinct causal efficacy. For non-reductionists would be reserved about scaling down the actional autonomy of a group agent to a simple authorised individual action with collective content, a merely ‘loaded’ low-level intervention. And they would be reserved about extending intentional states and reasoning capabilities to complex social entities just to rob them of the ability to act in a similarly independent fashion. Beside elevating groups to novel centres of judgment, representation, and intention, non-reductionists tend to admit that groups must ‘intervene suitably in the environment’, ‘act in the manner of an agent’, or ‘seek the realisation of motivations in the world’ (List and Pettit (2011), p. 20, 32).⁸

In reply, critics of non-reductive theories typically argue that groups fail the [TEA] clause (and consequently cannot be genuine agents), for higher-level efficacy entails a redundancy-identity dilemma:

Either the distinct causal efficacy of group agents competes with that of the networked individual members ($m_1 \cdot m_2 \dots m_n$) and systematically overdetermines or excludes their causal efficacy;

Or it is identical with (a subset of) the causal efficacy of its networked individual members ($m_1 \cdot m_2 \dots m_n$) and group agents have no distinct causal efficacy.

Noticeably, the dilemma extends to social entities the challenges faced by non-reductionisms about minds and material objects (Kim, 1993; Merricks, 2001), with ‘redundancy’ designating competing causes that might overload or nullify lower-level causal work (overdetermination or downward exclusion). Although prominent (see, e.g., the recent Elder-Vass (2010, 2014)—Wahlberg (2013, 2014) exchange), it is not a straightforward dialectical strategy. Several reasons come to mind. As regards overdetermination, it might be argued that it is not an exceedingly problematic notion (Sider, 2003); certain kinds of overdetermination come across as less problematic than others (Funkhouser, 2002)⁹; and overdetermining causes seem to

⁸ Some retain [TEA]’s provision (a) but deny endorsing a distinct causal edge (b). List and Pettit (2011) argue that supervenient group agency is compatible with methodological individualism, but this styled compatibility generates a functional sense of group agency and an epistemological sense of group autonomy, not an account of independent higher-level action. Likewise, Epstein (2017) holds that group agents are instances of a functional kind even when not performing that function, so they could still be agents even if they fail or are unable to act – but autonomous social structures are theoretically attractive mostly for their being systems of practical activity with genuine impact in the world.

⁹ In standard causal overdetermination (SCO) independent causes c_1 and c_2 are individually sufficient and work through separate mechanisms to bring an effect E about (e.g., firing squads [horizontal independence]; substance dualism with parallel physical causation [vertical independence]). In incorporated causal overdetermination (ICO), overdetermining causes share a connection and work through the same mechanism (e.g., stock market crashes when some investors sell off in a panic) (Funkhouser 2002). Although it is fair to say that SCO and ICO are problematic in different ways, SCO seems more problematic—independent overdetermination examples are outlandish (forest fires started by lightning and dropped cigarettes), and substance dualism seems more philosophically puzzling than supervenience/grounding theories (but others might disagree). So if group agents are fully autonomous, one needs to explain how they operate through their own independent devices; if they operate through their members, one wonders if there is an independent sense in which they make a difference in the world. (It is typically objected that it is mysterious how macro-causes coincide with micro-causes, or that parsimony justifies the exclusion of extra entities [Kim 1993; Merricks 2001]).

be, after all, ubiquitous (Schaffer, 2003). As regards upwards exclusion (when realisers exclude higher-level causes) (Kim 2007), it is (arguably) sound under a generative conception of causation, but (arguably) unsound under a difference-making conception of causation. As regards downward exclusion (when higher-level causes exclude their realisers) (Menzies & List, 2010), it is (arguably) unsound under a generative conception of causation, but (arguably) sound under a difference-making conception of causation (causation as difference-making allows ‘realisation-insensitive’ higher-level correlations [see Section §4 for more details]). As regards identity, it might be argued that group agents are not identical with but constituted of their suitably networked members. Higher-level agents turn out to be distinct from their constituents, while at the same time their causal efficacy at t is indistinguishable from the suitably networked members’s efficacy at t (Elder-Vass, 2014: p. 794).¹⁰ In addition, perhaps some non-reductionists would be tempted to avoid Ockham’s Razor-concerns and postulate group agents in a Quinean fashion, as indispensable entities of social ontology (List, 2018).¹¹ Or commit to group agents in order to avoid the spectrum of mereological nihilism.¹² Or hold that higher-level causation merely *redescribes* lower-level causation (Elder Vass, 2010; Wimsatt, 2006).¹³ Since all these options generate complications that ramify exponentially, the redundancy-identity dilemma forces critics of non-reductionism into a complex, multi-layered argumentative procedure that requires caution in the assumptions about overdetermination, the Eleatic Principle, and Ockham’s Razor; in the treatment of supervenience, epiphenomenalism, identity, or exclusion principles; and in the defence of stimulating but fairly counterintuitive positions such as mereological nihilism. However, there is a conceivable way around this argumentative approach. I believe that critics of non-reductionism may resort to a simpler, undeviating, and more persuasive reductio strategy, one that does not focus solely on whether groups *are* or *are not* causally efficacious.

¹⁰ However, if a possible world W where group agents exist alongside appropriately networked members were to be contrasted with a possible world W^* where no group agent emerges, we could not show that our actual world A is not in fact similar to W^* (Wahlberg 2014: p. 803). No causal test could substantiate it, for causal efficacy is alike, so by Ockham’s Razor there is no reason to posit group agency (and by the Eleatic Principle no imaginable W).

¹¹ However, a methodologically-motivated concept of group agency—an instrument for understanding, interacting with, and making demands on groups—undermines genuine higher-level causation and deviates from the robust idea non-reductionists initially advertised.

¹² However, mereological nihilism is not an indefensible position (Dorr and Rosen, 2002; Merricks, 2001; Unger, 1979). Intuitively, it is also not the position the *onus probandi* rests with. As Fodor’s remarks, “why there should be (...) macrolevel regularities *at all* in a world where, by common consent, macrolevel stabilities have to supervene on a buzzing, blooming confusion of microlevel interactions (...) [W]hy is there anything except physics?” (Fodor 1997: p. 161, original emphasis).

¹³ However, redescription amounts to iterative overdetermination and leads to impalement in the first horn of the dilemma. Iterative overdetermination typically concerns second-order properties overdetermining their first-order correlate realisers, e.g., a pill’s ‘dormitivity’ and its sleep-inducing chemical properties, both cause A ’s falling asleep (Funkhouser, 2002).

3 Mapping Causation

This new approach grants social entities independent causal efficacy and charts its implementation by appealing to analytic¹⁴ how-questions: (a) how is causal efficacy carried out, (b) how it relates to the causal efficacy of individual determiners, and (c) how it connects to the causal relevance of background structural factors. Presumably, a forthright ‘how it works’ anatomising strategy that locates causal efficacy in the architecture of complex aggregates will clear up non-reductive claims. The strategy partially overlaps with process-tracing understood as the search for intervening variables that connect independent with dependent variables in the modelling of social phenomena, and with process-observation as a complement to quantitative correlations and comparative methods for theory-testing or theory-building. The specific difference consists in mapping’s metaphysical focus, its concern with deep social structure, and its role as an alternative to the redundancy-identity dilemma, viz. it aims to evince that higher-level causation is minimisable and avertible rather than tautological or reducible.

Let a causal *why*-question about the occurrence of an effect E (event or state-of-affairs) be described as inviting a range of answers offering potential reasons for the occurrence of E. Call these answers reason-giving analyses of E. Each reason-giving analysis specifies a cause C or set of causes $\langle C_1, C_2, \dots, C_n \rangle$ as reasons for E. Each cause or set of causes are stipulated following a counterfactual conditional: Why did E occur? Because had C/ $\langle C_1, C_2, \dots, C_n \rangle$ not occurred, E would not have occurred. Example: *Why* did the HMS Eagle sink? Reason-giving analysis: It was torpedoed by a U-73 (cause), and had it not been torpedoed by a U-73, HMS Eagle would not have sunk (counterfactual conditional). Of course, it remains possible to offer alternative reason-giving analyses: it was at war; it was heavy and had a damaged hull, etc., and had HMS Eagle not been at war, had it not been heavy and damaged, it would not have sunk. And so on. Stated in a more compact way, for any why-queried effect E there is a set of possible causal answers that constitute reason-giving analyses for E following a counterfactual rule.

Let a causal *how*-question about the occurrence of an effect E (event or state-of-affairs) be described as inviting a range of answers that specify the process (or mechanism, or method) leading to the occurrence of E. Call these answers process-specification analyses of E. Each process-specification analysis stipulates a sequence of causal determinants $\langle d_1 \oplus d_2 \oplus \dots \oplus d_n \rangle$ leading to E, which describes gradually and systematically the emergence of E. How did E occur? It was preceded by a string of causally connected determining factors (events/states/actions/etc.), which eventually led to E. Example: *How* did the HMS Eagle sink? Process-specification analysis: A U-73 initiated a torpedo launch-sequence TS ($\langle x$ loaded the torpedo into the tube, shut and locked the breech door, turned on the power to activate the torpedo, flooded the tube, opened the pressure equalising valve, opened the muzzle door, operated the

¹⁴ Analytic how-questions target means, methods, and mechanisms (Jaworski 2009), in contrast to how-questions of manner (‘How did you do it?—Carefully.’) and cognitive resolution (‘How do we solve Molyneux’s problem?’).

water ram to eject the torpedo, etc. >) that lead to the sinking of HMS Eagle. Again, it remains possible to offer alternative process-specification analyses: for instance, the number of steps specified may be reduced, leading to a shorter analysis; or the analysis could be phrased in a different vocabulary, etc. In sum: for any how-queried effect E there is a set of possible causal answers that constitute process-specification analyses for E .¹⁵

Why-questions are typically associated with dependence intuitions and counterfactual accounts of causation [c causes e iff $\sim O(c) \rightarrow \sim O(e)$ (if c had not occurred, e would not have occurred)]. How-questions are typically associated with production intuitions and process accounts of causation [c causes e iff $O(c) \& P \rightarrow O(e)$ (if a process P —energy flow, momentum transfer—relates the occurrence of c to the occurrence of e)] (Hall, 2004). As mapping focuses on causal how-questions about processes and structure within macro-agency (and some conditions external to it), two caveats follow:

(i) Mapping arguments and the connected verdict on independent higher-level efficacy are more compelling under process or mechanistic accounts of causation. Not only that processes are natural answers to analytic how-questions, but, as an inquiry into implementation, mapping makes sense if there is ‘derivativeness’ or ‘causal fluid’ from c to e to be spelled out via analytic how-questioning. Process accounts are also well-motivated in this context, as paradigmatic agency appears to require the ‘real connectedness’ of a generative conception of causation (Kim, 2007: p. 236)—performing actions and bringing about changes in the world calls for spatiotemporally continuous sequences of causal determiners (what Hall (2004) calls ‘locality’).

(ii) If higher-level causal efficacy is defended under a counterfactual difference-making view (as it standardly is), mapping presses to the fore the intuition that it is still sensible to ask analytical questions about how that works (e.g., how higher-level claims such as ‘NATO’s intervention caused protracted instability in the Middle East’, ‘Right-wing media stokes social division’, ‘Human activities cause climate change’, or ‘State fragility causes civil war’ come to hold). How-questioning increases analytical depth. If there is higher-level causation, it makes sense to ask how it carries through to arrive at actual results (otherwise high-level causal capacities would come across as indistinguishable from epiphenomenal powers reaching results by way of action at a distance). It also seems sensible to eschew a fully ‘detail-disinterested holism’, as it would lead to an equally flat ontology (as the one holism tries to avoid). These may appear as distinct issues: whether *there is* higher-level efficacy is distinct from what lower-level mechanisms *implement* it (having implementation-mechanics is not per se a challenge to higher-level efficacy). However, if we discovered a) that higher-level efficacy operates through and obtains in

¹⁵ Beside ‘horizontal’ sequences, analytical how-questions may specify vertical, lateral, or diagonal processes (upward-running causal sequences must not be confused with the argument that causation is a kind of building relation, similar to realisation, constitution, fact-grounding, or micro-based determination [Bennett 2017]). Also, how-questions are directed at uncovering causal structures rather than explanatory sequences.

virtue of lower-level realisers, and b) that a detailed, structured geography of efficacy might be enough, or preferable, or more revealing—one could be deliberately agnostic about taking the further reductive step that would identify higher-level efficacy with its realisers and reject the ‘non-reductive’ element in ‘non-reductive efficacy’—then mapping could be seen as a kind of low-key challenge. Of course, if one is not persuaded by broader locality intuitions and the added analytical questioning, mapping remains contingent on a) generative (or hybrid) accounts of causation, or b) causal foundationalism, the view that difference-making facts depend on facts about physical connections (Ney, 2009). At the same time, the suggestion that it is reasonable to ask how group agents work *beside* questions regarding their identity, persistence, or parthood conditions remains independent of one’s causal theory.

Here are several reasons why mapping focuses on how-questions. First, looking at how causation carries through appears to be a better indicator of efficacy when compared to external attribution based on extended agential features. External attribution black-boxes implementation and renders current theories of group agency vulnerable to counterexamples involving coextensive membership. Imagine two higher-level agents A_1 and A_2 , based on a coincident set of members $\{m_1, m_2, \dots, m_k\}$, bringing about an outcome Ω . Should causal efficacy be attributed to A_1 ? A_2 ? A_1 and A_2 ? A_1 or A_2 ? External attribution alone will not suffice, only further information about causal set-up might clarify efficacy ascriptions.¹⁶ Second, unless one is solely interested in surface-correlations and outcomes in a purely consequentialist sense, one thinks that counterfactual dependence relations (determined via why-questions) cannot be central to the moral evaluation of social action. It is not enough to know what effects counterfactually depend on higher-level agents; we also need to know details about how things come about, the means, the methods, the mechanisms. Third, it makes sense to look at how efficacy is articulated given that higher-level agency absorbs fragmented circumstances, distinct individual constituents, incompatible perspectives, and different degrees of knowledge, understanding, and contribution. Although collective entities seem unitary, compact architectures of roles, they are controversially so in terms of causal efficacy. Fourth, the focus on how-questions fends off objections that one might oversimplify structure, action, precise history, and distinct moral considerations by choosing to look at an abstract, detail-free level. Fifth, since deficient causal claims and explanations can be either too abstract or too fine-grained, ‘just right’-level characterisations seem more likely if analysis were focused on implementation. Finally, how-questions should be interesting for redundancy-identity dilemma theorists, as overdetermination arguments are more powerful when they emphasize how redundant causal lines arrive at the same effect (in contrast to simply mentioning counterfactual dependence on higher-level agency).

¹⁶ Ditto for cases where powerful boards of trustees, specialized subcommittees, or influential individuals control the causal efficacy externally attributed to groups.

Now here is how it works. Consider the following group agency claim:

(1) The Roman Senate executed¹⁷ the leaders of the Catilinarian conspiracy.

Under a simple why-question—why were the leaders of the Catilinarian conspiracy executed?—the Roman Senate appears as causally efficacious: had the Senate not acted, the leaders of the Catilinarian conspiracy would not have been executed. Having assumed the causal efficacy of group agents, imagine the following question being asked:

Q1: *How did the Roman Senate actually carry out its decision to execute the leaders of the Catilinarian conspiracy?*

Under **Q1**, (1) converts into a sequence of more definite claims:

(2) The members of the Roman Senate decided to execute the leaders of the Catilinarian conspiracy.

(3) Cicero decided to execute the leaders of the Catilinarian conspiracy.¹⁸

(4) Cicero had the leaders of the Catilinarian conspiracy strangled in the Tullianum (an ancient building in the Forum).

(5) Executioners¹⁹ strangled with a cord the five leaders of the Catilinarian conspiracy (Lentulus, Cethegus, Statilius, Gabinius, and Caeparius), one at a time, in the Tullianum dungeon, a dark room with stone arches sunk about twelve feet underground.

(6) A hypoxic state of the brain following ligature strangulation caused the death of the leaders of the Catilinarian conspiracy.

If (1) confirms non-reductionist intuitions, **Q1** consistently yields individualist responses (2)–(5) [(6) illustrates non-agential conversion]. Although the Roman Senate is a relatum in common causal talk, actual causal efficacy—the actual strangling—is associated with those who made or implemented the decision.²⁰

Now one might worry that how-questioning relies on unjustified paraphrasing assumptions and artificially pressures for rephrasing perfectly sensible claims like (1). But this would be to mistake attention to how agency is exercised for an argument about the translation of collective action assertions into individual action assertions. It is one thing to inquire whether paraphrasing procedures support reduction, and another to understand how causal capacities are effectively carried out. Second, one might argue that answers (2)–(5) are compatible with (1) rather than

¹⁷ I take ‘executed’ to be a causative verb and understand (1) as ‘The Roman Senate **caused** the execution of the leaders of the Catilinarian conspiracy’ in virtue of the equivalence thesis, viz. ‘For any causative ϕ , $X \phi$ -ed = X caused a ϕ -ing to exist’ (Moore 2009: pp. 5–19). Ditto for the causatives in (2)–(5).

¹⁸ As consul, Cicero delivered four excellent orations against Catilina and convinced members of the Senate, including Caesar who proposed less excessive measures, that execution was the right course of action.

¹⁹ ‘Certain men, to whom orders have been given’ (Sallust 44-40BC/1899, Ch 55).

²⁰ By ‘actual strangling’ being associated with those who made/implemented the decision I do not mean to imply that ‘The Roman Senate strangled the leaders of the Catilinarian conspiracy’ is unsound. Rather, (a) it makes sense to say ‘yes, but how did the Roman Senate do it?’, and (b) it is counterintuitive to say that, e.g., ‘The Roman Senate wiped the sweat from its forehead and meandered down the dark, railless stairs of the frightful dungeon, where it quietly took a cord from its pocket and brutally strangled Lentulus.’ Lastly, ‘actual strangling’ is associated with those who made/implemented the decision *if we asked* analytical how-questions.

undermine its truth. But rather than evidencing explanatory compatibility amongst levels, the point of how-questions is to trigger analyses of why-answers with the aim of tracing the sources and mechanics of causal powers. Once all aspects of the processes underpinning higher-level actions are gradually exposed, group agents blend in with their causal surroundings. Reminiscent of Nagel's evanescent agency observation (Williams & Nagel, 1976), just as individual agency appears to dissipate in an event-causal framework, group agency appears to fade under **Q1**. The analogy conveys a change in perspective rather than fully erasing higher-level difference-making relations; it means seeing things through 'a more complete and precise account of the facts' (Williams & Nagel, 1976: p. 139), that is, Q1-type how-questioning might lead to sufficient, preferable, clearer alternatives—a perspective-change that comes with understanding structure and fundamentality. Third, one might differentiate between statements about agency [(1)] and statements about participation or contribution to an action [(2)–(5)]—presumably, an individual could participate or contribute to a collective action without being its agent. Yet (2)–(5) look like straight agential claims (the members of the Senate, Cicero, or executioners pass intentional and causal conditions for agency), so it is at least equally plausible to hold that several agents rather than a collective entity were causally efficacious in executing the leaders of the Catilinarian conspiracy. Lastly, perhaps it is not always possible to trace the individual actions or patterns of decision-making behind group-level events. Perhaps implementation is deliberately anonymised for confidentiality reasons, automatised for efficiency reasons, encrypted for security reasons, or simply too complex to pin down – so the Q1-type inquirer might not always have better, more accurate answers. Nonetheless, one may still argue that it is certain individuals who designed and oversee and may intervene to alter such coding procedures. Or one may still argue that it makes sense to ask how agency is exercised *independently* of whether implementation ultimately proves to be transparent or inaccessible. Or one may insist that, in most cases, group-level events are likely to be sensibly scalable to individualist terms rather than to some super-complex, untraceable, blind realisation. At this point imagine a further question being asked:

Q2: *How is the causal efficacy of the Roman Senate related to the causal efficacy of individual determiners (members, Cicero, executioners)?*

The question is purely rhetorical to those who associate causal efficacy with the executioners or with a complex causal sequence coordinating executioners' actions with Cicero's actions. But non-reductionists might answer **Q2** as follows:

(i) Higher-level efficacy is grounded in lower-level efficacy – the former obtains in virtue of the latter and the two cohere together in a layered configuration. The grounded cannot reduce to what grounds it, for irreflexivity (if p grounds q , then $p \neq q$) clashes with identity (if p reduces to q , then $p = q$). However, mapping theorists do not have to counter this proposed layered configuration by insisting that causal efficacy in (1) is *identical* with or *reduces* to that specified in (2)–(5). Two avenues of approach are open: first, instead of being linked to identity claims, grounding can simply convey a priority or essentiality message. Instead of holding that the causal efficacy of the Roman Senate (1) is identical with the causal efficacy in (2)–(5), it is enough to argue that (2)–(5) are more fundamental, or, following Rosen (2010), that (2)–(5) provide a 'real definition' of (1). If so, how-questioning

sidesteps identity claims and shows that (1) may be reduced to a more fundamental formula, or to its essentials via a more adequate analytic definition. Second, instead of reduction, grounding can simply favour an unallied stance. One may remain officially neutral on whether there really is higher-level causal efficacy. It might well be that groups are causally efficacious, but there is comparatively little concern for this view once sufficiently reasonable individualist answers are available. (1) might be true but given (2)–(5) the issue may simply be left undecided.

(ii) Both higher-level and lower-level determiners are causally relevant, as they both *control* one and the same action. List and Pettit (2011) offer an analogy: in a closed flask, water is brought to the boil. When the flask breaks, both the fact that the water was boiling *and* the position and momentum of a molecule triggering a break in the surface of the flask can be said to cause the breaking. Stated differently, the boiling of the water ‘programs for’ or makes it ‘inevitable’ that some molecule will trigger the breaking. By analogy, group agents control for the performance of the relevant individuals—and make their action ‘inevitable’—through the maintaining of procedures, role hierarchy, and task divisions, so both are causally relevant (pp. 161–164). This ‘co-authorship’ analogy is nonetheless deficient.

First, the flexibility to describe the breaking by citing events, facts, individual entities, etc. does not mean that a higher-level determiner (the boiling) and an individual determiner (molecule *x*) *both* do the causal work. Levels of *description* do not entail levels of *efficacy*. Second, the notion of *control* rests on an equivocation: if a certain kind of control is needed to maintain the internal structure of procedures, hierarchies, and role divisions, a completely different kind is involved in carrying things out.²¹ Third, the water boiling makes it inevitable that some molecule will trigger the breaking, but maintaining procedures or role divisions can be done *without* actually carrying things out. What is more, members might disagree with, resign from, or act against organisations.²² Fourth, why ‘control’? If there is a sense in which groups can be said to *cause* an outcome, it does not necessarily mean they also *control* it—not everything groups cause is controlled. If higher-level control is an independent notion, it is bound to raise the same issues as higher-level causation. Also, control is a causal notion—‘to control *x*’ = ‘to cause *x* to be controlled’—so the argument appears overtly circular.

Finally, imagine a further question being asked:

Q3: *How is the causal efficacy of the Roman Senate connected with the causal relevance of other background structural factors?*

Q3 makes sense in light of a more recent anti-individualist theme: actions undertaken by groups depend on much more than the actions of individual members, so the former cannot reduce to the latter. Current theories, Epstein remarks,

²¹ In addition, it might be argued that ultimately it is individuals who design and maintain procedures, hierarchies, or role divisions.

²² Alternatively, maintaining procedures or role divisions *raises the probability* that individuals will carry some action out, whereas the water boiling *makes it necessary* that some molecule will trigger the breaking.

(...) overlook all the other factors surrounding groups – infrastructures, external constraints, membership requirements, existence conditions, assignment of positions of power, and so on – which do not depend on the members and yet play key roles in their performance of practical activities. Once we abandon the dogma that facts about groups are exhaustively determined by facts about their members, groups and their attitudes take on a whole new look (Epstein 2016; pp. 133–134).

On Epstein's theory, if facts about individual members and their attitudes are amongst the *grounding* conditions of group agency, they are so in virtue of external social facts that *anchor* certain social rules that frame the grounding relation. It is one thing to ask what are the grounding conditions for the Roman Senate to be a group agent, and another to ask what sets up ('anchors') the grounding conditions for *x* to be a group agent. Not only that members are mere elements of a significantly larger grounding base, but group agency presupposes two distinct kinds of metaphysical determination.

Indeed, individual actions (2)–(5) are not a context-free, bare set of endeavours. But peculiarly enough, the argument makes it harder for non-reductionists to substantiate higher-level causal efficacy. It is more difficult to extract an independent sense of causal efficacy from a twice as sophisticated metaphysical framework, originally intended for explaining the emergence of social objects/kinds. Not only that anchoring supplements an already contentious grounding relation, but the verdict is still out on whether anchoring is a species of grounding (Epstein, 2019; Schaffer, 2019), or whether both dependence relations are instances of metaphysical causation. So as far as causal power is concerned, more work is needed to show why doubling on determination relations is helpful; what are the features of anchoring (is it an irreflexive, asymmetric, transitive relation?); and how anchoring is relevant to causal efficacy. Thus, although actions undertaken by groups depend on more than individuals, non-reductionists are not yet in a position to answer **Q3** decisively. Strengthened accounts of anchoring and its role may of course be advanced in light of the worries listed above. Until then, however, the connection between anchors and higher-level efficacy remains undertheorized and debatable.

In contrast, individualists have a simpler answer to **Q3** and a more straightforward narrative about anchoring features (for now, at least): spotlighting the causal relevance of salient determiners within a social structure such as the Roman Senate—positions occupied by members; constraints (social, biological, psychological, etc.); codes; resources; stakes; social rules and practices, and so forth—lends plausibility to a complex set of determiners rather than to an independent, monolithic collective entity. At a closer look, the social infrastructure around individuals is designed and coordinated by us, so individuals remain the structural and functional units within collective social entities and the genuine bearers of efficacy. Beside the simple claim that individuals and their interactions ground social agency, individualists might point out that facts anchoring practices, conventions, or rules are essentially dependent on the minds and actions of socially situated persons. Thus, anchoring conditions complement structures of individuals; and individuals design

the aforementioned anchoring conditions. Whether anchors are to be treated as mere grounds may well remain a topic for a later time.

Drawing the strands together, the *reductio* argument at the centre of the second strategy may be briefly articulated as follows:

- (1) Assume that groups are causally efficacious agents ‘in their own right’.
- (2) If genuine causal efficacy tends to be associated with complexes of structural and lower-level determiners under how-questioning in the Roman Senate case, then causation tends to be similarly articulated in every instance of group-level efficacy.
- (3) If how-questioning shows that causation is associated with complexes of structural and lower-level determiners in every instance of group-level efficacy, then groups fail the [TEA] test—groups are not causally efficacious agents ‘in their own right’ (or, in more moderate terms, the case for higher-level efficacy is weakened).
- (4) Causation tends to be associated with complexes of structural and lower-level determiners under how-questioning in the Roman Senate case.
- (5) Therefore, groups fail the [TEA] test—groups are not causally efficacious agents ‘in their own right’ (or, in more moderate terms, the case for higher-level efficacy is weakened).

Stated differently: let us assume that groups are causally efficacious ‘in their own right’. But when we ask how their causal efficacy is carried out, how it relates to the causal efficacy of individual determiners, and how it connects with the causal relevance of background structural factors, higher-order causation appears to collapse into, be attenuated to, or simply give way to lower-level causation. At minimum, there seems to be no substantial reason to endorse higher-level causal efficacy (despite persuasive instrumental reasons). Therefore, if causal efficacy tends to be associated with complexes of structural and lower-level determiners—intuitively or with a higher probability—groups cannot be said to be causally efficacious ‘in their own right’. In mapping the structure and scope of collective action, analytical how-questioning weakens the case for independent higher-level causation.

An advantage of the strategy is its simplicity: how-questioning reveals the implementation of causal efficacy independently of overdetermination- and Eleatic Principle-related arguments. Another advantage is its depth. The focus exclusively on whether groups *are* or *are not* causally efficacious black-boxes implementation, while attention to how causation is articulated increases the visibility of individualist arguments and countenances structure-oriented explanations. And yet another advantage is that it shifts attention from the agential capacities of social entities (and ontological building relations) to the way collectives actually perform—it is sensible to ask analytical questions about how *x* works, independently of questions about *x*’s identity, persistence, or parthood conditions. If much of the recent social ontology literature is centred on the constraints set forth for group agency, the metaphysics of social entities should give more consideration to causation. It is, after all, the only way collectives impact the world.²³ In what follows I look at a central, recurrent

²³ One of the consequences of focusing on causation is an increased attention to complexity and indeterminacy in the social world. Most large-scale social structures (nations, governments, markets) are not complicated, function-preforming, machine-like social entities, but *complex*: a large number of actors, in continuous activity, acting and reacting to each other and reacting back again, and defining a highly connected, constantly fluid state between structure and entropy. An effect of complexity is that top-down

argument for higher-level efficacy as a study case for the application of mapping principles.

4 Insensitive difference-making

Here is an alternative anti-reductionist route. Start from the idea of *multiple realisability*: a higher-level property ϕ is multiply realisable iff ϕ -instantiations can be realised by lower-level instantiations of two or more distinct types. Since ϕ has multiple lower-level realisations, it cannot be identical to lower-level properties, and has causal capacities distinct from and irreducible to those of its realisers. Second, add an account of causation compatible with the autonomy of special-science properties, viz. a counterfactual theory (C causes E iff C makes a counterfactual difference to E [$\sim O(C) \rightarrow \sim O(E)$ —if C had not occurred, E would not have occurred]). Counterfactual conditionals allow non-reductionists to claim that higher-level properties make a difference to some effect while their lower-level realisers do not. For example, Menzies and List (2010) argue that higher-level causal claims are *realisation-insensitive* i.e., irreducible to and true regardless of their lower-level realisation. Consider a causal relation $C \rightarrow E$, where C is a higher-level state realised by a lower-level state C_1 , but which could have been realised by any other C_1, C_2, \dots, C_n lower-level states. $C \rightarrow E$ is realisation-insensitive just in case E obtains in some close $\sim C_1$ -worlds that are still C worlds, i.e., E obtains if C is differently realised. Third, adapt the argument to support the causal efficacy of social entities. Imagine a well-structured group in which multiple distinct configurations of individuals H_1, H_2, \dots, H_n correspond to the same higher-level action ϕ – e.g., Royal Bank of Scotland’s takeover of the Dutch bank ABN Amro. In the actual world, RBS is described by a particular complex lower-level configuration H_{289} , which leads to takeover action ϕ . However, had its lower-level complex configuration not been H_{289} , it would have been in one or other possible configurations instead—say H_{467} or H_{672} —so it would still have achieved takeover ϕ (ϕ is realisation-insensitive). Thus, RBS as a collective entity makes a difference in the world independently of its contingent lower-level organisation. Although this line of reasoning is contentious, it provides a prime model case for the application of mapping principles.

(i) First, analytic how-questioning exposes the partial suitability of counterfactual conditionals, as one might argue that, despite realisation-insensitivity, the higher-level difference-making potential is derived from the foundational- and structural-level efficacy (its actual extension). At minimum, **Q1**-type how-questions show that higher-level counterfactual conditionals need to be complemented with details about causal processes (how causation carries through), while **Q2**-type how-questions

Footnote 23 (continued)

authority does not always work and may have unexpected and fortuitous consequences. Another is that many actions attributed to group agents are causally indeterminate and outcome responsibility cannot be accurately identified. Another is that major effects are difficult to address and adjust (conflicts, climate change, global shifts in migration patterns, etc.). And yet another is that individual actions may have significant causal force when a system reaches a critical state (e.g., protesters, whistle-blowers, etc.).

convey the priority of the latter or neutrality towards the first. At maximum, **Q1-Q2**-type how-questions highlight the tension between distinct, level-specific theories of causation—a counterfactual account that vindicates higher-level efficacy and a process account that details lower-level structures of events.²⁴ Consider a related example:

[E]: Deficiencies in RBS’s management and governance caused its collapse.

[E] is a lightly unpacked version of the more abstract ‘RBS caused its own collapse’. Under **Q1**-type how-questioning—how did deficiencies caused the collapse? we learn that several factors share causal prominence: weaknesses in capital position (due to bad management and weak global regulatory framework); uncertainties about RBS’s underlying assets (due to minimal analysis by the UK regulator); the aggressive takeover of ABN Amro (traceable to bad management); reliance on risky short-term wholesale funding (due to an inadequate approach to the regulation of liquidity, so bad management again); underestimation of losses in credit trading (by both management and regulators). Under **Q2**-type how-questioning—how is the force of deficiencies related to the efficacy of relevant individual determiners?—the collapse is traced to poor individual decisions by management and regulators, resulted from flawed analysis, judgment, and execution in specific circumstances. Under **Q3**-type how-questioning—how is the force of deficiencies related to the causal relevance of other background structural factors? – the collapse is placed in the more general context of the 2008 global liquidity crisis, the ‘light touch’ approach of the UK regulator, and the limited global regulatory framework. But only *some* banks failed, under equivalent regulations and crisis conditions, so the collapse of RBS comes down to decisions made by the RBS management (chief executive Sir Fred Goodwin, chairman Sir Tom McKillop, and board). The Financial Service Authority concurs in its ex post facto analysis: “These were decisions for whose commercial consequences the RBS executive and board are ultimately responsible.”²⁵ Throughout, the analysis does not minimise the causal force of structural conditions or dictate a ‘moralised’ selection of top individuals for legal responsibility purposes. If anything, it locates causal efficacy and articulates the partial and non-specific character of higher-level causal claims.

(ii) Second, analytic how-questioning shows that realisation-insensitivity relies on a limited understanding of social entities as architectures of roles rather than systems of individuals in roles. Insensitive difference-making obtains if change of individuals does not induce change in the identity and causal powers of social entities (e.g., RBS remains RBS regardless of its executive, board, and personnel). But it often does. Football clubs target valuable players and managers precisely because changes in team and management make a difference to their profile and value. Departments look for talented faculty members precisely because a change in staff makes a difference to raising their academic profile. Similarly, changes in governments make a fundamental difference to the character of states. Conversely, we find

²⁴ Current literature contains no explanation of why higher-level causal efficacy relies on dependence—rather than production-intuitions about causation.

²⁵ FSA report at <https://www.fca.org.uk/publication/corporate/fsa-rbs.pdf>.

it concerning when higher-level identity and causal powers must change but individuals do not, e.g., when the same officials who planned and implemented a hostile policy are tasked with designing a compensation programme for the victims of that policy. Replacing individuals in a well-defined system of roles matters and it does not amount to the blind realisation we often associate with contingent micro-physical configurations. It therefore makes sense to correlate RBS's collapse with a particular lower-level realisation involving Goodwin-McKillop-board, and argue that under an alternative realisation (different management) RBS would not have collapsed. Rather than being a multiple-realised single organisation, RBS's profile is essentially distinct in the two cases. The lesson is that although social entities can be abstractly described as autonomous configurations of roles, they are in fact complexes of individuals in roles plus wider structural factors.²⁶

(iii) Third, analytic how-questioning exposes the tendency of realisation-insensitivity supporters to run together causation (extensional) and causal explanation (intensional). For instance:

- 'It is entirely possible that social properties, not individual-level ones, display the most systematic *causal relations* in some social phenomena. The relation between interest rates and inflation, for example, may well be more robust than any particular individual-level transmission mechanism. This *explanatory holism*, however, does not imply any more radical form of holism that gives some kind of meta-physical priority to social structures over and above the individuals.'; '[We] defend (...) *causal-explanatory holism*: Some *causal relations* (of the kind that a social scientific explanation would describe) are distinct from (and not re-describable as) any individual-level *causal relations*' (List & Spiekermann, 2013: p. 629, 634 respectively, emphases added).

- 'A good way to see that macrosociological entities can stand in *causal relations* is to take well-known notions of causation and showing that they make sense of *macrolevel causation*—this strategy can be used to shake intuitions that causation at the social level is otiose or otherwise objectionable. So this is an argument against [methodological individualism] because it claims to show that *macrosociological causal explanations* are possible' (Zahle & Kincaid, 2019: p. 663, emphases added).²⁷

Q1–Q3-type how-questions show that higher-level efficacy is determined by and works through its 'realising' grounds and deeper causal armature. Although explanations are better if pitched at the 'right' level of specificity—e.g., it would be scientifically inappropriate to describe the kinematic behaviour of lenticular galaxies

²⁶ An aspect I cannot take up for reasons of space concerns the tension between special duties and role-based duties. On the one hand, acceptance of official roles entails prioritizing role-associated duties and responsibilities over personal interests. On the other hand, an institution cannot require its officers to do something that would be wrong for them all things considered (e.g., sacrifice themselves and their families if necessary). Finding a balance here is difficult. A sound institution cannot demand its role-bearers to perform actions that are morally wrong for them. But it cannot be sound if its principles were systematically overridden. It is important to note that action in official roles is bound up with action as a role-bearer and there is no possible description in which one acts in a role without acting as an individual.

²⁷ Further examples of running together causation and causal explanation appear in Sawyer (2003), List and Menzies (2010), List and Pettit (2011).

by citing the position and velocity of each star—explanatory holism and higher-level causal efficacy are distinct notions. Specificity-infelicitous lower-level explanations allow higher-level entities, properties, or actions to enjoy *explanatory* distinctness, not *causal* independence. Why the confusion? One reason could be that ‘causes’ is semantically ambiguous between the ‘real’ causal relation and the sentential connective specific to causal explanations (Davidson, 1980). However, supporters of realisation-insensitivity offer neither evidence for ambiguity, nor criteria for controlling disambiguation. Perhaps another reason could be adherence to relata liberalism – the view that causal relata are as diverse as common causal talk suggests and may include entities from different ontological categories²⁸ – but again, no defence of relata liberalism is offered. And yet another reason could be the appeal to higher-level interventionist counterfactuals. Although interventionism is basically a theory of explanation built on a version of a counterfactual theory of causation,²⁹ it appears to draw no explicit distinction between causation and explanation – causation is defined in relation to practical notions such as manipulation and control, which in turn underpin experimentation and good scientific explanations:

I draw no sharp distinction between providing a causal explanation of an outcome (hereafter the *explanandum-outcome*) and providing information about the causes of that outcome. According to the interventionist conception, when we provide such causal information we provide information that can be used to answer a *what-if-things-had-been-different* question: we identify conditions under which the *explanandum-outcome* would have been different, that is, information about changes that (...) might be used to manipulate or control the outcome. (Woodward, 2010: p. 291, original emphases).

However, since both causation and explanation are construed as systematic patterns of counterfactual dependence between chosen variables, higher-level causal efficacy only means that a higher-level variable (rather than a ‘realiser’-variable) happens to be, in a contingent context, specificity-felicitous and explanatory optimal. It does not mean that higher-level entities, properties, etc. cause things independently. It may be countered that specificity-felicitousness is sufficient for causal power on a ‘minimalist (...), relatively thin, unmetaphysical conception of causation’ (Woodward, 2015: 313) connecting changes in relevant variables. But supporters of realisation-insensitivity have yet to show why they need such a thin conception (barring its obvious suitability to realisation-insensitivity arguments) and how deliberately embracing it is supposed to generate non-circular, legitimate reasons for

²⁸ Relata liberalism offers flexibility as regards immanence (the question whether relata are intensional or extensional) and individuation (the question whether relata are coarse-grained [events] or fine-grained [facts, tropes, states of affairs]).

²⁹ According to interventionism, causal relata (c , e) are variables and causation relates changes in c -variables to changes in e -variables: c causes e just in case if the value of c were to change as a result of an idealized manipulation (called *intervention*), then the value of e would also change, *ceteris paribus* (Woodward 2003).

collapsing the causation–causal explanation distinction.³⁰ In addition, as Franklin-Hall (2016) had argued, the lack of constraints on variables in the interventionist causal-explanatory programme cannot recommend high-level claims as explanatory optimal (criteria such as proportionality, exhaustivity, and stability are unsuitable guides to explanatory optimality). So freedom to shift frameworks and adjust the granularity of causal claims for purely instrumental reasons (explanatory economy, division of labour, etc.) does not equate with independent efficacy. Recurrent examples of irreducible higher-level correlations (monetary policy–inflation, poverty–delinquency, carbon emissions–climate change, technological change–unemployment, etc.) are in fact claims that a certain explanatory level is deemed optimal. They do not strike us as false causal assertions because they are appropriate explanatory assertions. And it makes sense to press further how-questions all the same—were actual causal detail not needed, why is it natural and legitimate to ask for it in the analysis of efficacy (i.e., *after* highlighting the specificity and explanatory role of higher-order properties like robustness, probability, or function)? Finally, higher-level causal claims are seldomly self-standing. As the RBS example shows, they are frequently supplemented by or translated as more precise causal claims.³¹

(iv) Fourth, analytic how-questioning allows a closer look at the specifics of multiple realisability. Here are some examples. If ‘insensitive’ higher-level efficacy obtains in pluri-realizable social entities, and the number of realising configurations gradually reduces in less sophisticated organisations, then higher-level efficacy seems to be directly proportional with the complexity of groups. But when is a social entity sufficiently complex to sustain insensitive efficacy? Assuming that one higher-order property realised by two lower-level properties is insufficient (two grains do not make a heap), at what point does mentioning realisation-insensitivity and higher-level difference-making switch from inappropriate to appropriate? Non-reductivists offer no theoretical treatment of vagueness issues in the definition of insensitive social entities. Furthermore, multiple realisability seems compatible with some sort of reduction. Just as paperweights—realised by bricks, cast bronze whales, marble cuboids etc.—are reducible in principle to one causal factor (the property of

³⁰ Running together causation (extensional) and causal explanation (intensional) constitutes a compelling criticism contingent on one’s theory of causation, as Woodward may have legitimate reasons to collapse the distinction. However, if there are such reasons, they should figure in the realisation-insensitivity theorist’s argument for higher-level efficacy. At the same time, it is worth noting that most difference-making views appear to draw such a distinction – counterfactual theories target a concept of causation untainted by the pragmatic considerations specific to explanatory angles (Lewis 1973); contrastivism about explanation (Lipton 1990) is different from contrastivism about causation (Schaffer 2005); and interventionists tend to acknowledge that ‘[c]ausal relationships are features of the world: they are “out there” in nature. By contrast, explanation is an activity (...) having to do with the discovery and provision of *information*, information about causal relationships.’ (Woodward 2003: p. 23). Thus the ‘running together’ criticism would be in general worth signalling when higher-level efficacy is defended with difference-making views.

³¹ Another approach takes explanatory facts as prior to causal claims (Strevens 2008). If causal claims are typically given priority over the explanatory – explaining Ω requires citing causes of Ω – Strevens holds that x caused y is a shorthand for x [partly] explains why y occurred: ‘[C]ausal claims (...) are causal-explanatory claims’ (Strevens 2008: p. 4). If correct, however, we end up with a theory of explanation calling for optimality criteria (but not with independent higher-level causal efficacy).

exerting a force strong enough to prevent papers from being scattered), higher-level efficacy might be equated with more fundamental, ‘essentialised’ action principles (for instance, RBS’s efficacy could be theoretically condensed to root principles of retail and investment banking). Lastly, Q2–Q3-type how-questions offer an exploded axonometric view of realisation which highlights lower-level and structural, rather than higher-level, causal claims—social entities are realised by an individualist grounding base (horizontal realisation) *and* by further subsystems and wider structural conditions (three-dimensional realisation).

These points suggest that when applying mapping principles to insensitive difference-making arguments, the idea of independent higher-level efficacy begins to look less persuasive. Complex social entities retain contextual explanatory relevance, but sequential analytic how-questioning weakens insensitive difference-making by breaking down and clearing up what causes what. Naturally, this is not to imply that difference-making ceases to play a central role in ordinary causal reasoning, or that counterfactual dependence cannot drive an analysis of causation, or that abstract correlations are suddenly extravagant.³² The aim is not to deliberately adopt a process-premised approach and exploit conflicting intuitions in our concept of causation to dispute higher-level efficacy. Rather, the specific message is that ascriptions of group-level causation presuppose pragmatic elements and disguise explanatory claims. The broader message is that it is inaccurate to construe causation in the social domain as exclusively about uncovering systematic dependence relations between macro-level independent variables (causes) and dependent variables (effects) via why-questions (particularly in the light of issues with variable-choice and variable-calibration gripping interventionism and causal modelling). Most elaborate social theory is about accessing implementation, i.e., about determining means, detecting methods, revealing processes.

5 Concluding remarks

Standard accounts of higher-level causal efficacy are typically countered with redundancy-identity dilemmas. Nonetheless, critics of non-reductionism may resort to a simpler and more persuasive *reductio* strategy to show that collective social entities fail the test of independent efficacy: allow that group agents are efficacious and chart (a) how their causal efficacy is carried out; (b) how it relates to the causal efficacy of individual determiners; (c) how it connects to the causal relevance of background structural factors. Instead of appealing to redundancy-identity dilemmas, critics only need to show that progressive analytic how-questioning tones down the case for higher-level efficacy (somewhat similar to the way a magic trick loses

³² The intuition that it worth asking how things work is not meant as an expression of radical dissatisfaction with counterfactual accounts—the view defended here is not the basic reductionism that sees e.g., fish as swimming collections of living cells interrelated fish-like. Beside the hope for a methodological realignment towards efficacy, seeing groups as ‘structured efficacious pluralities’ aims to achieve an elevation that comes close to the one espoused by non-reductive theorists, but stops short at perceiving and committing to independent higher-level efficacy.

its compelling character once we understand how it was done). In particular, non-reductive accounts attempting to make light of the causal-explanatory divide are still pressured by process intuitions and the imperative to spell out implementation. Once spelled out, attention to how causation unfolds increases the visibility of individualist arguments and encourages structure-oriented explanations. While it offers reasons for favouring individualism,³³ mapping is not meant as a definitive tactical manoeuvre in the offensive against anti-reductionism—as mentioned in the opening section, further anti-reductionist moves in response to analytical how-questioning are of course possible. The gauntlet is down. Until then, however, analytic how-questioning lends plausibility to describing higher-level action as a matter of there being individual agents interlocked in complex infrastructures of roles, augmented by codes, constraints, resources, stakes, social rules and practices, all sharing specific relevance in the concatenation of causal paths to target events. Two key corollaries follow. One is a different methodological stance in social theory: mapping procedures advance current debates by motivating a general shift in analysis from the constraints set forth for group agency to examining how collectives impact the world. Another is a fictionalist stance towards group agency talk: mapping procedures imply that complex social entities are instrumental abstractions with a pivotal role in social explanation, prediction, and design. Of course, it may be that realism about group agency and causal efficacy is ultimately semi-fictional and special, like realism about viruses, numbers, quantum wave functions, forces of selection, or phonemes of Proto-Indo-European,³⁴ however, further argument would be required to defend this view.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article

³³ In this context, is individualism a negative claim (group-level facts *are not* causally efficacious), or a positive claim (individual-level facts *are* causally efficacious)? Since mapping looks at the how-terrain beyond group-level facts and disputes the idea that social entities have independent, novel, overriding causal capacities, it is mainly a negative claim. At the same time, given a disjunctive framing of the holism-individualism debate, a negative claim about the former entails a positive claim about the latter. If holists take higher-level efficacy to supervene on individual-level efficacy, then by tollendo ponens ($p \vee q \ \& \ \neg q \rightarrow p$) individualism emerges as a positive claim: ‘The hangman, not the state, executes a criminal’ (von Mises 1949: p. 41). Agency and action are positively associated with the individual-level if it is more accurate, persuasive, or revealing to say ‘[t]he hangman, not the state, executes a criminal’ (cf. fn. 20). Should a positive claim trigger concerns about explanatory regress and the redeployment of mapping against individual-level causation? Not exactly. Although a comprehensive discussion of regress would go beyond this paper’s scope, it is important to note that a) full-scale regress arguments are themselves negative claims that offer no support to holistic views; b) how-questioning individual-level causation is in certain contexts natural, as individuals have their own causal histories (e.g., courts might analyse potentially exonerative neurological evidence when assessing criminal responsibility). If the point is not to overlook individual agency, then surely one must not overlook how it is shaped by biology, society, chance, or God.

³⁴ Examples borrowed from Christopher Hitchcock (interviewed by Richard Marshall, <https://316am.site123.me/articles/causation-probability-and-philosophy>).

are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bennett, K. (2017). *Making Things Up*. Oxford University Press.
- Colyvan, M. (1998). Can the Eleatic Principle be Justified? *Canadian Journal of Philosophy*, 28(3), pp. 313–336.
- Copp, D. (2006). On the Agency of Certain Collective Entities: An Argument from 'Normative Autonomy.' *Midwest Studies in Philosophy*, 30(1), pp. 194–221.
- Davidson, D. (1980). *Essays on Actions and Events*. Clarendon Press.
- Dorr, C., & Rosen, G. (2002). Composition as a Fiction. In R. Gale (Ed.), *The Blackwell Guide to Metaphysics* (pp. 151–174). Blackwell.
- Elder-Vass, D. (2010). *The Causal Power of Social Structures: Emergence, Structure and Agency*. Cambridge University Press.
- Elder-Vass, D. (2014). Redescription, Reduction and Emergence: A Response to Tobias Hansson Wahlberg. *Philosophy of the Social Sciences*, 44(6), pp. 792–797.
- Epstein, B. (2015). *The Ant Trap*. Oxford University Press.
- Epstein, B. (2016). Précis of The Ant Trap. *Journal of Social Ontology*, 2(1), 125–134.
- Epstein, B. (2017). What are Social Groups? Their Metaphysics and How to Classify Them. *Synthese*, 196(12), pp. 4899–4932.
- Epstein, B. (2019). Anchoring versus Grounding: A Reply to Schaffer. *Philosophy and Phenomenological Research*, 99(3), pp. 768–781.
- Fodor, J. A. (1997). Special Sciences: Still Autonomous After All These Years. *Noûs*, 31(11), pp. 149–163.
- Franklin-Hall, L. R. (2016). High-Level Explanation and the Interventionist's 'Variables Problem.' *The British Journal for the Philosophy of Science*, 67(2), pp. 553–577.
- Funkhouser, E. (2002). Three Varieties of Causal Overdetermination. *Pacific Philosophical Quarterly*, 83(4), pp. 335–351.
- Gilbert, M. (1989). *On Social Facts*. Princeton University Press.
- Gilbert, M. (2013). *Joint Commitment: How We Make the Social World*. Oxford University Press.
- Hall, N. (2004). Two Concepts of Causation. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and Counterfactuals*. (pp. 181–276). MIT Press.
- Hindriks, F. (2008). The Status Account of Corporate Agents. In H. B. Schmid, K. Schulte-Ostermann, & N. Psarros (Eds.), *Concepts of Sharedness—Essays on Collective Intentionality*. (pp. 119–144). Ontos Verlag.
- Hindriks, F. (2013). The Location Problem in Social Ontology. *Synthese*, 190(3), pp. 413–437.
- Jaworski, W. (2009). The Logic of How-Questions. *Synthese*, 166(1), pp. 133–155.
- Kim, J. (1993). *Supervenience and Mind*. Cambridge University Press.
- Kim, J. (2007). Causation and Mental Causation. In B. P. McLaughlin & J. D. Cohen (Eds.), *Contemporary Debates in the Philosophy of Mind*. (pp. 228–242). Blackwell.
- Kincaid, H. (2009). Causation in the Social Sciences. In H. Beebe, C. Hitchcock, & P. Menzies (Eds.), *The Oxford Handbook of Causation*. (pp. 726–743). Oxford University Press.
- Lewis, D. (1973). Causation. *The Journal of Philosophy*, 70(17), 556.
- Lipton, P. (1990). Contrastive Explanation. In D. Knowles (Ed.), *Explanation and its Limits, Royal Institute of Philosophy Supplement*. (Vol. 27, pp. 247–266). Cambridge University Press.
- List, C., & Pettit, P. (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press.
- List, C., & Spiekermann, K. (2013). Methodological Individualism and Holism in Political Science: A Reconciliation. *American Political Science Review*, 107(4), pp. 629–643.
- List, C. (2018). What is it Like to be a Group Agent? *Noûs*, 52(2), pp. 295–319.

- Ludwig, K. (2007). Normative Autonomy for Collective Agents. *Journal of Social Philosophy*, 38(3), pp. 410–427.
- Mackie, J. L. (1980). *The Cement of the Universe: A Study of Causation*. Clarendon Press.
- Menzies, P., & List, C. (2010). The Causal Autonomy of the Special Sciences. In C. McDonald & G. McDonald (Eds.), *Emergence in Mind*. (pp. 108–128). Oxford University Press.
- Merricks, T. (2001). *Objects and Persons*. Oxford University Press.
- Moore, M. S. (2009). *Causation and Responsibility*. Oxford University Press.
- Ney, A. (2009). Physical Causation and Difference-making. *The British Journal for the Philosophy of Science*, 60(4), pp. 737–764.
- Pettit, P. (2003). Groups with Minds of their Own. In F. Schmidtt (Ed.), *Socializing metaphysics: the nature of social reality*. (pp. 167–193). Rowman and Littlefield.
- Rosen, G. (2010). Metaphysical Dependence: Grounding and Reduction. In R. Hale & A. Hoffman (Eds.), *Modality: Metaphysics, Logic, and Epistemology*. (pp. 109–136). Oxford University Press.
- Sallust (44–40 BC/1899). *Conspiracy of Catiline*. Rev. John Selby Watson, M.A. New York and London. Harper & Brothers.
- Sawyer, R. K. (2003). Nonreductive Individualism part II: Social Causation. *Philosophy of the Social Sciences*, 33(2), pp. 203–224.
- Schaffer, J. (2019). Anchoring as Grounding: on Epstein's the Ant Trap. *Philosophy and Phenomenological Research*, 99(3), pp. 749–767.
- Schaffer, J. (2003). Overdetermining Causes. *Philosophical Studies*, 114(1–2), pp. 23–45.
- Schaffer, J. (2005). Contrastive Causation. *Philosophical Review*, 114(3), pp. 327–358.
- Searle, J. R. (1995). *The Construction of Social Reality*. The Free Press.
- Sider, T. (2003). What's So Bad About Overdetermination? *Philosophy and Phenomenological Research*, 67(3), pp. 719–726.
- Strevens, M. (2008). *Depth: An Account of Scientific Explanation*. Harvard University Press.
- Tollefsen, D. P. (2003). Participant Reactive Attitudes and Collective Responsibility. *Philosophical Explorations*, 6(3), pp. 218–234.
- Tuomela, R. (2005). We-intentions Revisited. *Philosophical Studies*, 125(3), pp. 327–369.
- Wahlberg, T. H. (2013). Elder-Vass on the Causal Power of Social Structures. *Philosophy of the Social Sciences*, 44(6), pp. 774–791.
- Unger, P. (1979). There Are No Ordinary Things. *Synthese*, 41(2), pp. 117–154.
- von Mises, L. (1949). *Human Action: A Treatise on Economics*. William Hodge and Yale University Press.
- Wahlberg, T. H. (2014). Causally Redundant Social Objects. Rejoinder to Elder-Vass. *Philosophy of the Social Sciences* 44 (6), pp. 798–809.
- Williams, B.A.O. and Nagel, T. (1976). Moral Luck. *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vol. 50, pp.115–152.
- Wimsatt, W. C. (2006). Reductionism and its Heuristics: Making Methodological Reductionism Honest. *Synthese*, 151(3), pp. 445–475.
- Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.
- Woodward, J. (2010). Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation. *Biology and Philosophy*, 25(3), pp. 287–318.
- Woodward, J. (2015). Interventionism and Causal Exclusion. *Philosophy and Phenomenological Research*, 91(2), pp. 303–347.
- Zahle, J., & Kincaid, H. (2019). Why be a Methodological Individualist? *Synthese*, 196(2), pp. 655–675.