Journal of Ethics and Social Philosophy
VOL. 30, NO. 7 · DECEMBER 2025
https://doi.org/10.26556/jesp.v30i7.4429

ELICITORY STRUCTURAL POWER AND AGENTIAL POWER

AN OUTLINE AND DEFENSE

Arash Abizadeh

THE THESIS that some intentional agents have structural power clashes with two powerful intuitions: that the power of agents operates only by way of their intentional actions and that it consists in the capacity to play some causal role in effecting outcomes. These intuitions suggest that notions of passive and noncausal power illicitly conflate power with prospects for success. As Brian Barry argues, cases in which "desired outcomes occur with no intervention" by agents do not evidence their power because one should "distinguish between those who do well by exercising power and those who are the passive beneficiaries of the activities of others."[1] Similarly, Keith Dowding asserts that being the passive beneficiary of social structures is a matter of "systematic luck," not power.[2]

Yet these intuitions clash with familiar phenomena. We often recognize that it is precisely because some are so powerful that they can satisfy their wishes without having to lift so much as a finger. Picture a mafia godfather whose henchmen, with nary a word from him, anticipate and carry out what they take to be his wishes. His power is the envy of every mafioso reduced to monitoring the details himself—and reduced to pleading, cajoling, and bribing others to secure their service. Or think of a charismatic prophet who so inspires admirers that they anticipate and serve her every wish even before she herself becomes conscious of them. The power she wields over people is the dream of every

1 Barry, "The Uses of 'Power,'" 348.

2 Dowding, *Rational Choice and Political Power*. Terence Ball suggests that exercising power conceptually implies causally efficacious intentional action ("Power, Causation, and Explanation," 211). For the view that power must operate by way of actions, see Laswell and Kaplan, *Power and Society*, xiv; Simon, *Models of Man*, 11; Dahl, "Power," 410; Goldman, "Toward a Social Theory of Power," 225–26; and the overview in Ball, "Models of Power." For the view that power implies causation, see Simon, *Models of Man*, 11; Dahl, "Power," 410; March, "An Introduction to the Theory and Measurement of Influence," 437; Nagel, *The Descriptive Analysis of Power*, 11; and Isaac, *Power and Marxist Theory*, 74.

would-be spiritual leader. Or consider the structural position of men in our patriarchal societies: a man can often satisfy his preferences or aims in a way a woman cannot, sometimes thanks not to any intentional actions of his own but to society's gender structure, regulated by norms and epistemic frames that dispose others to defer to men's preferences or aims. Something similar is true of white people in European and European-settler societies.

My agenda here is fourfold: to show that active or *agential* power can be efficaciously exercised by way of intentional actions even when the outcomes would obtain without those actions; to defend a notion of *elicitory* power as a type of nonintentional or passive power by which agents elicit welcome outcomes, but not by way of their intentional actions; to defend a notion of *structural* power as a type of elicitory power that agents have in virtue of their positions in social structures; and to defend a *noncausal* category of power—whether agential or elicitory—at stake when outcomes obtain in virtue of an agent's power but without the agent having played any causal role in producing them.[3]

Many actors at the bottom of our societies' power hierarchies intuitively recognize structural modes of power, which they frequently articulate by drawing on the vocabulary of "privilege."[4] I argue, however, that these phenomena are often properly understood as instances in which privileged actors have and wield social *power*. This matters because privilege implies inegalitarian, hierarchical power relations, and while many instances of structural power—those grounded in gender or racial hierarchies, for example—fit this mold, other instances may be widely distributed, equally shared, or reciprocal. The recognition of nondecisive, elicitory and structural, and noncausal categories of power is significant because it serves the practical purpose of identifying those over whom it would be useful to wield power; the moral purpose of assigning responsibility and blame; and the evaluative purpose of critically appraising social arrangements in light of their distribution of power.

Readers be forewarned: precisely because moral responsibility is often premised on causal responsibility, I deploy a series of assassin cases below to tease out and clarify latent intuitions about the latter by appealing to clearer judgments about the former. But a drawback is that assassin-style cases can give the misleading impression that power is inherently power over others and hence hierarchical, or that *power-over* is inherently evil. Neither is the case. The concept of power defended here is not a moralized one: having or wielding power is not inherently wrong or evil, and even power-over may be welcome or beneficial to those over whom it is wielded.

3 On active versus passive power in the sense I employ here, see Morriss, *Power*.

4 Harris, "Whiteness as Property."

I have three terminological notes and one methodological comment before proceeding. First, the concept of power concerning us here is not the concept to which metaphysicians refer in phrases such as "the causal powers of entities," which we might call *entity power*. We are here exclusively concerned with the power *of* agents *qua* intentional agents, which we might label *agent power*—the core notion of which is the capacity to obtain what one might want. Both agential and elicitory power are species of agent power in this sense. I therefore use the term 'agent' strictly in the action theory sense of an *intentional* agent, not, as some metaphysicians do, to refer to any entity that produces changes in other entities' kind, structure, causal powers, or intrinsic properties (as in the phrase 'causal agent'). Concomitantly, by elicitory or "passive" power I mean the agent power that intentional agents have independently of their intentional actions, not, as these metaphysicians do, the power of an entity to suffer fundamental changes.[5]

Second, the concept of structural power defended here is distinct from three other similarly labelled concepts. 'Structural power' is sometimes used to refer to the power *of* social structures.[6] Social structures might be said to have causal power, for example, insofar as they constitute what Fred Dretske calls "structuring causes," i.e., "background conditions that enable one thing to cause the other" were the former to occur.[7] (Dretske contrasts structuring causes, which establish potential causal pathways, to "triggering causes," which are events that cause the first element of a causal process to *occur*—now.) I set this alternative usage aside, not because the concept is unimportant but because my concern here—whether in the case of agential or structural power—is with the social power of agents *qua* intentional agents. By analogy with agent power, I call the power of structures *structure power*. 'Structural power' is also sometimes used to refer to the power of agents *over* social structures, to create or shape them.[8] By analogy with structuring causes, I call this *structuring power*. Finally,

5 Marmodoro, "Aristotelian Powers at Work"; and Kuykendall, "In Defense of the Agent and Patient Distinction." The question of whether intentional agents have elicitory power does not turn on whether there is a viable ontological distinction between "active" and "passive" causal powers (and between causal agents and patients) in this metaphysical sense. Cf. Heil, *The Universe as We Find It*, 74; and Ingthorsson, "Causal Production as Interaction." Nor by the causal power of agents do I mean, as defenders of "agent causation" do in metaphysical debates about free will, agents' capacity directly to cause action-triggering intentions. See O'Connor, *Persons and Causes*.

6 Hayward, *De-Facing Power* and "On Structural Power"; Elder-Vass, *Causal Power of Social Structures*; Forst, *Normativity and Power*; and Hasan, "Republicanism and Structural Domination." Cf. Dowding, *Rational Choice and Political Power*, 8–9.

7 Dretske, *Explaining Behavior*, 42.

8 Strange, *States and Markets*; and Roy, *Socializing Capital*.

some use the term to refer simply to the power that an agent has in virtue of their position in a social structure.[9] This is a comprehensible sense of structural power—call it the "broad" sense. But my focus is on a more restricted sense because we are here concerned with *social* power, and almost all social power is ultimately structural in this broad sense.[10] I take *social structures* to be constituted by social relations to some extent stabilized by a set of background expectations, rules, norms, schemas, or practices enacted by agents operating within a certain *habitus*—agents whom Thomas Wartenberg calls "peripheral agents" and Nicholas Vrousalis calls "regulators."[11] Few instances of social power do not depend on social structures—and hence are not structural to some degree—in this sense. The broad sense is therefore largely redundant. (Even the power to lift boulders is, in a social context, a power one has partly in virtue of the fact that other agents are not disposed to prevent one from lifting boulders—because of the property regime, for example.) Thus two conditions must be satisfied for power to be structural in the strict sense at stake here: it must be in virtue of the agent's social-structural position, yes, but it must also be elicitory power, i.e., not operate by way of one's intentional actions.

Third, I do not mean the verb in phrases such as 'to effect an outcome' to be a synonym for *causing* an outcome: to *cause* is to *affect* what happens, but to *effect* what happens is, as I use the term, to realize, or, as the *Oxford English Dictionary* puts it, "accomplish" a preference, intention, objective, or plan.[12] Causally affecting an outcome is neither sufficient nor necessary for accomplishing one's intentions: one may cause unwelcome outcomes; and, as I argue, one may accomplish one's intentions or objectives without causing them. I restrict the terms 'effecting' and 'exercising power', moreover, to the case of *agential* power. By contrast, when outcomes are obtained by way of an agent's *passive* power (and so not by way of their intentional actions), I say the agent *elicits* (rather than effects) the outcome—which is why I label this type of power *elicitory*. And I use the expression 'to *obtain* an outcome' indifferently between actively effecting or passively eliciting. Finally, I use the verb 'wield'—as in the phrase 'wielding power'—indifferently between the intentional, active and the nonintentional, passive modes of power: one may "wield" agential or elicitory power, but one

9 Marsh, "Interest Group Activity and Structural Power"; Isaac, *Power and Marxist Theory*; Gädeke, "Does a Mugger Dominate?"; and Vrousalis, "The Capitalist Cage." Cf. Stone, "Systemic Power in Community Decision-Making"; Ward, "Structural Power"; and Haslanger, "Oppressions Racial and Other."

10 Isaac, *Power and Marxist Theory*.

11 Bourdieu, *La domination masculine*; Wartenberg, *The Forms of Power*; and Vrousalis, "The Capitalist Cage."

12 Morriss, *Power*, 29–30.

may "exercise" agential power only. (In using 'wield' in this way, to include a passive sense, I am resurrecting its obsolete meanings, which, according to the *OED*, include "to have the . . . advantage of" and to "accomplish" or "achieve, attain.")

These are of course stipulations on my part—terms of art used in the service of articulating a theory of social power. So my defense of that theory does not rely on appealing to readers' linguistic intuitions about such terms. Rather, my method is to present paradigmatic cases that recognizably instantiate agents' social power, to guide our considered judgments about the concept; and to interpret those cases in light of what I take to be the core notion of agent power. To this core we now turn.

1. The Welcome, Intentional Action, and Linkage Tests

Consider a victim who would not have been robbed had he not displayed his wallet.[13] The victim's action of pulling out his wallet to give alms is one of the robbery's antecedent causes: but for his actions, the mugging would not have occurred. Yet although the victim helped cause the outcome, he did not *exercise power* (over the thief, for example) to effect it. Similarly, consider an aspiring leader who, having commanded her would-be subjects to march, is served their loud refusal. Although the aspiring leader has *caused* them to do something—loudly refuse—she has *failed* to efficaciously exercise power: she did not effect the refusal.[14]

Why? Because the notion of power concerning us here is not the metaphysical concept of entity power (*qua* causal entity) but rather agent power (*qua* intentional agent). At the heart of this latter notion is the *capacity to obtain what one might want*. The distinctive feature of efficaciously wielding agent power—whether agential or elicitory—is that the outcome must be welcome to the agent: she must favor it in some sense. The hapless victim has not *ex post exercised* power to effect his own mugging because the agent-outcome relation fails what I call the *welcome test*. To be sure, given the presence of muggers, *ex ante* the agent does *possess* such power, and on other occasions, he might exercise it. Imagine the "victim" were a police officer conducting an undercover sting on a known mugger: then the officer would indeed be exercising his power to get the mugger to mug him.

The welcome test comes in both an *ex post* and an *ex ante* version. When we retrospectively inquire whether an agent has efficaciously *wielded* power to obtain an outcome, we apply an *ex post* test, to wit: Was the outcome favored by

13 Morriss, *Power*, 29.

14 Ball, "Power, Causation, and Explanation," 205.

the agent? The *ex ante* version, by contrast, pertains to whether an agent prospectively *has* the power to obtain a potential outcome. We therefore apply a counterfactual test, to wit: If the agent were to favor the outcome, would it obtain?

The welcome test applies to all forms of agent power: it articulates the core notion underlying both agential and elicitory power. It might be objected that this construal fails to account for *unwelcome* power. The objection rests on a mistake. Consider John Stuart Mill, who lamented the arbitrary power Victorian legal structures gave him over his wife.[15] Although *possessing* this power was unwelcome, its *nature* consisted in a capacity to obtain outcomes concerning his wife should he favor them. Mill may have deplored being able to get his way—and perhaps refused to exercise a power he deemed unjust—but his power over his wife consisted in being able to do so. Of course, in one respect, he was powerless: he could not free her from conjugal subjection without systemic legal reform. But the fact that he was powerless to relinquish his power over her does not imply he lacked it.

If the welcome test is common to both agential and elicitory power, the former's distinctive feature is its intrinsic link to exercising intentional *agency*. This has two aspects. First, agency is manifested in intentional actions, i.e., actions constituted by an intention-in-action.[16] Second, intentional action is responsive to one's intentional states, i.e., subjective mental states with representational content. Exercising agential power therefore requires satisfying two corresponding conditions beyond the welcome test. First, the outcome must obtain by way of one's intentional actions. Call this the *intentional action test*.

Second, the fact that the outcome is welcome must be appropriately linked to one's intention-in-action. In particular, one's intention-in-action, which renders intentional action responsive to one's intentional states, must be explained by one's favoring the outcome (i.e., by the favoring attitude satisfying the welcome test). This *linkage test* connects the welcome test to the intentional action test and hence to the exercise of agency. The link is tightest, of course, when one intends the outcome, i.e., when the intentional object of one's intention-in-action is the same as the object of the favorable attitude that satisfies the welcome test and explains one's intention. However, to restrict agential power to intended effects, as many propose, would be to construe exercising agency too narrowly.[17] That the agent intends the outcome is *sufficient* for satisfying the linkage test, but it is not necessary, for two reasons.

15 Mill, *The Subjection of Women*.

16 Searle, *Making the Social World*, 33.

17 Russell, *Power*, 23; Ball, "Power, Causation, and Explanation"; Debnam, "Nondecisions and Power"; Wrong, *Power*; Searle, *Making the Social World*; and Forst, *Normativity and Power*.

First, agents sometimes produce unintended outcomes, as by-products of their actions, which they nevertheless favor in ways sufficiently linked to their exercise of agency to count as their having effected them. Consider the fanciful case of Bumbling Master, inspired by Plautus's comedies.[18] The domestic Lar has set things up so that whenever Master intentionally acts to treat Slave badly, he unintentionally treats her well (and vice versa). So far, so bumbling: whatever Master does, he never satisfies the welcome test and so never exercises his agential power. But now imagine that Master has caught on to Lar's setup and so begins deciding strategically: when he prefers to treat Slave badly, he decides to treat her well—and thus adopts and acts on the intention to treat her well (and vice versa). The welcome test is now satisfied, insofar as he effects the outcomes he prefers. The intentional action test is also satisfied: he effects his preferred outcomes by way of his actions. The outcomes, however, are *unintended*: he intends to treat Slave badly but ends up treating her well. Nevertheless, his intention-in-action is sufficiently linked to his preference: the former is directly explained by the latter.

Consider now a less fanciful case.

> *Predatory Movie Mogul*: A movie mogul holds tremendous power over the careers of women hoping to star in his films. He uses this power to coerce them into sex. His predatory actions have numerous unintended by-products. First, women in the industry, aware of his willingness to abuse his power, come to fear him, resulting in a culture of deference to his artistic judgments on set. Second, after years of apparent impunity, the predator is arrested, convicted, and imprisoned. Although both by-products were unintended, when acting, he had a preference for a culture of subservience, but not for incarceration.

The mogul clearly effected the rapes in virtue of exercising his agential power: the outcomes were intentionally caused by him. But although he caused his own imprisonment, he did not effect it: the outcome was neither intended nor favored by him in acting. By contrast, it seems the culture of subservience was effected by his exercise of agential power because, although unintended, in acting, he favored the particular outcome in question—he preferred it—in a way sufficiently linked to his actual intention-in-action. How so? One way in which his preference (for a culture of subservience) might have been linked to his intention (to subordinate sexually) would have been if that preference had *directly* explained his intention. This would establish a sufficient link, to be sure, but the typical way in which a preference explains an intention is by bringing

18 I owe this example to Niko Kolodny.

about an intention with the same object as one's preference—which simply means the outcome was intended. (So, for example, he might have intended to coerce sex partly in order to foster a culture of subservience, in which case he also intended to foster that culture as well.) But in the mogul's case, his preference (for a culture of subservience) is sufficiently linked to his intention (to sexually exploit) in a weaker way: the former is an instance of a more general attitude—for example, a preference for subordinating women—which in turn helps explain his actual intention-in-action. Thus, the mogul exercises his power in effecting the culture of subservience: while his intention (to sexually exploit) is not caused by his preference for this culture, this preference is an instantiation of a more general preference (for subordinating women) that does cause the intention. Moreover, a hypothetical intention to effect a culture of subservience is consistent with his actual intention. His favoring attitude explains his intention only in this weaker, more extended sense.

Second, people often exercise agency via actions whose intentional objects are not the particular outcomes extrinsic to the actions themselves. (The action event itself is the intrinsic consequence of acting.) Consider a prime minister facing uncertain circumstances.[19] She does not reliably conjecture any of the potential extrinsic consequences of her available courses of action and consequently, in acting, does not intend any particular extrinsic consequence. But this does not imply she cannot exercise her agency or effect outcomes. If the linkage test were reduced to intended effects, then it would follow, absurdly, that the prime minister has no agential power in virtue of her high office to effect outcomes extrinsic to her action.

If the prime minister thought that despite the unpredictability of her actions' consequences, she could nevertheless increase the likelihood of satisfying her general objectives and plans, her actions may have been caused and constituted by a *general* intention to fulfill those objectives; if so, then the intended effects test might suffice to explain her agential power and its efficacious exercise, without resorting to a weaker linkage condition. The "general intention" retort, however, is often unavailable. Perhaps she was acting out of habit, unreflectively applying a rule of thumb in uncertain circumstances, or out of a sense of duty to formal procedural considerations independently of outcomes.[20] In that case, even if the extrinsic, downstream consequences of her action fulfilled her objectives or plans, they would not have fulfilled any of her intentions when acting—even a general intention. Yet it seems that even here, she might have efficaciously exercised power.

19 White, "Power and Intention," 751–52.

20 White, "Power and Intention," 756.

We can gain analytical clarity by assuming that from among the menu of possible extrinsic outcomes, the prime minister is in fact indifferent between them. How could she have exercised power in effecting the particular set of outcomes that ensue despite not favoring them? Imagine that in acting, the prime minister favors the circumstance that whichever extrinsic outcomes ensue do so *by way of her agency*—that is, she favors acting as she does and favors the efficacy of her agential power in so acting, i.e., she favors her intentional actions counting as an instance of her exercising power to effect those outcomes. Call this favoring her own *power efficacy*. The upshot is that although she may not favor the particular extrinsic consequences of her actions in and of themselves—whatever they may be—she does favor them indirectly insofar as they are effected by way of her exercise of intentional agency—that is, she favors the set of *overall* comprehensive outcomes: the combination of her actions' extrinsic consequences, her action itself (the intrinsic consequence of her action!), and her power efficacy in effecting those particular outcomes. Then her favoring attitude towards the overall set of consequences, including the extrinsic consequences, is appropriately linked to her exercise of agency because it explains her actual intention-in-action. We can therefore explain why, when the prime minister acts using the powers of her office, she *ex post* successfully effects the extrinsic outcomes—despite not intending or even favoring them in particular.

This analysis can similarly explain the following case.

> *Devout Tyrant*: A fanatically devout tyrant rules over his kingdom with the sole objective of carrying out the directives of his religious adviser. The tyrant has no other objectives or cares and so is completely indifferent to his subjects' plight. He consequently is not motivated by and, in acting as ruler, does not hold in view the effects of his actions on his subjects. Yet his subjects' well-being wholly turns on his actions.

Despite the facts that none of his actions' extrinsic outcomes are intended by him—they are all by-products!—and that he does not favor any of them in particular, he nevertheless exercises immense power in effecting them. The outcomes are appropriately linked to his agency because two conditions are met: first, given that he favors his own actions and their power efficacy, he favors the *overall* set of consequences of his intentional actions (which include his actions' intrinsic consequence—namely, the actions themselves and his power efficacy in effecting the extrinsic consequences); and, second, favoring his own power efficacy is appropriately linked to his actual intention-in-action because it helps explain it. Like the prime minister, the extrinsic consequences are part of a set he favors, and he favors the power efficacy of his actions. The same

could be said of a high-ranking civil servant who cares and intends to ensure only that she follow formally correct procedure in effecting outcomes and who consequently undertakes numerous actions that significantly affect many lives.

Indeed, even if there were some extrinsic consequences the prime minister, tyrant, or civil servant *disfavors*—but not enough to outweigh the extent to which they favor acting efficaciously—they would nevertheless effect even the regrettable by-products. Members of the upper classes whose actions foreseeably help perpetuate lower-class misery exercise power in effecting those outcomes, even if those outcomes are unintended by-products that they would prefer to avoid, given that their reticence is outweighed, all things considered, by the considerations prompting them to continue as they do. Or consider a lieutenant who unintentionally causes subordinates to adopt his outlook but deems such an outcome regrettable because it deprives him of advice from diverse viewpoints.[21] He might nevertheless be exercising his power insofar as this regrettable outcome is part of an overall package caused by his actions and towards which he is favorable, and he favors his actions being a mode of exercising power over his subordinates. Or consider ruthless industrial "disrupters" whose aim is just to muck around and see what happens: even if the extrinsic results are not in themselves welcome to them in particular, those results may nevertheless satisfy both the welcome and linkage tests insofar as they result from the exercise of their agential power.

To sum up, according to the linkage test for agential power, an agent's favoring the outcome (which favoring satisfies the welcome test) must be appropriately linked to the intention-in-action that constitutes the action (by way of which the outcome is effected, satisfying the intentional action test). This linkage test can be met in two ways. First, if the agent favors the *particular* outcome in question, then the test is met if her favoring attitude either itself directly explains her intention-in-action or is an instance of a more general favoring attitude that explains her intention (so that her favoring attitude explains the intention in an extended sense). The case in which the agent straightforwardly intends the outcome is only one way of instantiating this. Second, if the agent favors an *overall* set of outcomes of which the particular outcome is a part and favors her own power efficacy in effecting this set of outcomes, and her favoring attitude explains her intention-in-action, then the linkage test is met for the set of outcomes—including the particular outcome in question, even if, on its own, the agent disfavors it.

21 Partridge, "Some Notes on the Concept of Power," 114.

2. AGENTIAL CAUSAL POWER

A paradigmatic case of having agential power and efficaciously exercising it is provided by the following case.

> *Unique Assassin*: A victim has taken refuge in a location inaccessible to anyone except one assassin. If not for the unique assassin, the victim would live until old age. But the assassin accesses the location and shoots, intentionally killing the victim.

Ex ante, the unique assassin *has* the agential power to kill the victim (and is the only one with the power to do so herself). *Ex post*, the assassin has efficaciously *exercised* her power to effect the victim's killing. I wish to provide a preliminary analysis of these two phenomena.

What explains the assassin's power to kill the victim? Although she has efficaciously *exercised* her power to effect the killing, *ex ante* she *has* the power to do so regardless of whether she chooses actually to exercise it. The assassin's possession of power could be explained by two sets of facts. First, if she were to favor the victim being killed, then she would consequently act, and the victim would be killed. Second, if she were to act on such a favoring attitude, the victim would die, whereas if she were not to so act, the victim would live. Hence, *ex ante*, a set of actions available to her is, within the given social-structural context, causally both necessary and sufficient for the outcome; and *ex post*, she has efficaciously exercised power insofar as her actions have caused the death. Insofar as the assassin favors the outcome, the welcome test is met; insofar as her action is necessary and sufficient for and hence causes the outcome, the intentional action test is met; and insofar as she intends the outcome, the linkage test is satisfied.

Yet such an explanation cannot cover all cases; in particular, it does not cover cases of causal overdetermination. Neither the necessity nor the sufficiency of an agent's actions is required. Consider the following.

> *Three Small Assassins*: A victim is tied up in a car's trunk. There are three assassins, none of whom is strong enough to push the car over the cliff by himself, but any two together are sufficient. All three push, intentionally killing the victim.

No single assassin's action is either necessary or sufficient for the killing: if he were to have not pushed, the victim would still have been killed by the other two; if he had pushed on his own, the victim would not have been killed at all. Yet each plays a causal role in the killing. How is this exercise of agential causal power to be explained?

There exist several related approaches to explaining the causal role of individual conditions in cases of overdetermination.[22] For our purposes, we do not need to decide between them; we need simply to note that causation can come in degrees and that in causally overdetermined cases, some causal conditions are merely partial causes. A simple approach for explaining this (for cases without preemption) is to analyze a condition's causal role in terms of a NESS test: a condition plays a causal role in case it (is not preempted and) is a *necessary element of a sufficient set* of conditions for the outcome, which set is a subset of the actual set of conditions on that occasion.[23] Here, the actual set of conditions comprises the pushing by the first, second, and third assassins. This set of three action events has three proper subsets sufficient for the outcome: the actions of the first and second, of the first and third, and of the second and third assassins. If the first assassin's action were sufficient for the outcome, it would be the *sole* necessary element of at least one sufficient subset of the actual set of conditions. And if it were necessary for the outcome, it would be an element of *all* sufficient subsets; as such, it would be fully causally efficacious. The fact that his action is insufficient does not imply he has no power; it implies he does not have strictly *unilateral* power to kill the victim—whatever power he has is a "power-with" others.[24] And the fact that his action is unnecessary does not imply he played no causal role: it implies his action was not decisive or a "full" cause. But since his action is a necessary element of two and only two of the three sufficient proper subsets, it is a *partially* efficacious cause.[25] *Ex ante*, the first assassin has some degree of power-with to effect the victim's killing.

Recognizing partial efficacy exposes the inadequacy of the widespread view that agents efficaciously exercise power only if they are *decisive*, i.e., only if, but for their action, the outcome would not occur. This view is presupposed by all who follow Robert Dahl in claiming that an agent exercises power over another only if she causes him to do something he "would not otherwise do" but for that exercise.[26] This view ignores the power one may exercise *with* others to effect an outcome—including power exercised *over* someone, causing them to alter

22 McDermott, "Redundant Causation"; Ramachandran, "A Counterfactual Analysis of Causation"; Hitchcock, "The Intransitivity of Causation Revealed in Equations and Graphs"; Schaffer, "Overdetermining Causes"; Halpern and Pearl, "Causes and Explanations: Part I" and "Causes and Explanations: Part II"; Braham, "Social Power and Social Causation"; Braham and van Hees, "Degrees of Causation"; and Pearl, *Causality*.

23 Wright, "Causation in Tort Law" and "Causation, Responsibility, Risk."

24 Allen, "Rethinking Power"; and Abizadeh, "The Grammar of Social Power."

25 Braham and van Hees, "Degrees of Causation."

26 Dahl, "The Concept of Power," 202–3. See Forst, *Normativity and Power*, 40.

their behavior—even when, as in overdetermined cases, one could not have unilaterally scuttled the outcome.[27]

We now have a preliminary analysis of the most intuitively straightforward type of power: agential causal power. *Power*, in that the outcome is (or depends on being) favored by the agent. *Agential*, in that the outcome obtains (or would obtain) by way of the agent's *intentional actions*, where the agent's intention-in-action is appropriately *linked* to, because explained by, her favoring attitude. And *causal*, in that the agent effects (or would effect) the outcome via actions helping to cause it.

3. AGENTIAL NONCAUSAL POWER

The cases considered so far are cases of agential *causal* power. But casual efficacy—whether full or partial—is not only insufficient; it is also not necessary for effecting outcomes. It is not necessary because one can help ensure or see to an outcome without actually causing it.[28] Consider the following cases.

> *Preempted Poisoner*: Assassin *A* injects a victim with a fatal poison. But before it takes effect, Assassin *B* shoots the victim, intentionally killing him instantly.
>
> *Preempted Shooter*: A victim is sleeping. Many people want him dead before sunrise, but no one can tell whether he is sleeping or dead. To ensure he is dead, Assassin *A* will shoot the victim at 3 AM no matter what; her shot will be sufficient for the kill. But at 2:30 AM, Assassin *B* shoots the victim, intentionally killing him. At 3 AM, Assassin *A* shoots the victim, who, unbeknownst to Assassin *A*, is already dead.

In neither case is Assassin *A*'s action necessary: even if she were not to poison or shoot, the victim would be killed. Nor indeed is Assassin *B*'s action necessary in either case. However, Assassin *A*'s action, like Assassin *B*'s, *does* pass the NESS test: because her action is sufficient for the killing, it is the sole necessary element of a sufficient subset of the actual set of conditions. Yet in neither case does Assassin *A*'s action *cause* the outcome: it operates too late to do so. These examples—of what David Lewis calls *late preemption*—show that the NESS test is insufficient for demonstrating causal efficacy because it is satisfied by not just preempting causes but also preempted potential causes.[29] A *complete causal test* must therefore add a further set of conditions ruling out preempted potential

27 Abizadeh, "The Power of Numbers."

28 Morriss, *Power*, 30–31.

29 Lewis, *Philosophical Papers*, 200–7.

causes. I do not defend a particular approach to fleshing out these further conditions; readers should supply whichever approach they deem most successful.[30]

Ex ante, the preempted assassin has the power to see to it that the victim is killed; insofar as she exercises this power, she ensures the victim is killed. That the preempted assassin's action ensures the outcome is precisely why the preempting assassin's action is not necessary either on this occasion. The preempted assassin's action *effects* the victim's killing insofar as it satisfies the NESS test. But insofar as it fails the complete causal test (which rules out preempted potential causes), her action effects the killing without *causing* it.

A similar analysis can be provided of some types of invigilation.[31] Consider the following case.

> *Invigilating Rainmaker*: A farmer has a rainmaking machine. The machine cannot prevent rain, but on otherwise rainless days, it can be used to make it rain. On days in which it would rain naturally anyway, the natural causal process leading to rain renders the machine causally inert even if used. The farmer needs rain today and intends to use the machine if necessary. But it rains naturally today, so she leaves the machine idle.[32]

Ex ante, the invigilating rainmaker has the power to ensure or see to it that it rains today. By acting on the conditional intention to use the rainmaking machine if necessary and refrain from using it if not, she exercises this power when she acts with that intention. It is true that the rainmaking machine gives her the power to cause it to rain in this general context. But today, on this particular occasion, given that nature causes it to rain, she does not have the power to cause it to rain. (If she were to try, she would fail, preempted by nature.) Therefore, *ex post*, the rainmaker effects the outcome on this occasion without causing it.

It might be objected that the invigilator cannot be said to effect the outcome *ex post* at all—on the grounds that, unlike the preempted poisoner and shooter above, she does not undertake any action. But invigilation is relevantly analogous to cases of late preemption. The difference is that the invigilator undertakes an intentional action whose intention has a *conditional* form. The point is perhaps clearer in the following case.

> *Invigilating Assassin*: Two rival assassins want the same victim dead. Assassin *A* knows that Assassin *B* intends to kill the victim by 3 AM but

30 See Lewis, "Causation" and *Philosophical Papers*; McDermott, "Redundant Causation"; Ramachandran, "A Counterfactual Analysis of Causation"; Hitchcock, "The Intransitivity of Causation Revealed in Equations and Graphs"; and Pearl, *Causality*.

31 Pettit, "Freedom and Probability" and "Republican Freedom."

32 See Goldman, "Toward a Social Theory of Power"; and discussion in Morriss, *Power*.

> has doubts about his rival's efficacy so he secretly invigilates the killing, adopting the following conditional intention: if his rival succeeds by 3 AM, he will hold his fire, but if not, he will shoot the victim himself. Assassin *B* kills the victim before 3 AM. The invigilating assassin carries out his conditional intention by holding fire.

Forbearing from shooting is an intentional action that instantiates the invigilating assassin's conditional intention. The assassin has the power to ensure and hence *effect* the victim's killing and does ensure it by acting conditionally as he does, but given his rival's actions, he does not *cause* the outcome *ex post* on this occasion.

The similarity with late preemption stems from the fact that these invigilation cases are also ones of preempted causation: cases in which a potential causal process is preempted by another causal process. The difference is that in invigilation cases, the potential causal process is preempted not by preventing the preempted potential cause from causing the effect but by preempting the potential cause itself; they therefore instantiate not late but so-called *early* preemption.[33]

It might be objected that neither late nor early preempted agents can ensure that *they themselves* kill their victim or make it rain and therefore cannot be said efficaciously to effect the relevant outcome. The premise is correct, but the conclusion does not follow. If the agent welcomes events only if they are caused by her, then the relevant outcome is an event-caused-by-her. True, the preempted shooter or invigilating rainmaker does not effect that outcome: she cannot ensure that she herself kills the victim or makes it rain. But if, as stipulated, her intention-in-action is not agent relative, she does effect the favored outcome: she ensures death or rain.

4. ELICITORY CAUSAL POWER

It is of course possible to cause outcomes without undertaking any intentional actions at all. Consider the following case.

> *Mafia Godfather*: A mafia godfather's henchmen understand his overall objectives and therefore accurately anticipate his particular wishes. With nary a word from him, they anticipate and carry out what they believe to be his wishes. Given his objectives, were the mafioso cognizant of the threat posed by a police officer building an actionable case against him, he would want the officer assassinated. His henchmen, without so much

33 Lewis, *Philosophical Papers*, 200–7.

> as apprising him of the situation, assassinate the officer. Whatever the godfather were to want, his henchmen would seek to realize.

The mafioso's objectives are fulfilled, but not by way of his intentional actions. If exercising power implies doing so by way of intentional actions, then, because the agent-outcome relation fails the intentional action test, the mafioso cannot be said to be exercising agential power here. Yet I take it that, recognizably, the mafioso's objectives are nevertheless fulfilled in virtue and by way of his power, and not as a mere side effect unrelated to his preferences or objectives: the outcomes are not only welcome to him (satisfying the *ex post* welcome test); they are caused by and counterfactually vary with his preferences (satisfying the *ex ante* welcome test). Indeed, the mafioso may very well prefer to wield power passively in this way—not only to save himself the effort but also because it maintains plausible deniability!

What explains the mafioso's power to obtain his preferences without acting? Let us begin by filling out some implicit background details. There is, firstly, a hierarchical social relation between the godfather and each henchman: they all recognize him as boss, with the power to issue orders to them, etc. This relation, secondly, is also structural: a set of "regulators"—not just he and any given particular henchman but also the other henchmen, others in the organization, indeed other crime groups and the police—also recognize his position over the henchman, and their recognition and treatment of the former as the boss, and the common knowledge that others will do so as well, stabilize and reinforce the relation. So the mafioso clearly has *agential* power in virtue of his structural position: he can order his henchmen to do things. And insofar as his power to obtain outcomes—for example by ordering them around—is due to the position he occupies over his henchmen, his power is "structural" in the broad sense I set aside.

Our question is how his position translates into elicitory power. We can consider several variants of the Mafia Godfather case, each with its own explanation. On the first, which we can call the Vengeful Mafioso variant, what explains his elicitory power is his henchmen's anticipation of his reaction should they not fulfill his objectives: if they were to not assassinate the officer, the mafioso may kill them for incompetence or disloyalty. This is what Carl Friedrich calls the "rule of anticipated reactions," referring to the reactions of the *power holder* himself.[34] The Friedrichian interpretation may lead some to question whether the Mafia Godfather case supports the notion of elicitory power. If what instills fearful anticipation and loyalty is the mafioso's current and known disposition to exact vengeance, then it might be objected that he

34 Friedrich, *Constitutional Government and Politics*, 16–18, and *Man and Government*, 199–215.

does exercise agential power via his intentional actions after all. There are three versions of this objection.

First, the reason why the vengeful mafioso's henchmen anticipate his wishes and vengeance might after all be due to the mafioso's *past* intentional actions and ensuing reputation.[35] However, to have intentionally acted in the past to create a reputation in virtue of which one now passively elicits outcomes is not *now* actively to exercise agential power. It is to have exercised agential power in the past to effect a structure in which one now passively elicits outcomes. Others' knowledge of one's disposition for unwanted reactions, moreover, is not always caused by one's past actions. Sometimes it is known because one is a token of a certain type, for example, by occupying a certain social-structural position. Consider the following case.

> *Capitalist Giant*: A capitalist owns and controls a giant corporation. His aim is to maximize profits; to fulfill this aim, he must invest in the jurisdiction with the lowest corporate taxes. If the corporation's current jurisdiction raises corporate taxes, he will move the corporation elsewhere, with catastrophic economic consequences for his current jurisdiction. The capitalist would not intend to inflict these harms; he would merely intend to maximize profits. The incentive structure of capitalist corporations is well known to the government; to avoid catastrophe, it invariably sets corporate taxes at a low rate.[36]

The capitalist giant elicits outcomes in virtue of the government's anticipation of his hypothetical future actions should it act contrary to his preferences. However, the government knows his preferences and consequently his disposition, not because of his past actions but because of the capitalist's type and position in the economic structure.

A second version of the objection is that the mafioso's disposition to avenge failures reflects a conditional standing intention to do so, in which case he *does* undertake a negative intentional action: he is an invigilator. I concede that *if* so, then the case of anticipated reactions would be one in which the mafioso exercises agential power. But although some mafioso might have such a standing intention, another might not: he might have a disposition without a standing intention, for example, if had not yet decided or adequately considered what to do if his henchmen betrayed his trust. The vengeful mafioso I have in mind is of the latter type.

35 Dowding, "Resources, Power and Systematic Luck," 316.

36 See Lindblom, "The Market as Prison"; and Barry, "Capitalists Rule Ok?"

Third, some might worry that the vengeful mafioso is poor evidence for elicitory power because whatever power he has seems parasitic on his *agential* power to exact revenge. Although he does not fulfill his preferences by way of his *actual* intentional actions, he does seem to realize them by way of *hypothetical* intentional actions—which he would undertake if he were betrayed. And it might be argued that the category of agential power should be expanded to include wielding power by way of hypothetical intentional actions of this kind, without invoking a separate category of elicitory power.[37] The problem with this suggestion is twofold. First, the vengeful mafioso would wield the same power even if he did not have or was not disposed to exercise the agential power to exact revenge, just as long his henchmen believed that he has and is disposed to exercise such power. Second, the objection fails to account for the other variants of the Mafia Godfather case, to which we now turn.

Sometimes agents have elicitory power in virtue of the anticipated reactions of other agents, not their own reactions. Consider the case of the Figurehead Mafioso. The figurehead has no capacity to exact vengeance on his henchmen, and they know this. Their loyalty is explained rather by the larger social context of rivalry between mafia gangs: if the figurehead mafioso were to fall, his henchmen would fall with him. What motivates them is the anticipated actions of third-party "peripheral actors" or "regulators," not the godfather's. The henchmen act on the basis of what they take to be his preferences, not necessarily because they deem his judgment about such matters superior to their own but because they use him as a focal-point coordinating mechanism: their gang's standing requires an efficient way to coordinate their actions without acting at cross-purposes.[38] This mafioso elicits outcomes in virtue of his social-structural position in the network of mafia hierarchies. He does not have the agential power to effect the outcomes nor even effectively to retaliate against failures, but nevertheless elicits those outcomes by his structural power.

Some cases of elicitory power do not depend on the anticipation of *anyone's* future actions. Consider the case of the Old Mafioso, whose gang has successfully wiped out all rivals and is at little risk of failing. He is so old and decrepit now that he cannot leave his bed and no longer has the power to extract deadly vengeance against his henchmen. But his henchmen are so used to anticipating and carrying out his wishes—except extracting vengeance against fellow henchmen—they can no longer imagine doing otherwise: they cannot fathom their life's purpose or meaning without serving their godfather. Serving has become second nature. The old mafioso no longer has and hence could not

37 I thank David Estlund for this suggestion. See Barry, "Capitalists Rule Ok?" 178–80.

38 Schelling, *The Strategy of Conflict*.

even hypothetically exercise the agential power to avenge the outcomes, but he still has the passive power to elicit them: outcomes vary with his preferences because his preferences help cause them.

Similarly, consider the following case.

> *Charismatic Prophet*: A prophet has such a charismatic presence that many fall under her spell merely from being in her presence. She so inspires her admirers that they anticipate and seek to fulfill her every wish even before she herself has become conscious of them, not because they seek future reward or approval but because they take joy in serving her.

The charismatic prophet elicits outcomes thanks to her elicitory causal power: some features of her besides her intentional actions—such as her physiognomy, odor, and unconscious demeanor—cause others to anticipate her preferences and fulfill them. In a social context, such power also at least partly exists in virtue of the agent's position within a social structure—for example, the structure of norms and mental schemas according to which a woman's physiognomy, odor, and demeanor tend to be interpreted. But her preferences' causal role in eliciting outcomes is not due to anticipated reactions.

Here is another example.[39] Imagine we live in a society in which, thanks to the racial structure, people are very attentive and sympathetic to how light-skinned people wish to be treated: most are disposed to detect and satisfy light-skinned people's preferences for how to be treated. You already have light skin, so this is quite pleasant for you. But I have dark skin and, predictably, have not been treated as I had wished. However, my skin is not too dark, I have money, and I can and do buy effective skin-lightening creams. My prior intention and my intention-in-action in buying and applying these creams are to get people to treat me as I wish, which is exactly what happens. So I have efficaciously exercised my agential power to effect outcomes satisfying my preferences for how others treat me. Yet what I have done is merely, by way of my intentional actions, to try to mimic a power you already have and wield but not by way of any intentional actions. You fulfill the same type of preferences for yourself without having to do anything. I have to expend energy exercising agential power just to mimic your elicitory, structural power.

In sum, when agents' preferences help cause outcomes independently of their intentional actions, they elicit those outcomes. They wield elicitory power.

39 I thank Lucas Stanczyk for suggesting an example like this.

5. ELICITORY NONCAUSAL POWER

The most controversial type of power I defend here is perhaps elicitory noncausal power. But once we acknowledge that one can effect outcomes by exercising agential *noncausal* power, on the one hand, and elicit outcomes via *elicitory* causal power, on the other, then there is no reason to deny the combination of elicitory and noncausal power. Not only can social power take this form; it is one of the most important kinds of social power.

How can agents passively wield power to elicit outcomes by way of neither intentional actions nor any causal role? One obvious way (analogous to noncausal agential power) is if their preferences help to ensure the outcome but are preempted. Consider the case of the Preempted Mafioso, whose preferences are not even the outcome's *actual* cause. Imagine the henchmen sometimes cannot discern his preferences as easily and swiftly as they can his son's. Since the son shares his father's objectives and invariably wants the same, the henchmen treat (their sense of) what the son wants as a kind of oracle for his father's wishes. Now imagine they think the decision whether to assassinate is urgent but that trying to discern the father's preference risks too much delay; so they act instead on the basis of discerning the son's preference. Here it is the son's preference, not the father's, that causes the outcome, but had the son not preferred it, the henchmen would have later discerned the father's preference and carried out the assassination anyway. The son's preference preempted the father's potential causal role, but the preempted mafioso nevertheless has (passively) elicited the outcome, in a manner similar to how invigilators (actively) effect outcomes: his preference passively invigilated the outcome, so to speak, by ensuring it would happen.

Sometimes, moreover, agents have the power to fulfill their preferences thanks to the causal role of others' actions or preferences rather than their own, even without their own preferences being preempted causes. Consider the following case.

> *Immobile Little Capitalist*: A capitalist owns and controls a small business. Because his aim is to maximize profits, his incentive structure normally propels him to invest in the jurisdiction with the lowest corporate tax rates. But due to family responsibilities, he is now immobile and will stay in his current jurisdiction whatever the tax rate. There are, however, hundreds of other little capitalists in his jurisdiction who are mobile: if the government raises corporate taxes, they will all move their businesses elsewhere, with catastrophic economic consequences for the current jurisdiction. The incentive structure of capitalists is well known to the government; to avoid catastrophe, it sets corporate taxes at a low rate.

The outcome is welcome to the little capitalist, but his preferences (and anticipated reactions) play no causal role in effecting it: the outcome is caused by the preferences and anticipated reactions of other, mobile little capitalists. Nothing about the immobile little capitalist—not even the position he occupies in the social structure as a capitalist business owner—helps cause the outcome. But in virtue of the position he occupies, the immobile little capitalist obtains the tax outcomes he wishes. He has the power to fulfill his preferences thanks to his position and what *others* systematically prefer and would do, and thanks to the systematic correlation of their preferences with his—but independently of his intentional actions or any causal role (or even preempted causal role) for his preferences.

Why should this count as having and wielding structural power rather than merely as good luck and benefitting? Because the relevant outcomes systematically vary with the agent's preferences. Imagine the economic situation changes, such that the profits of little capitalists come to depend on a highly educated, highly skilled workforce. And imagine that producing such a workforce requires public state investment in early education and training that can be funded only via higher corporate taxes.[40] Now the little capitalists all prefer jurisdictions with significant educational investment funded by higher taxes. To prevent the little *mobile* capitalists from leaving, the state raises corporate taxes, just as the little mobile *and* immobile capitalists have come to prefer. This is no mere benefitting. It is benefitting that counterfactually varies with the mobile *and* immobile capitalists' preferences—even though the preferences of only the former actually *cause* the varied benefits. (The preferences of the latter, in turn, vary with those of the former because of their similar position in the economic structure.) The difference between merely benefitting versus wielding elicitory, structural power turns on this key point.

Therefore, not all beneficial outcomes may be imputed to an agent's elicitory power. Consider a neighbor who just happens to splash water on the plants in your parched garden. That the outcome is beneficial and welcome to you is not sufficient for showing you have elicited it in virtue of your power. We cannot discern elicitory power just by observing that the outcome is *ex post* welcome on a particular occasion: what matters is whether such outcomes would vary with hypothetical variations in your preferences. For there to be such variation in principle—for the splashing to count as elicited by your power—then at least one of two conditions must hold. Either your preferences were a full, partial, or preempted cause of the splashing; for example, your neighbor knew you wanted your plants watered and was therefore moved to splash water because,

40 I thank Jacob Levy for suggesting this example.

due to your standing in the community, she wants to curry your favor. Or, if your preferences did not cause or help ensure the splashing, then, at least on some occasions in which the splashing might have happened, if you had counterfactually preferred that it not happen, then it would not. Otherwise, you have merely benefitted. The welcome watering is elicited by your structural power only if your neighbor's actions somehow counterfactually track your preferences thanks to your position in a social structure (such as the property regime).

More generally, whereas agential power requires satisfying the welcome, intentional action, and linkage tests, elicitory power requires satisfying only the welcome test. But to have efficaciously wielded elicitory power *ex post*, it is necessary *but not sufficient* to satisfy the *ex post* welcome test. The possession and hence efficacious wielding of elicitory power requires also satisfying the *ex ante* welcome test: not only must the actual outcome be welcome *ex post* on this particular occasion; outcomes must also systematically vary counterfactually with the agent's preferences in this context in general. This is what explains the difference between efficaciously wielding structural power and mere benefitting.

It also explains the difference from *systematic* mere benefitting. Consider an incumbent infrastructural regime such as an existing clean water system whose continuing functioning requires no maintenance during my lifespan.[41] I "systematically" benefit from it in the sense that my actual preferences for clean water are continually and habitually fulfilled. But I do not thereby have elicitory or structural power to fulfill my preference for clean water: the provision of the benefit does not counterfactually vary with my preferences. Were I counterfactually to prefer clean kombucha rather than water, the infrastructure would still furnish water.

6. CONCLUSION

Why is it important to recognize nondecisive, noncausal, and elicitory and structural categories of power? Peter Morriss argues that the concept of power serves three purposes: the practical purpose of determining how to fulfill our aims; the moral purpose of allocating responsibility and blame; and the evaluative purpose of appraising social arrangements.[42] Each of these contexts illuminates the significance of recognizing the types of power defended here.

Begin with the practical purpose of trying to identify the agents over whom it would be useful to wield power. If someone has *agential* power over issues

41 I thank Matthew Noah Smith for the example.

42 Morriss, *Power*, 37–42.

that concern you, then of course it makes perfect instrumental sense to wield power over them—to get them to do something they otherwise they might not do, to reprise Dahl's classic formulation. But it *also* makes good instrumental sense to wield power over agents with elicitory, structural power: shaping their preferences would be another way to obtain or ensure your preferred outcomes—albeit without getting them to *do* anything.

Next consider moral theory. It is a platitude that power and responsibility go hand in hand. It is also widely held that to be normatively responsible and appropriately blamed, two conditions must be met: first, one must be normatively *competent*; and second, that for which one is responsible must be under one's *control*. A plausible account of normative competence equates it with *rational agency*: the capacity to grasp concepts and recognize facts as reasons; to reflectively assess what reasons one has; and to respond, in one's intentional states and actions, to those reflectively assessed reasons.[43]

The widespread restriction of power to agential power has often been motivated by a particular gloss on the second, control condition: that one can bear normative responsibility and appropriately be blamed only for what is under one's *voluntary* control. On this view, one can be (derivatively) responsible only for the consequences of one's choices and intentional actions.[44] Thus, if power is linked to responsibility and serves to identify those responsible, then it seems the notion of power at stake must be agential—operating by way of intentional actions.

This "voluntarist" theory of responsibility and blame, however, is highly controversial at best; indeed, I believe it to be mistaken, but I limit myself to showing why we should not burden a theory of power with it.[45] The problem is that it sits poorly with our actual responsibility and blaming practices. There is nothing more pedestrian than resenting or feeling indignation towards those who take pleasure in or desire others' pain, bear evil hopes or desires for others, fail to take others' well-being into account in deliberation, or believe that members of our social group do not merit respect or are less capable than we are.[46] And there is no denying that people resent friends for regularly forgetting a special occasion, failing to notice ethically salient features of the current circumstances, or involuntarily betraying indifference or contempt. Yet unlike

43 Wallace, *Responsibility and the Moral Sentiments*, 157–58; and Skorupski, *The Domain of Reasons*, 21–23, 59.

44 Wolf, *Freedom Within Reason*; and Wallace, *Responsibility and the Moral Sentiments*.

45 I leave aside incompatibilist accounts equating control with metaphysical *free will*, which, combined with the supposition of causal determinism, would evacuate any talk of normative responsibility and blame. See Fischer and Ravizza, *Responsibility and Control*.

46 Adams, "Involuntary Sins."

bodily movements, we typically cannot directly will or choose to believe, feel, or desire: many thoughts, emotions, and desires are spontaneous, involuntary responses to the world. In short, we often blame people and hold them responsible for involuntary states (and behaviors). True, sometimes our reactive blaming responses target agents' past, character-shaping choices, but often we resent the intentional state itself, not its genetic history.[47]

A plausible explanation for these practices is that we appropriately hold agents responsible for intentional states insofar as these are rationally responsive to agents' own reflective evaluations or judgments. On this construal of the control condition—which fits much more neatly with the first, rational agency condition—we can be blameworthy even for involuntary states or behaviors insofar (and only insofar) as they are judgment sensitive.[48] What "controls" or guides our states is not our will but our evaluative judgments or the mechanisms constituting our rational agency.[49]

The upshot is that there is no responsibility-related reason for restricting the power concept to agential power. Indeed, there is every reason for it to encompass elicitory power: the preferences on the basis of which the welcome test analyzes elicitory power are judgment-sensitive evaluative states.[50] In other words, one normative payoff of the theory defended here is that it opens the door for power structure analyses that more closely track actual practices and better align with more plausible theories of responsibility and blame.

Another payoff is to provide a plausible framework for thinking about how power may be related to the different forms of holding people responsible. Here I can merely sketch out this possibility. Consider the blaming emotions that, if we follow the Strawsonian tradition, are constitutive of blame and holding people morally responsible. The content of anger, disdain, or shame, on the one hand, and resentment, indignation, or guilt, on the other, is that its target has violated an expectation, standard, or requirement. But unlike the first trio, the second trio specifically concerns requirements *owed* to others; as such, they are "reactive" or vindicatory, second-personal emotions in the sense that they inherently

47 Smith, "Responsibility for Attitudes."

48 Smith, "Responsibility for Attitudes."

49 Fischer and Ravizza, *Responsibility and Control*.

50 Hausman, *Preference, Value, Choice, and Welfare*. True, some *sources* of preferences may be judgment insensitive. But because the preferences in question are total or *all-things-considered* evaluative rankings incorporating all relevant sources, they remain judgment sensitive. Even if my preference for obtaining vanilla over chocolate sorbet were currently based solely on brute taste, if I were now to become aware of a normative reason against obtaining vanilla, I may, all things considered, come to prefer obtaining chocolate (despite the partial preference for vanilla on taste grounds).

demand a response from their target—whether justification, excuse, apology, compensation, or measures to prevent future violations.[51] Expressly to hold someone morally accountable for a perceived blameworthy violation is, in the first instance, to demand such responses by way of expressing reactive blaming emotions (or the belief that they are appropriate) and/or to impose sanctions.

There is no reason to attribute blanket immunity against these responses to agents whose power falls short of decisive, causal agential power. Nevertheless, we can now see how, depending on the type of power, some modes of holding responsible may be inappropriate. Due to collective action problems, for example, agents with nondecisive power may be unable to implement prevention measures and hence not be liable. Their nondecisiveness may also furnish an excuse for violations, alongside standard excuses like ignorance. Agents with causal structural power may, in virtue of their causally efficacious preferences, be open to all the accountability demands and even to interpersonal or social sanction but, because choice is missing, not be open to coercively imposed political sanction. And whether those with noncausal (because preempted) power—agential or structural—are deemed immune from demands for compensation may depend on which views about circumstantial moral luck are justified. Moreover, I suspect that one whose noncausal structural power stems solely from sharing preferences with others is not the appropriate target of reactive blame at all—but may very well be the appropriate target of some of the *demands* associated with them (such as prevention measures) and perhaps even nonreactive forms of blame.[52]

Finally, from an evaluative perspective, why is it important to recognize structural power, including noncausal structural power, as a category of power—rather than as merely good luck or privilege? True, sometimes "privilege" amounts to mere systematic benefitting, not power. But sometimes privilege amounts to more, and when it does, we should recognize it as structural *power* and not merely "privilege" because structural power is a wider category. It is not inherently hierarchical, inegalitarian, or evil: it can be widely dispersed, equally shared, reciprocal, and in many cases of normative value. Consider a society governed by an effective norm that others offer their seat to any person facing physical hardship; a Good Samaritan norm that others nearby give aid, shelter, or protection to any imperiled person; or laws providing for basic health care to any person in need. All persons in such a society would have, in virtue of their position in the corresponding social structures, the structural power to have these fundamental needs met—independently of their own intentional actions.

51 Strawson, *Freedom and Resentment and Other Essays*; Wallace, *Responsibility and the Moral Sentiments*; and Darwall, *Second-Person Standpoint*.

52 Watson, "Two Faces of Responsibility."

Furthermore, many egalitarians prize political equality, frequently understood to demand equal power over binding political decisions, while neorepublicans seek to minimize relations of domination, characterized as arbitrary relations of power-over.[53] But if the theory defended here stands to reason, then political equality and nondomination may require attending to inequalities or arbitrariness in noncausal and structural forms of power as well. To miss these phenomena as instances of power is to obscure one of the central concerns of social theory and normative political philosophy: the distribution of power.[54]

McGill University
arash.abizadeh@mcgill.ca

53 For discussion of the former, see Beitz, *Political Equality*; Wilson, *Democratic Equality*; and Wodak, "What Is the Point of Political Equality?" For the latter, see Pettit, *A Theory of Freedom*; and Lovett, *A General Theory of Domination and Justice*.

54 Funding was provided by Social Sciences and Humanities Research Council (SSHRC) of Canada grant 435-2025-0832 and Fonds de recherche du Québec – Société et culture (FRQSC) grant 2023-SE3-318011. For comments on previous drafts, I am grateful to Arthur Applbaum, Federica Berdini, Yuna Blajer de la Garza, Adrian Blau, Matthew Braham, Garrett Cullity, Jorah Dannenberg, Robin Douglass, Keith Dowding, Yiftah Elazar, David Enoch, David Estlund, Mollie Gerver, Pablo Gilabert, Valéry Giroux, Carole Gould, Alex Gourevitch, Amanda Greene, Allen Habib, Alon Harel, Sally Haslanger, Sean Ingham, Dongxian Jiang, David Johnston, Steven Klein, Kyuree Kim, Yunhyae Kim, Niko Kolodny, Thomas Krödel, James LaBelle, Holly Lawford-Smith, Chad Lee-Stronach, Jeffrey Lenowitz, Jacob Levy, Éliot Litalien, Catherine Lu, Karuna Mantena, Mara Marin, Shal Marriott, Jonathan Leader Maynard, Victor Muñiz-Fraticelli, Josh Ober, Armando Jose Perez-Gea, Ryan Pevnick, Veronica Ponce, Samantha Puzzi, Will Roberts, Andrea Sangiovanni, Shlomi Segal, Joshua Simon, Matthew Noah Smith, Jules Solomone-Sehr, Nic Southwood, Lucas Stanczyk, Jordan Walters, Leif Wenar, Daniel Weinstock, Vanessa Wills, Yves Winter, Yi Yang, several anonymous referees, and participants at the McGill Research Group on Constitutional Studies (RGCS) workshop on February 18, 2021; the Stanford Political Theory Workshop on February 19, 2021; the Columbia Political Theory Workshop on March 31, 2021; the Centre de recherche en éthique (CRÉ) lunch series at Université de Montréal on September 21, 2021; the Kings College London Department of Political Economy Seminar on September 29, 2021; the Jerusalem Seminar in Moral and Political Philosophy on April 24, 2022; the RGCS/CRÉ/Groupe de recherche interuniversitaire en philosophie politique manuscript workshop on August 15–17, 2022; City University of New York (CUNY) Center for Global Ethics and Politics on October 13, 2022; the Annual Meeting of the Philosophy, Politics, and Economics Society in New Orleans on November 3–5, 2022; the Brown University Philosophy Department manuscript workshop on March 25, 2023; the Australian National University Conference in Honour of Robert Goodin on August 8–9, 2023; the University of Sydney Philosophy Seminar Series on August 14, 2023; the Hamburg University Collective Decision-Making Program Interdisciplinary Research Seminar on April 4, 2024; and the Hamburg University Collective Decision-Making Program manuscript workshop on April 8, 2024.

REFERENCES

Abizadeh, Arash. "The Grammar of Social Power: Power-to, Power-with, Power-despite and Power-over." *Political Studies* 71, no. 1 (2023): 3–19.

———. "The Power of Numbers: On Agential Power-with-Others Without Power-over-Others." *Philosophy and Public Affairs* 49, no. 3 (2021): 290–318.

Adams, Robert Merrihew. "Involuntary Sins." *Philosophical Review* 94, no. 1 (1985): 3–31.

Allen, Amy. "Rethinking Power." *Hypatia* 13, no. 1 (1998): 21–40.

Ball, Terence. "Models of Power: Past and Present." *Journal of the History of the Behavioral Sciences* 11, no. 3 (1975): 211–22.

———. "Power, Causation, and Explanation." *Polity* 8, no. 2 (1975): 189–214.

Barry, Brian. "Capitalists Rule Ok? Some Puzzles About Power." *Politics, Philosophy, and Economics* 1, no. 2 (2002): 155–84.

———. "The Uses of 'Power.'" *Government and Opposition* 23, no. 3 (1988): 340–53.

Beitz, Charles R. *Political Equality: An Essay in Democratic Theory*. Princeton University Press, 1989.

Bourdieu, Pierre. *La Domination Masculine*. Seuil, 1998.

Braham, Matthew. "Social Power and Social Causation: Towards a Formal Synthesis." In *Power, Freedom, and Voting: Essays in Honour of Manfred J. Holler*, edited by Matthew Braham and Frank Steffen. Springer, 2008.

Braham, Matthew, and Martin van Hees. "Degrees of Causation." *Erkenntnis* 71, no. 3 (2009): 323–44.

Dahl, Robert A. "The Concept of Power." *Behavioral Science* 2, no. 3 (1957): 201–15.

———. "Power." In *International Encyclopedia of the Social Sciences*, edited by David L. Sills. Macmillan, 1968.

Darwall, Stephen. *The Second-Person Standpoint: Morality, Respect, and Accountability*. Harvard University Press, 2009.

Debnam, Geoffrey. "Nondecisions and Power: The Two Faces of Bachrach and Baratz." *American Political Science Review* 69, no. 3 (1975): 889–99.

Dowding, Keith. *Rational Choice and Political Power*. Edward Elgar, 1991.

———. "Resources, Power and Systematic Luck: A Response to Barry." *Politics, Philosophy, and Economics* 2, no. 3 (2003): 305–22.

Dretske, Fred. *Explaining Behavior: Reasons in a World of Causes*. Massachusetts Institute of Technology Press, 1988.

Elder-Vass, Dave. *The Causal Power of Social Structures: Emergence, Structure and Agency*. Cambridge University Press, 2010.

Fischer, John Martin, and Mark Ravizza. *Responsibility and Control: A Theory*

of Moral Responsibility. Cambridge University Press, 1998.
Forst, Rainer. *Normativity and Power: Analyzing Social Orders of Justification*. Oxford University Press, 2017.
Friedrich, Carl J. *Constitutional Government and Politics*. Harper and Brothers, 1937.
———. *Man and Government*. McGraw-Hill, 1963.
Gädeke, Dorothea. "Does a Mugger Dominate? Episodic Power and the Structural Dimension of Domination." *Journal of Political Philosophy* 28, no. 2 (2020): 199–221.
Goldman, Alvin I. "Toward a Social Theory of Power." *Philosophical Studies* 23, no. 4 (1972): 221–68.
Halpern, Joseph Y., and Judea Pearl. "Causes and Explanations: A Structural-Model Approach. Part I: Causes." *British Journal for the Philosophy of Science* 56, no. 4 (2005): 843–87.
———. "Causes and Explanations: A Structural-Model Approach. Part II: Explanations." *British Journal for the Philosophy of Science* 56, no. 4 (2005): 889–911.
Harris, Cheryl I. "Whiteness as Property." *Harvard Law Review* 106, no. 8 (1993): 1707–91.
Hasan, Rafeeq. "Republicanism and Structural Domination." *Pacific Philosophical Quarterly* 102, no. 2 (2021): 292–319.
Haslanger, Sally. "Oppressions: Racial and Other." In *Racism in Mind*, edited by Michael P. Levine and Tamas Pataki. Cornell University Press, 2004.
Hausman, Daniel M. *Preference, Value, Choice, and Welfare*. Cambridge University Press, 2012.
Hayward, Clarissa Rile. *De-Facing Power*. Cambridge University Press, 2000.
———. "On Structural Power." *Journal of Political Power* 11, no. 1 (2018): 56–67.
Heil, John. *The Universe as We Find It*. Oxford University Press, 2012.
Hitchcock, Christopher. "The Intransitivity of Causation Revealed in Equations and Graphs." *Journal of Philosophy* 98, no. 6 (2001): 273–99.
Ingthorsson, Rognvaldur D. "Causal Production as Interaction." *Metaphysica* 3, no. 1 (2002): 87–119.
Isaac, Jeffrey C. *Power and Marxist Theory: A Realist View*. Cornell University Press, 1987.
Kuykendall, Davis White. "In Defense of the Agent and Patient Distinction: The Case from Molecular Biology and Chemistry." *British Journal for the Philosophy of Science* 75, no. 2 (2024): 443–63.
Laswell, Harold, and Abraham Kaplan. *Power and Society*. Yale University Press, 1950.
Lewis, David. "Causation." *Journal of Philosophy* 70, no. 17 (1973): 556–67.

———. *Philosophical Papers*, vol. 2. Oxford University Press, 1986.
Lindblom, Charles E. "The Market as Prison." *Journal of Politics* 44, no. 2 (1982): 324–36.
Lovett, Frank. *A General Theory of Domination and Justice*. Oxford University Press, 2010.
March, James G. "An Introduction to the Theory and Measurement of Influence." *American Political Science Review* 49, no. 2 (1955): 431–51.
Marmodoro, Anna. "Aristotelian Powers at Work: Reciprocity Without Symmetry in Causation." In *Causal Powers*, edited by Jonathan D. Jacobs. Oxford University Press, 2017.
Marsh, David. "Interest Group Activity and Structural Power: Lindblom's *Politics and Markets*." *West European Politics* 6, no. 2 (1983): 3–13.
McDermott, Michael. "Redundant Causation." *British Journal for the Philosophy of Science* 46, no. 4 (1995): 523–44.
Mill, John Stuart. *The Subjection of Women*. Vol. 21 of *The Collected Works of John Stuart Mill*, edited by J. M. Robson. Toronto University Press, 1984.
Morriss, Peter. *Power: A Philosophical Analysis*. 2nd ed. Manchester University Press, 2002.
Nagel, Jack H. *The Descriptive Analysis of Power*. Yale University Press, 1975.
O'Connor, Timothy. *Persons and Causes: The Metaphysics of Free Will*. Oxford University Press, 2002.
Partridge, P. H. "Some Notes on the Concept of Power." *Political Studies* 11, no. 2 (1963): 107–25.
Pearl, Judea. *Causality: Models, Reasoning, and Inference*. 2nd ed. Cambridge University Press, 2009.
Pettit, Philip. "Freedom and Probability: A Comment on Goodin and Jackson." *Philosophy and Public Affairs* 36, no. 2 (2008): 206–20.
———. "Republican Freedom: Three Axioms, Four Theorems." In *Republicanism and Political Theory*, edited by Cécile Laborde and John Maynor. Blackwell, 2008.
———. *A Theory of Freedom: From the Psychology to the Politics of Agency*. Oxford University Press, 2001.
Ramachandran, Murali. "A Counterfactual Analysis of Causation." *Mind* 106, no. 422 (1997): 263–77.
Roy, William G. *Socializing Capital: The Rise of the Large Industrial Corporation in America*. Princeton University Press, 1997.
Russell, Bertrand. *Power: A New Social Analysis*. Routledge, 2004.
Schaffer, Jonathan. "Overdetermining Causes." *Philosophical Studies* 114, no. 1 (2003): 23–45.
Schelling, Thomas C. *The Strategy of Conflict*. 2nd ed. Harvard University Press,

1980.
Searle, John R. *Making the Social World: The Structure of Human Civilization*. Oxford University Press, 2010.
Simon, Herbert A. *Models of Man*. Wiley, 1957.
Skorupski, John. *The Domain of Reasons*. Oxford University Press, 2010.
Smith, Angela M. "Responsibility for Attitudes: Activity and Passivity in Mental Life." *Ethics* 115, no. 2 (2005): 236–71.
Stone, Clarence N. "Systemic Power in Community Decision Making: A Restatement of Stratification Theory." *American Political Science Review* 74, no. 4 (1980): 978–90.
Strange, Susan. *States and Markets*. 2nd ed. Continuum, 1994.
Strawson, P. F. *Freedom and Resentment and Other Essays*. Routledge, 2008.
Vrousalis, Nicholas. "The Capitalist Cage: Structural Domination and Collective Agency in the Market." *Journal of Applied Philosophy* 38, no. 1 (2021): 40–54.
Wallace, R. Jay. *Responsibility and the Moral Sentiments*. Harvard University Press, 1996.
Ward, Hugh. "Structural Power: A Contradiction in Terms?" *Political Studies* 35, no. 4 (1987): 593–610.
Wartenberg, Thomas E. *The Forms of Power: From Domination to Transformation*. Temple University Press, 1990.
Watson, Gary. "Two Faces of Responsibility." *Philosophical Topics* 24, no. 2 (1996): 227–48.
White, D. M. "Power and Intention." *American Political Science Review* 65, no. 3 (1971): 749–59.
Wilson, James Lindley. *Democratic Equality*. Princeton University Press, 2019.
Wodak, Daniel. "What Is the Point of Political Equality?" *Philosophical Review* 133, no. 4 (2024): 367–413.
Wolf, Susan. *Freedom Within Reason*. Oxford University Press, 1990.
Wright, Richard W. "Causation in Tort Law." *California Law Review* 73 (1985): 1735–828.
———. "Causation, Responsibility, Risk, Probability, Naked Statistics, and Proof: Pruning the Bramble Bush by Clarifying the Concepts." *Iowa Law Review* 73 (1988): 1001–77.
Wrong, Dennis H. *Power: Its Forms, Bases, and Uses*. Routledge, 2017.