

Dealing with Trespassers in the Kingdom of Ends

Ryan Davis

Princeton University

Prepared for: Northwestern University Society for Ethical Theory and Political Philosophy, 2nd Annual Conference, May 15-17, 2005

A tyrant threatens to starve his own people unless international sanctions against his regime are lifted. A resistance organization claims it will deploy terrorist tactics against a state's civilian population if its demands are not met. A South American militia captain warns that he will execute twenty inhabitants of a local village unless an errant botanist shoots one.

These are cases of *moral blackmail*. Moral blackmail occurs when one agent tries to manipulate another by threatening to commit some heinous moral crime if the other does not perform a comparatively less severe moral indiscretion.¹ Being the victim of moral blackmail puts anyone in a bind, but the resulting practical puzzle seems especially to plague Kantian moral theory. There is a simple reason for this: Kantian rightness is decided by a test of the agent's intention. Kantian morality is a formal constraint on action, rather than a substantive theory of value.² As a result, Kantianism displays a "though-the-heavens-may-fall" stubbornness against the unfortunate effects of acting rightly.

Sometimes this is a virtue. Kantians need not rummage around the empirical data looking for a good reason why torture should *never* be employed, to name just one example. But the case of moral blackmail pits a Kantian's theoretical absolutism against her better judgment. No matter how bad things get, orthodox Kantianism will never

¹ Cf. Terrence C. McConnell, "Moral Blackmail," *Ethics* 91:4 (1981): 544-567.

² This is not a neutral description. See Robert Johnson, "Value and Autonomy in Kantian Ethics," *Oxford Studies in Metaethics* Vol. 2.

advise yielding to moral blackmailers. What gives Kantianism its distinctive response is its denial of negative responsibility, or responsibility for what others do in response to one's action. Some count this as quixotic dogmatism. I think something better can be said for it.

In particular, I will advance three claims. First, revisiting Bernard Williams's classic case of moral blackmail, I try to show the reasons a Kantian might have for resisting—even when it brings about less value in the world. My argument will fall somewhat short of rendering the Kantian reply palatable to the faint of heart. The second section of this paper will try to do this work by accepting a modification of the theory proposed by David Velleman. Traditionally, what might be called *purist Kantianism* held that the only reasons for action were those that were produced through the formal procedure of the categorical imperative. What Velleman calls *concessive Kantianism* allows that reasons for action may have others sources, such as in the practical commitments and identities of agents.³ Concessive Kantianism succeeds at explaining how we should respond to moral blackmail while keeping intact most of what is good about Kantian absolutism, or so I will argue in section III. In short, Kantians have good reason to be wary of cooperating with blackmailers because giving in to blackmail creates a unique moral risk of disrespecting everyone involved. But Kantianism can still give an agent reasons to avoid terrible outcomes without becoming a pawn in the hands of evil.

³ J. David Velleman, "Willing the Law," in *Self to Self* (Cambridge, Cambridge University Press, 2006): 284-311.

I.

Blackmail refers to a demand paired with a threat to expose a person in some way if she does not comply. When a person's moral standing is threatened to be undermined, the blackmail is moral. Jones morally blackmails Smith by threatening to commit a greater moral wrong if Smith does not perform a lesser, otherwise immoral act.⁴

A canonical case of moral blackmail is Williams's example, Jim and the Indians.⁵ In Williams's story, Jim happens upon a local militia captain, Pedro, who is about to execute 20 villagers. As an honored guest, he is presented with the chance to shoot one of the villagers, allowing the rest to go free. If he refuses, Pedro will tell his men to shoot all of them.

Williams imagines that Jim will feel considerable anxiety about his decision. This is worth thinking about, Williams believes, because the solution to Jim's practical problem might not seem puzzling at all. He has a chance to prevent 19 deaths. "How much better off the village will be because I chanced upon it," Jim might reason, "I will be a hero!" But if Jim is morally sensitive at all, he will not count his chance to shoot a villager as a happy opportunity to make the world a little better. What would pain him is not just that a person would needlessly die. Unless Jim is a hopeless idealist, we might suppose he is aware that people needlessly die everyday (and unless he is hopelessly depressed, these facts probably do not consume his inner life). Rather, what distresses him is the thought that he will *be the person* who brings about the needless death.

⁴ I say that the demand is to perform an 'otherwise' immoral act to leave open the question of whether the fact that Jones issued the threat changes the moral situation to render the action permissible, or even obligatory.

⁵ See his contribution to *Utilitarianism: For and Against* (New York: Cambridge, 1973).

Central to Jim's dilemma is whether he is morally responsible for the actions of others. A principle of strong negative responsibility says that he is. It might be stated:

Strong Negative Responsibility: An agent is morally responsible for all outcomes that she can reasonably predict will follow from her action.

Many species of consequentialism accept some version of this thesis. If Strong Negative Responsibility is true, Jim can just as much be blamed for 19 deaths if he does not shoot as he can for one death if he does. Kantianism, however, rejects the thesis, locating rightness in formal conditions of Jim's will rather than in the outcomes of his decision.

Kantians might accept a version of Personal Responsibility, roughly:

Personal Responsibility: An agent is morally responsible only for the action she intends.

Personal Responsibility need not be accepted if Strong Negative Responsibility is denied.⁶ Still, resistance to some form of a principle allowing negative responsibility tends to reverse one's perspective on Jim's dilemma. While the consequentialist has difficulty grappling with the badness of shooting, Kantian theory provides no reason *to* shoot. To say that Pedro's actions would be Jim's fault would just be a mistake.

What motivates the Kantian's resistance? Williams famously worried that requiring Jim to shoot would be an attack on his integrity.⁷ It might seem that the badness of shooting is that Jim must kill the villager—that it is something he does rather than just allows.⁸ But this fails to capture the difference. At issue is not just that the

⁶ My purpose here is just to describe two broadly different approaches.

⁷ See, for example, Elizabeth Ashford, "Utilitarianism, Integrity, and Partiality," *The Journal of Philosophy* 97:8 (2000): 421-439; Damian Cox, "Integrity, Commitment, and Indirect Consequentialism," *The Journal of Value Inquiry* 39 (2005): 61-73.

⁸ There are many ways of making the exact distinction. See Samuel Rickless, "The Doctrine of Doing and Allowing," *The Philosophical Review* 106:4 (1997): 555-569; Kai Draper, "Rights and the Doctrine of Doing and Allowing," *Philosophy & Public Affairs* 33:3 (2005): 253-280.

villagers will die if Jim does not shoot, but that they will be killed by another agent—in particular—an agent who is acting with the intention that Jim shoot. Resisting blackmail differs from just allowing a harm to occur; not acting in the face of blackmail has something to be said for it in a way that not acting in the face of famine does not. In acting to avoid a bad outcome, the agent must *be intentionally blackmailed*.

The relevance of intention suggests another explanation. Perhaps it is worse for Jim to intend the death of one villager than to foresee but not intend the deaths of twenty.⁹ Looking a villager in the eye and pulling the trigger is especially horrible. However, it also misses the unique badness of being blackmailed. Jim’s problem is not that the same action will bring about desirable and undesirable effects. To see this point, imagine a world in which shooting the single villager was a necessary means to saving the other 19. (For instance, if the villager was about to expose the village to a horrible disease that would kill all of them, and Jim had only enough time to shoot to prevent his entry). Were this the case, Jim would face a problem similar to the (philosophically) ordinary dilemma of the person who must decide whether to throw a fat man in front of an oncoming trolley.¹⁰ But this is false. The same action that kills the one will not save the multitude. Rather, killing the one is only a means to saving the multitude because Pedro has warned that he will kill all of the villagers if Jim does not shoot.

It is tempting to think this is a distinction without a difference. Jim might happen to know that Pedro is exceptionally strong willed. Jim might be as confident that the

⁹ Christine Korsgaard discusses Jim’s predicament in this way in “Reasons we Can Share,” *Creating the Kingdom of Ends* (Cambridge: Cambridge University Press, 1996): 275-310, as does David McCarthy, “Intending Harm, Foreseeing Harm, and Failures of the Will,” *Nous* 36:4 (2002): 622-642. For a general defense of the doctrine of double effect, see Alison Hills, “Intentions, Foreseen Consequences and the Doctrine of Double Effect,” *Philosophical Studies* 133 (2007): 257-283.

¹⁰ Much more could be said. For possible variations, see Frances Kamm, “Responsibility and Collaboration,” *Philosophy & Public Affairs* 28:3 (2000): 1-36.

villagers will die if he does not shoot as the person behind the fat man is that more people will die if she does not push. Still, I believe there are three reasons that suggest something else significant is going on.¹¹

First, notice that we can only think of Jim's predicament as parallel to the trolley case if we envision Pedro as the trolley—a mere force of nature whose next movements will be propelled without any volitional control. That is not what Pedro is. Pedro will have to form an intention of his own to take the lives of the villagers. The blackmail's success hinges on Jim's taking Pedro's next move for granted: if Jim did not believe that Pedro's action would necessarily follow from his own, he would not have much reason to kill an innocent bystander. The more clearly Jim sees Pedro's action as volitional—as authorized by Pedro's will—the more difficult it might be for him to see why he should follow Pedro's demand. Consider a case:

Enemy: Your enemy believes that you are a moral anti-expert. That is, she believes that everything you do is the opposite of what ought to be done. Every time you make a right choice, your enemy will determine that she should commit to doing the opposite action. As it happens, she has the capacity to correctly ascertain what she ought to do. She is responsible for deciding to act contrary to how you act. She will take your right actions as reasons to act wrongly.

It is certainly possible to think that the right thing to do in *enemy* is to act so as to promote the best consequences overall, even if this means doing a wrong thing so that your enemy will be led to do the right thing. Still, when the volitional action of the enemy is treated as a subject of her own choice, rather than as a causal fact about the

¹¹ My approach here is related to Thomas Hill, "Moral Purity and the Lesser Evil," in *Autonomy and Self-Respect* (Cambridge: Cambridge University Press, 1991), but I eschew reasons that appeal to values.

universe, it becomes more plausible to deny a strong doctrine of negative responsibility.¹² “It’s just not my fault!” you might respond. Taking the point of view of your potential victim sets this point into relief. If you decide to act badly to bring it about that your enemy acts well (supposing this will be best), your explanation likely be of little comfort to the victim of your action. He would likely still resent you, and his resentment would be warranted by facts about your intention.

Now the first consideration dividing moral blackmail from trolley-type cases can be seen more readily. When the runaway trolley will cause the moral disaster, you are treated in figuring that into your set of considerations simply as a causal fact. But there is something amiss in treating another person’s will in this way. It regards her as a piece of furniture in the case, bound to act causally in systematic ways. Thinking of a person as a mere causal force to be dealt with fails to treat the blackmailer as an agent too—as an end in herself. Although it sounds strange to say that you owe it to the blackmailer not to be blackmailed, there is a sense in which—Kantians and their sympathizers at least—may feel this pressure.

Second, caving to blackmail treats yourself as a mere means. If you have to respond to a trolley about to hit five people by diverting it to kill one, we might describe you as simply responding to unhappy circumstances. Although you are in a bind, you are the agent who authorizes your own actions to get out of that bind. When you act as a blackmailer wishes, you are not just letting unhappy circumstances decide what you will do, you are letting another person decide your actions for you. In Christine Korsgaard’s

¹² For a suggestive analysis of these themes, see Samuel Scheffler, “*Doing and Allowing*,” *Ethics* (2004): 215-39.

evocative phrase, you become the “tool of evil.”¹³ It is not just that the blackmailer uses you as his tool. It is worse than that. In deciding to form your intention to act in the way directed by the blackmailer, you—yourself—decide to treat yourself as tool. You become a tool of evil of your own fashioning.

None of this demands accepting the Kantian line. Jim certainly can still accept strong negative responsibility. It would be wrong, on this view, for Jim to not shoot if he knows this is the only way to save lives, given what others will do. If so, Pedro demands that Jim shoot, and—surprisingly—morality does too. While there are no two ways about the fact that accepting negative responsibility does bring about the best outcome, this also gives up something worth pondering. To believe that morality requires Jim to shoot means that moral blackmail can be successful. That is, it can succeed not just as a matter of fact, but as a matter of *moral* fact. Jim wants to minimize death because he is sensitive to a moral prohibition of killing, and it is his very sensitivity to this moral rule that Pedro uses against him. Pedro aims to change what Jim should do by manipulating his moral commitments. In effect, he attempts to use the moral prohibition of killing as a means to induce killing. And if strong negative responsibility is accepted, Pedro succeeds. He *actually does* use the moral law as a means to moral evil.

II.

Of course, what I have been able to offer so far is an outline of why Jim might be right to feel reluctant to shoot. The considerations offered presume some sympathy with the

¹³ See her “The Right to Lie: Kant on Dealing with Evil,” in *Creating the Kingdom of Ends* (Cambridge: Cambridge University Press, 1996).

Kantian project to begin with.¹⁴ But no one would suggest that mere recognition of the profundity of Jim's decision betrays a moral error. Rather, that it provides evidence against shooting counts in Kantianism's favor (the opposite of Williams's point against consequentialism). What is unnerving is that Kantianism doesn't stop there. The startling thing about the orthodox Kantian advice to Jim is how unperturbed it is.¹⁵ Jim not only should not shoot the villager, but he has *no reason* to shoot. This is just the kind of thing that provides fodder for an attack against Kantian theory: it seems aloof, insensitive to human suffering, and fixated on the moral uprightness of the agent at the expense of others.

To see the problem, it's important to be clear about why Kantianism produces this extension. Only a sketch can be provided here. A basic point is that Kantian theory holds that it is never rational to intend a morally impermissible action. That morality marches in step with rationality is a particularly Kantian slogan. Here is why: for Kantians, practical reasons arise from autonomous exercises of the will. Only a principle that I freely adopt is one that provides me with reasons to act. That is, any consideration can be counted as a reason for acting so long as it could be willed as such with universal validity.¹⁶

So far, this is just a statement of the categorical imperative. Usually a discussion of the categorical imperative focuses on what actions it rules out, like lying or (normally)

¹⁴ An important problem, which I will not confront here, is in deciding whether the first section presented considerations against shooting, or constraints forbidding it. As it will become clear, my solution will not treat them as constraints.

¹⁵ The view I have so far referred to as "orthodox Kantian" is what I will term "purist Kantianism."

¹⁶ The categorical imperative is the test of whether a consideration can be counted as a reason for acting. An action is permissible if the principle of practical inference connecting the reasons for the intention to its end could be held in common by all similarly situated reasoners. This is a way of describing Kant's idea that one's maxim must be universalizable. See Velleman, "The Voice of Conscience," in *Self to Self*, pp. 110-128.

coercing. For present purposes, it is more important to see is what the categorical imperative shows about candidate practical reasons for impermissible actions—considerations that count in favor of acting in a way that is non-universalizable. On the Kantian view, reasons for action—just to be reasons at all—can only be those that have universal validity. In other words, there is no normative reason to act wrongly.

This fact brings good news and bad news. The good news is that Kantianism shores up moral absolutes that we intuitively believe are important to hold without exception, but otherwise have difficulty justifying as absolutes. The bad news is that it is unclear why any permissible action should be chosen above any other, equally permissible action. When all that provides a reason to perform an action is that its intention could be autonomously willed, the reasons for any given course are just that choosing that course could be an instance of autonomous authorization. Unfortunately, this is not enough for a theory of practical reasoning, simply because it does not offer any guidance in deciding what to do. David Velleman describes the set of potential actions:

Even if we believe that this set would exclude any morally impermissible actions, we must doubt whether the agent can will distinctions of correctness among the remaining, permissible alternatives. Within the constraints of the Categorical Imperative, the agent appears to face an arbitrary choice among various universal rules, which would specify various actions as correct in light of the circumstances, variously considered, thereby constituting different considerations as reasons for taking different actions.¹⁷

Now we have the materials in place to revisit Jim and Pedro. Jim must decide whether he will comply with Pedro's demand that he shoot a villager. A Kantian adviser might tell him something like this: "You can either shoot or not. Shooting the villager involves taking up an intention directed to you by Pedro, whose intention could not be held with

¹⁷ Velleman, "Willing the Law," p. 294.

universal validity.¹⁸ Maybe, just maybe, your intention would be sufficiently different from Pedro's that it could be permitted, but this is uncertain. Shooting is at best gambling with your good will, but not shooting involves no risks at all. Not shooting is an intention that could certainly be held by everyone. All told, there seems to be no reason to shoot."

This advice is from what might be called purist Kantianism, the version of the theory considered so far. Notice that the same feature that impoverishes purist Kantianism as a theory of practical reasoning also blinds the theory to the considerations in favor of shooting. It says:

Purist Kantianism: Only considerations ratified by the categorical imperative count as normative reasons.

To solve the problem of Kantian practical reasoning, Velleman proposes what he terms concessive Kantianism. It might be distinguished:

Concessive Kantianism: Considerations arising from one's practical identity *and* considerations ratified by the categorical imperative count as normative reasons.

Of course, these principles do not represent all of Kantian theory; they just highlight one contrast. With it, we can disentangle how these two Kantianisms would advise Jim.

Purist Kantianism's advice. An action that fails to treat another as the author of her actions will fail the test of universal validity. Shooting someone seems paradigmatic of denying another's choice, but this need not always be so. The issue hinges on what the villagers want Jim to do. If the villagers are all pacifists who would denounce Jim's decision to shoot one of them even when they knew

¹⁸ On acting directly on another's intention, see Abraham Sesshu Roth, "Shared Agency and Contralateral Commitments," *The Philosophical Review* 113:3 (2004): 359-410.

his abstention would result in twenty deaths, Jim would be wrong to override their wills in the name of their welfare. Alternatively, one villager might volunteer to be shot to preserve the lives of the others, offering to forgive Jim in advance. While the villager would still have a right to resent his plight, in the “smaller moral world” comprised just of the villager and Jim, the villager would be appropriately respected.¹⁹ When the reasons run out, Jim can permissibly shoot or permissibly not shoot.

Purist Kantianism never finds an obligation for Jim to shoot, or even a practical reason for Jim to shoot. At best, it allows Jim to shoot, and this is only possible when additional information about the case is provided. Contrast with an alternative approach:

Concessive Kantianism’s advice. As before, actions must have universal validity to be morally permissible. However, Jim’s various practical identities can also be considered. A practical identity is an unselfconscious source of reasons.²⁰ For example, suppose Jim discovers that two expeditions are leaving: one to search for a new plant species, the other to search for a rare bird. If Jim values himself under the identity of a botanist (but not an ornithologist), he will find that he has reason to attend the first expedition. His practical identity provides a deliberative shortcut in deciding what to do. Perhaps another of Jim’s practical identities is his solidarity with Andean villagers (he might attend rallies, or advocate their rights in the academy back home). If nothing else, Jim surely values himself as a person who desires to alleviate the suffering of others. Given these practical

¹⁹ Both of these specifications of the case are from Christine Korsgaard, “The Reasons We Can Share,” in *Creating the Kingdom of Ends*, p. 296.

²⁰ The reason for this, I believe, is that questions about practical identity are closely related to questions about personal identity and time. See J. David Velleman, “Self to Self,” in *Self to Self*, pp. 170-202. My use differs with Korsgaard, *The Sources of Normativity* (Cambridge: Cambridge University Press, 1996).

identities, Jim will find that he has great reason to preserve the village, if he can permissibly do so. He will have reason to shoot.

What should Jim do? Suppose—as might be expected in more real-worldly cases of moral blackmail—that he cannot determine what the villagers want him to do. According to purist Kantianism, Jim does best to simply keep his hands clean by not shooting. On the concessive Kantian model, things are not so clear. Jim is in a real dilemma—at least in an epistemic sense. If, as in Williams’s original telling, the villagers unanimously want Jim to shoot, then he has decisive reason to shoot, given his configuration of practical identities. But if they are pacifists, as in Korsgaard’s telling, shooting is impermissible. Jim’s situation is an epistemic moral dilemma because there is a fact of the matter about what Jim has most reason to do, and Jim does know *that*, but he does not know what that fact is.

Jim is in a position of moral risk. A moral risk is a situation in which an agent must Φ or not, and the agent knows that Φ -ing is either morally impermissible or that there is decisive reason to Φ , but the agent does not know which. At this point, Jim must take morality into his own hands.²¹ He must either not shoot, knowing that his refusal could be a betrayal of the villagers and himself. Or he must shoot, knowing that if his action is impermissible, the villagers would appropriately resent him.

Concessive Kantianism is less decisive in advising Jim than either simple-consequentialism or purist Kantianism. But in this case, underdetermination is a virtue. The problem with other theories was never that the case gave them trouble, but that they were not sufficiently troubled by it. Concessive Kantianism represents the moral world

²¹ Cf. Christine Korsgaard, “Taking the Law into Our Own Hands: Kant on the Right to Revolution,” in Korsgaard, Reath, and Herman [eds.] *Reclaiming the History of Ethics: Essays in Honor of John Rawls* (Cambridge: Cambridge University Press, 1997): pp. 297-328.

more authentically, if less decisively. Admitting as much concedes that Kant was not altogether right in denying the gap between theory and practice. Sometimes, persons experience moral dilemmas because they really do stand in positions of moral risk—which is exactly what non-philosophers think anyway.

III.

Concessive Kantianism deals well with moral blackmail. One worry about my proposed solution is that modifying Kantianism to be less absolute when it seems right to be puzzled will also make Kantianism less absolute when it seems right to be stalwart.

One implication of concessive Kantianism is that morality and rationality no longer march in step. If it is possible that real normative reasons oppose acting rightly, an agent may occasionally find that she has most reason to *not* do what morality recommends. In such times, Velleman suggests that the concessive Kantian can still maintain that the agent only found her way into that position by developing practical identities in an irrational way. For example, a Mafioso may have most reason to kill an informant, but only because he irrationally acquired the practical identity of a Mafioso.²²

Can Velleman's defense cover all rational departures from morality? There is reason to fear it might not. Recall the example of torture. Many believe—and I agree—that torture should be absolutely off the table of options.²³ Let us suppose that torture is never morally permissible. Imagine the typical case where a government officer has a detained subject with knowledge of a bomb that will soon destroy an entire city. There is

²² See Velleman, "Willing the Law," pp. 302-309.

²³ Cf. David Sussman, "What's Wrong with Torture," *Philosophy & Public Affairs* 33:1 (2005): 1-23; Vittorio Fufacchi and Jean Maria Arrigo, "Torture, Terrorism and the State: A Refutation of the Ticking-Bomb Argument," *Journal of Applied Philosophy* 23:3 (2006): 355-373.

every reason to think that much in the officer's practical identity will provide her with reasons to do everything possible to save the city. She values herself as an officer of the law, as a government agent specially charged to protect others, and as a citizen. Certainly the balance of such reasons might rationally exceed countervailing moral reasons just as they do for the Mafioso, but not in virtue of any prior irrationality. This is the crucial difference. We cannot rationally criticize the officer for any of her acquired identities, so morality seems to have unfastened further from rationality than Velleman intended.²⁴

I will conclude by briefly noting one possible response. Torture is no ordinary activity. Even a single instance of torturing another human being will be apt to leave psychological scars, no matter how important the torturer's reasons.²⁵ Performing torture might change the way in which one unselfconsciously decides that some considerations are reasons and others are not, and it might do this in very fundamental ways. In other words, torture could destroy part of an agent's practical identity, and that might be irrational even when other moral indiscretions supported by non-moral reasons are not.²⁶ Our practical identities can take a few hits when we act against them and remain intact, but they will not withstand anything. Where it is important that Kantianism remain absolute, it is likely that the facts about who we are will be more vulnerable to change.

Here there is a parallel between the purist and concessive Kantianisms. In the purist version, a person rationally ought not perform any single action that brings about a contradiction in her will. In the action at hand, there should be no contradiction in her

²⁴ At least, this is so if Velleman maintains that for all agents, "[if] an agent can lack sufficient reason for doing the right thing, I would insist that such an agent is nevertheless irrational," *Self to Self*, p. 285.

²⁵ Assuming some plausible empirics. Cf. Jessica Wolfendale, "Training Torturers: A Critique of the Ticking Time Bomb Argument," *Social Theory and Practice* 32:2 (2006): 269-287.

²⁶ This admittedly controversial thought cannot be developed here, but I will note that is only about the agent's practical identity and not about personal identity. It may therefore be more palatable as a friendly amendment to the concessive Kantian view than related arguments about the self. For a discussion, see Velleman, "Identification and Identity," in *Self to Self*, pp. 330-360.

reasons. The concessive version is analogously concerned with a person's reasons, but in a temporally extended way. On this reading, a person rationally ought not act in a way that will change her practical identity from one in which she unselfconsciously takes morality to provide reasons for acting to one in which she does not. The emphasis on purity of will to the exclusion of other considerations is inadequate as a theory of practical reasoning, but a weaker rational standard survives. It would not make sense for an agent to value oneself under the identity of a moral person while also acting so as to destroy that practical identity.

Conclusion

This last thought gives one more piece of advice to Jim, as well as to other victims of moral blackmail. Jim faces moral risks in deciding how to respect others, and now we can see that Jim also faces moral risks in preserving his self-respect. If Jim yields to blackmail too often, or when too many of his practical commitments are at stake, he may risk losing a practical identity that he values. For instance, imagine a person routinely blackmailed into making minor moral concessions. No matter how important her reasons for so acting, it is hard to believe that after a while, she won't even think twice about giving up the moral high ground. Many dislike Kantianism's focus on clean hands, but few would be unconcerned with the person whose hands were so dirty that she had lost the desire to wash them. While caring only about moral purity betrays puritanical zeal, caring for it not at all betrays a lack of self-respect. Sometimes we should ignore the trespassers and live in our own Kingdom of Ends, much as Kant would recommend.²⁷

²⁷ My thanks to Charles Beitz, Brookes Brown, and Tristram McPherson.

