

In Defense of Hard Paternalism

Danny Scoccia
Department of Philosophy
New Mexico State University
Las Cruces, NM 88003
Phone: 505-646-4616
E-mail: dscoccia@nmsu.edu

“Hard” paternalism—understood (preliminarily) as restricting a competent adult’s liberty for his own good under conditions that “violate his autonomy”—is rejected as morally wrong by the majority of liberal theorists who write on the paternalism issue.¹ When the state engages in such paternalism, it imposes on its citizens a “conception of the good life” that some of them reject, thus violating the requirement (upheld by many of these theorists) that the state remain “neutral” on the question of the good life. An individual who engages in hard paternalism towards a competent adult is often faulted for arrogance, elitism, or lack of respect for others. The following discussion will ignore the objection to a hard paternalist coercive legislation that derives from the neutrality principle. Instead, it will focus on the claim—pressed most vigorously by Joel Feinberg—that the principle of autonomy is absolute or inviolable, and thus, trumps beneficence whenever the two values clash. If it is correct, then hard paternalism, whether practiced by the state or an individual, is always wrong. This paper argues that it, along with the often accompanying claim that any paternalism that is “soft” is unobjectionable, is not correct.

We begin with an analysis of what paternalism is and then turn to the “soft”/“hard” distinction. After criticizing Feinberg’s absolutist anti-hard paternalism (AAHP), we’ll consider a more pragmatic form of the view but reject it on the grounds that it cannot

¹ Opponents of hard paternalism include: Richard J. Arneson, “Mill versus Paternalism,” *Ethics* 90 (July 1980): 470-89; Gerald Dworkin, “Paternalism,” reprinted in Gerald Dworkin (ed.), *Mill’s On Liberty: Critical Essays* (Lanham, MA, Rowman and Littlefield: 1997), pp 61-82; Joel Feinberg, *Harm to Self* (New York: Oxford University Press, 1986); Alvin Goldman, “The Refutation of Medical Paternalism,” reprinted in John D. Arras and Bonnie Steinbock (eds.), *Ethical Issues in Modern Medicine*, 5th ed. (Mayfield Publishing, 2001); John D. Hodson, “The Principle of Paternalism,” *American Philosophical Quarterly* 14 (1977) 61-69; and Donald VanDeVeer, *Paternalistic*

accommodate liberal support for the legalization of physician-assisted suicide with eligibility limited to the terminally ill, and for a ban on prostitution. The key here is that soft paternalism cannot justify these measures; liberals who support them have no alternative but to defend them on hard paternalist grounds. Finally, this paper suggests a couple of strategies for explaining why hard paternalism seems clearly wrong in some cases (e.g. forced blood transfusions for adult Jehovah's Witnesses) but not others.

What Is Paternalism?

Paternalism is the use of certain means to induce others to behave in ways that we think are good for them or to stop behaving in ways that we think are bad for them. Which means? A common mistake is to define "paternalism" in a way that limits them to coercive ones. Probably coercion is the most common method that paternalism employs. So, for example, if I try to convince you that you will die of poisoning unless you swallow this antidote immediately, but I fail, I act paternalistically towards you if I tie you up and force feed it to you for your own good. But threats and restraints are not the only means by which paternalism can operate. Incentives expand rather than restrict one's options, and they belong in the paternalist's repertoire too. If I promise my teenage son monetary rewards if he earns high grades in school, I act paternalistically towards him.

Suppose that my dying friend asks me where his daughter and grandchildren are. Moved by the desire to spare him pointless grief, I tell him that their flight was cancelled at the last moment instead of the truth that their plane crashed and they all died. That's a

Intervention (Princeton, N.J.: Princeton University Press, 1986).

paternalistic lie. In addition to deception, there is withholding information from or failing to reveal it to another because one expects that he will act foolishly if he has it. For example, I might know that there will be a party with teenagers, drugs, and alcohol later tonight but decide not to tell my teenage son about it because I fear that he will want to attend and then succumb to peer pressure and do something foolish. I decide not to expand my son's options by not sharing with him information that I know he would like to have. My action here is paternalistic, though neither deceptive nor coercive.

There are a number of other noncoercive and nondeceptive methods available to the paternalist that involve structuring others' options in certain ways or taking advantage of the fact that they are already structured in certain ways. One sort of case is where a homeless person asks you for \$5 to buy a hamburger and you give him the hamburger rather than the money because you're afraid he would spend the money on booze. Still another case is where we change somebody's "default" options, so that instead of his having to decide to do something that will give him a good outcome, he would have to "opt out" of an arrangement that would give him the same outcome. For example, the state might require employers to set up a saving account for each of its workers and deposit a small percentage of his or her earnings into the account (with some matching funds) *unless* the worker opts out of the plan. It turns out that a smaller percentage of workers opt out than would go to the bother of setting up the accounts and contributing to them on their own initiative.² Thus, the arrangement succeeds in getting more people to save by exploiting their inertia. These are both examples of noncoercive paternalism.

² "A study by Brigitte Madrian and Dennis Shea [found that] changing the default rule in this way raised the enrollment rate from 49% to 86%." ("The Avuncular State,"

The possible means available to the paternalist are varied. They include deception, incentives, threats, physical restraints, the withholding of information, forms of trickery and manipulation, exploiting preexisting constraints, and structuring the person's options to elicit the desired behavior. What these tactics have in common is that they are *not* an attempt at reasoned persuasion focused on the merits of the case. This gives us the essence of paternalism. It is an attempt to induce others to do what's good for them or avoid doing what's bad for them, in a way that *doesn't* involve trying to persuade them to act or not act in the ways in question.³

Few would deny that paternalism is permissible in dealing with young children, the severely mentally retarded, or paranoid schizophrenics. Under what circumstances, if any, is it a permissible way of dealing with competent adults? The Kantian Principle of Humanity (PH), requiring that we treat the humanity in ourselves and others always as an "end" and never as a "mere means," can be interpreted as implying that the answer is almost never. The "humanity" in us is our capacities for free will and rationality, and respecting persons on this interpretation of the principle requires that we try to elicit prudent conduct from others *via* those capacities by showing them why the conduct in question is good or bad for them. Creating incentives or making threats does not *offer* reasons but instead attempts to *change the balance* of reasons so that now even someone who is fairly foolish or short sighted will see that he has good reason to act in the desired

The Economist, April 8, 2006, p. 67).

³ This analysis of "paternalism" differs from the one offered by Gerald Dworkin in his article "Paternalism" in the *Stanford Encyclopedia of Philosophy* [available online, at <http://plato.stanford.edu/entries/paternalism/>]. Dworkin assumes that a paternalistic act must "interfere with" the targeted person's "liberty or autonomy." On the view defended here, paternalism doesn't necessarily hinder autonomy/rational agency but does at least

way (namely, so as to obtain the reward or avoid the punishment). Treating competent adults as ends requires (on this reading of PH) that we try to persuade them to do what's best for them even if we think it likely that our attempts will not succeed. Accordingly, paternalism toward competent adults with the ability to appreciate and act on good reasons is always wrong, no matter which of the nonrational means it employs, whenever the opportunity to engage in persuasion rather than paternalism exists. I'll call this view "simple Kantianism."

Soft and Hard Paternalism

We'll return to simple Kantianism shortly. First, we should examine a couple of distinctions between types of paternalism that several writers have employed. The first is between "weak" and "strong" paternalism. Weak paternalism targets people whom it's thought have adopted means ill suited to achieving their own goals, while strong targets people whose goals are supposed to be mistaken or misguided in some way.⁴ A second distinction is between "soft" and "hard" paternalism. Several writers, including Feinberg, Arneson, and Dworkin, define the former as paternalism that's directed at substantially nonvoluntary choices, while the latter is paternalism directed at substantially voluntary choices. Note that neither distinction is concerned with the methods the paternalist

ignore or circumvent it.

⁴ Both Richard J. Arneson in "Paternalism, Utility, and Fairness" [reprinted in Gerald Dworkin (ed.) *Mill's "On Liberty": Critical Essays*] and Gerald Dworkin [in the "Paternalism" essay in *SEP*] draw the strong/weak distinction this way.

employs. Hence, weak and soft paternalism can use coercive methods, while strong and hard paternalism can use noncoercive ones.⁵

The point of drawing these two distinctions is to use them to do some normative work—to argue, for example, that only and/or all weak paternalism is permissible, or that only and/or all soft paternalism is permissible. Defenders of each of those claims end up appealing to the same set of intuitions. The problem that Feinberg sees with (coercive) criminal legislation based on “hard” paternalism is that it “imposes its own values and judgments on people.”⁶ It violates the person’s autonomy by subjecting him to the *alien* will, preferences, and values of another. Soft paternalism, by contrast, is interference on behalf of the person’s *own* will, preferences, and values, and thus does not violate the person’s autonomy. Surely the opponents of “strong” paternalism wish to lodge a similar objection to it.

If that intuition is the basis of our condemnation of paternalism (when we do condemn it), then it should lead us to reject simple Kantianism. Consider the earlier example where you refuse to swallow the antidote because of your mistaken belief that you have not been poisoned. If my attempts to convince you to swallow it fail—you suspect me, perhaps with good reason, of trying to play a practical joke on you—then according to simple Kantianism, it would be wrong for me to force feed it to you. But if

⁵ See Feinberg, *Harm to Self*, p. 12. Feinberg notes that he used the terms “strong” and “weak” in his earlier “Legal Paternalism” paper in the same he now uses the terms “hard” and “soft,” and he claims (probably correctly) that most writers on paternalism use the two pairs of terms interchangeably. Some writers (for example, those of the article in *The Economist* noted above) seem to identify “soft” with noncoercive and “hard” with coercive paternalism. This is either an alternative but confusing usage or a misunderstanding of a well-established distinction.

⁶ Feinberg, *Harm to Self*, p. 12.

coercing you in this way can be justified by appeal to your own values and preferences, then it does not subject you to the foreign will of another. The paternalism is both weak and soft, and it is supported by the intuition to which Feinberg appeals.⁷

“Hard” paternalism that uses deception, coercion, or the like is supposed to be paternalism that violates the “principle of autonomy.” This principle holds that competent adults should be left free to make their own decisions about how to live their lives based on their own preferences, religious beliefs, conceptions of justice and virtue, beliefs about honor and dignity, and views about what’s prudentially best for them.⁸ Paternalism in the poisoning example is “soft” because it does not violate this principle. By contrast, forcing a life saving blood transfusion on an adult Jehovah’s Witness who refuses to authorize one for religious reasons does. The view of commonsense morality is that the principle of autonomy sometimes conflicts with the principle of beneficence (the principle that motivates paternalistic interference with others’ actions), and in some of those cases beneficence trumps autonomy.⁹ Hence, commonsense morality holds that hard paternalism is sometimes justified. Feinberg wants to deny this, insisting that autonomy always trumps beneficence when they conflict.

⁷ The choice with which it interferes is nonvoluntary because based on ignorance/mistaken belief. “If a person playfully illustrates the game of Russian roulette with a fully loaded six-shooter, it utterly vitiates the voluntariness of his actions to show that he doesn’t know the gun is loaded, and any better-informed spectator owes it to him to intervene forcibly for his sake.” *Harm to Self*, p. 159.

⁸ If we prevent a religious fanatic from killing people whom he thinks have blasphemed his religion, then either that’s *no* violation of the principle at all (because the principle only applies to choices that don’t violate others’ rights) or it’s a violation that is justified by the need to protect others’ rights. Nothing turns on which alternative we adopt.

⁹ Dan Brock, “Paternalism and Autonomy,” *Ethics* 98 (April 1988) 550-565, p. 551.

If the “soft”/”hard” distinction is supposed to be related to the principle of autonomy in the way just described, then it turns out that Feinberg’s way of drawing it does not achieve that goal. It is possible for a choice to be substantially nonvoluntary (or nonautonomous) yet interference with it to violate the autonomy principle, and conversely, it’s possible for a choice to be substantially voluntary yet interference with it not to violate that principle.

An example involving “precommitment” shows that interference with fully voluntary or autonomous choices need not violate the autonomy principle. Sometimes the point of precommitment is to guard against future (voluntary) weakness of will or (involuntary) psychological compulsion where one expects no change in one’s values or preferences between now and then. The story of Ulysses and the Sirens illustrates this, as does a Missouri law that excludes “gambling addicts” from its 11 riverboat casinos.¹⁰ But it’s possible for one’s aim in precommitment to be not to protect oneself from any future weakness of will, but to ensure that one will not be able to act on preferences and values different from those one currently holds. Here are two examples. In the first, call it “super covenant marriage,” my fiancée and I are convinced for religious reasons that divorce for any reason is a sin, and accordingly, desire a marriage contract enforced by the state that makes our union indissoluble so long as we both remain alive. In the second I authorize the religious order to which I currently belong to kidnap and “reprogram” me should I ever convert to a different faith in the future. Such a conversion, I’m now convinced, would result in my eternal damnation, and I authorize the kidnapping and

¹⁰ Violators can be arrested and have their winnings confiscated. The law only applies to those who choose to put themselves on the list because they wish to overcome a

reprogramming as a precaution against such a spiritually disastrous eventuality. Enforcing my past precommitment wishes in these two examples would mean overriding my present, fully *voluntary* choice to act contrary to my old views about marriage and religion in accordance with my new ones. But it would not violate the autonomy principle, because it would only subject me to my own will, not someone else's.¹¹

Conversely, a choice could be “substantially nonvoluntary” yet interference with it still violates the autonomy principle. Consider the typical smoker, who admits that smoking is bad for him but makes little or no effort to quit, though he sincerely intends to (try to) quit at some later time. If the choice to continue smoking in spite of this admission is due to laziness or lack of resolve, then it remains substantially voluntary. If it's the result of psychological compulsion (“nicotine addiction”), then it's involuntary. Regardless of which of these alternatives is correct, it looks like the choice to continue smoking should count as involuntary for Feinberg insofar as it's based on ignorance/mistaken factual belief about the likelihood of the smoker's succeeding when she finally does decide to try to quit. Most smokers seem to underestimate how addictive nicotine is and/or overestimate how much willpower they will be able to bring to bear to the task of quitting.¹² But we cannot infer from this that a ban on smoking for smokers' own good is consistent with the autonomy principle. In all likelihood a ban would impose

gambling addiction. See *The Economist*, p. 67.

¹¹ It might be objected that when my past will differs from my current will, it is just as alien to me as another person's. Such an objection is unconvincing, failing to take seriously personal identity over time

¹² Thaddeus M. Pope [“Balancing Public Health Against Individual Liberty: The Ethics of Smoking Regulations,” *University of Pittsburgh Law Review* 61 (2000): 419-98] cites an FDA source from 1996 which claims that although “15 million smokers try to quit, fewer than 3% achieve one year of abstinence.” (p. 468)

on most of them values and preferences that they reject. Most smokers would indignantly oppose a ban, saying something like: “I’ll try and quit when I’m good and ready. But if I try and fail, that’s nobody else’s business. I’d rather be a smoker for the rest of my life who tries occasionally to quit and fails—even if that means cutting about ten years off of my life expectancy—than be forced by others to quit now. My life belongs to me, not other people.” This response is based on a strong desire to make decisions about how to live one’s life *oneself*. Probably nobody prefers making his own choices to interference by others in cases where noninterference would lead to imminent, extreme, and irrevocable harm to self, as in the poisoning example. But in cases where the harm to self is not imminent, extreme, and irrevocable—as with the choice to smoke another cigarette—such a preference may well loom large. If the typical smoker has this preference—he prefers to be a free nonsmoker to a free smoker, but also prefers to be a free, weak-willed smoker to a coerced (by others) nonsmoker—then a ban on smoking for smokers’ own good violates the autonomy principle and is hard rather than soft paternalism.¹³

The religious precommitment and smoking examples show that the “soft”/“hard” distinction does not coincide with the “nonvoluntary”/“voluntary” distinction. Hence, if the “soft”/“hard” distinction is to do the work expected of it, it has to be redrawn. “Soft”

¹³ Assuming that our smoker is addicted rather than merely weak-willed and that addiction involves psychological compulsion, then this is true of him: though he prefers not be coerced at all, he prefers “internal” to “external” coercion. The existence of a preference to make one’s own decisions, even if they are foolish as judged from one’s own point of view, is the reason why Gerald Dworkin is wrong when he says about weak-willed smokers that “[if we stop them from smoking] we are not imposing a good on someone who rejects it. We are simply using coercion to enable people to carry out their own goals.” (“Paternalism,” p. 80)

paternalism should be defined in terms of the satisfaction of either a “prior authorization” or “hypothetical consent” condition, while “hard” paternalism involves cases in which neither condition is satisfied. The reason why the paternalism in the poisoning example counts as “soft” is not that your refusal to ingest the antidote is substantially nonvoluntary, being based on ignorance, but rather that you *would* authorize the coercion if you were thinking clearly (calmly, deliberately) and didn’t suffer from any ignorance (of empirically demonstrable facts about your options). Whenever a hypothetical consent condition of this sort is satisfied, the paternalism can be justified by appeal to the person’s *own present preferences and values* and for that reason does not violate the autonomy principle. Note that this is *not* “hypothetical rational consent,” if by that is meant “what a person with perfectly or ideally rational values and preferences would consent to.” The hypothetical consent condition built into the definition of soft paternalism requires that we take the person’s preferences and values as given, “warts and all.”¹⁴

It’s fairly obvious in the poisoning example that the hypothetical consent condition is satisfied. Other cases may be more difficult. Suppose Jeffrey is trapped in a truck that has caught fire, and he will suffer the horrible fate of being burned alive if I do nothing. I could kill him quickly and painlessly, but he urges me not to.¹⁵ Would killing Jeffrey in spite of his denial of consent count as soft paternalism? It depends on what his reason is

¹⁴ This consent condition is similar to what VandeVeer calls “hypothetical individualized consent.” For VandeVeer this is what one would consent to if fully informed and “one’s normal capacities for deliberation and choice were not substantially impaired” (VandeVeer, p. 75). Though this condition does not count the irrationality of goals as an “impairment,” I’ll assume that it does count any failure in means-ends rationality as one. The idea that means-ends rationality is a “value” that some people might reasonably reject is probably not a coherent one

¹⁵ I believe that this example is H.L.A Hart’s.

for denying it. If he denies it only because he clings to the false hope that he will be rescued at the last moment when in fact there is no chance that help will arrive in time, then the hypothetical consent condition is satisfied and killing him for his own good is soft paternalism. But suppose that the reason why he won't consent is his belief in the "sanctity of life," one implication of which is that it's always wrong to intend one's death. Then the hypothetical consent condition is not satisfied. If I kill Jeffrey anyway because I believe that he's better off dying quickly and painlessly than being burned alive, that's hard, not soft, paternalism.

Hard paternalism, then, is the practice of deceiving, withholding information from, coercing, or using any of the other means distinguished earlier to induce another person to do what's best for him, where neither the prior authorization nor the hypothetical consent conditions are met. The hard paternalist insists that the absence of prior actual or present hypothetical authorization is of no moral consequence, because it is due to the person's having *mistaken values*. Sometimes the claim will be that the person has a mistaken conception of what is prudentially best for him. The hard paternalist who supports a ban on smoking has to say that smokers exaggerate the prudential value of making their own choices themselves and underestimate the prudential value of increased longevity and improved health. The example of Jeffrey trapped in the burning truck shows that the "mistake" could lie elsewhere. The mistake there, if there was one, would be Jeffrey's belief that a sanctity of life ethic is true and overrides the dictates of prudence. Or consider the case of Jack, who intends to commit suicide because his sacrifice would make it possible for a thousand wealthy people to receive free foot massages. Jack fully realizes that suicide would be very bad for him, but he believes that it's right on account of (what

he conceives as) the good it would do for so many others. Another possibility is that Jack wishes to commit suicide because he believes that he deserves death as punishment for being unfaithful to his wife. Preventing Jack's altruistic or retributive suicide for his own good would be hard paternalism. But it would be motivated by the idea that he has a mistaken view about the importance of his own prudential good in relation to other values, rather than the idea that he has a mistaken conception of what his good is.

Anti-Hard Paternalism

“Hard” paternalism can employ incentives rather than coercion, and whenever it does it will not violate the autonomy principle. Suppose that I think that my adult son would derive spiritual/cultural benefits from listening to Mahler symphonies that he cannot obtain from listening to his preferred Green Day, and he thinks otherwise. I'm convinced that his musical preferences are benighted and certain that any efforts to elevate his tastes via rational discussion would be to no avail. Offering him twenty dollars for every Mahler symphony that he listens to needn't be hard paternalism. It depends on whether he would, if rational and well-informed, consent to being offered the money to listen to music that he currently regards as “muzak for snobs.” Since the offer only increases his options, perhaps he would. Though he is now confident that the additional option is of little value, he might wish to leave open the possibility that he will change his mind about this in the future. In that case my offering him the incentives satisfies the hypothetical consent condition and thus counts as soft rather than hard paternalism. But suppose that he (like the religious believer in the precommitment case) is convinced that a

change in musical tastes would be “corruption” or “selling out,” and thus, regards the incentive as a temptation that he prefers not to face. In that case my action toward him would not satisfy the hypothetical consent condition and thus would be hard paternalism.

The important point is that even in the case where it is hard paternalism, it does not violate the autonomy principle because it does not “impose” on him any values that he rejects or “interfere” with his self-determination in any way. The same is true in the case where I attempt to convince someone to abandon his religion and convert to mine, and he regards my pleas as a Siren’s song to abandon the true faith. Of course in that case, unlike monetary incentives to listen to good music, there is no paternalism, because proselytizing uses persuasion rather than any of the nonrational methods in the paternalist’s arsenal. The point, though, is that to offer others either incentives or arguments to do what one thinks is best for them is surely a legitimate exercise of one’s *own* rights, no matter how offensive or objectionable those to whom they’re offered may find them. Assuming that the exercise by one person of her right to self-determination must be compatible with the exercise by others of theirs, it follows that attempting to elicit desired conduct from others in these ways cannot violate their autonomy.

Feinberg and other opponents of hard paternalism would surely accept this point. Hence, their objection is not to hard paternalism *per se*, but rather to a hard paternalism that violates the autonomy principle because it employs such methods as coercion, deception, withholding information to which the person in question has a right, and the like. I’ll call the view that *this* paternalism is *always* wrong “absolutist anti-hard paternalism” (AAHP).

Autonomy vs. Beneficence; When Soft Paternalism Isn't Paternalism

There is a way of characterizing hard paternalism so that it doesn't "really" violate the autonomy principle. It alleges that whenever someone on the basis of a false or unreasonable conception of the good opposes a paternalistic interference with his liberty, he doesn't "really" oppose the interference, because his "true" will is an "ideally rational" will that knows what's good for him, not his "empirical" will which is blinded by false ideologies. The problem with this sort of move is that it misrepresents what is really a denial of autonomy in the name of beneficence as though it were no denial of autonomy at all. It's dishonesty as been noted and protested by Isaiah Berlin, Feinberg, and others.¹⁶ The *honest* hard paternalist admits that what she is doing is overriding autonomy on the basis of (what she thinks is required by) beneficence.

The principle of respect for autonomy requires deference to other people's *past and present* values concerning how they should live their lives. The principle of beneficence, by contrast, requires that we reckon what's good for people in a way that is temporary neutral as regards *present and future*. That is, a good now doesn't count for more than a good in the future just by virtue of its being now. Because of this difference in the temporal focus of the two principles, there will always be the possibility of conflict

¹⁶ Feinberg calls it a "sham and an outrage" (*Harm to Self*, p. 187), while Berlin calls it a "monstrous impersonation" [Isaiah Berlin, "Two Concepts of Liberty," in *Four Essays on Liberty* (1969), p. 134]. A recent paper by Thaddeus M. Pope ["Monstrous Impersonation: A Critique of Consent-Based Justifications for Hard Paternalism," *UMKC Law Review* v. 73, no. 3 (Spring 2005)] argues that any attempt to defend hard paternalism by appeal to consent will disguise the autonomy vs. beneficence conflict that is at issue.

between them no matter which account of prudential value is built into beneficence.¹⁷

That is to say, there is no principle of beneficence that *cannot* conflict with the autonomy principle. For example, a principle of beneficence based on the view that autonomy alone has intrinsic value must support a violation of autonomy whenever it would maximize the person's autonomy in the long run.¹⁸ Even a principle of beneficence based on the subjectivist view that one's good is the fulfillment of whatever conception of the good one happens to hold can support a violation of autonomy in some cases. On this view I can be mistaken about what's best for me in the long run because my conception of the good might change in ways that I fail to anticipate. Insofar as others can predict that I will undergo such changes, they can prevent me from performing acts that will thwart my future conceptions of the good for my own good. If I would not (if well-informed and thinking clearly) now authorize their interference—perhaps I simply don't care about my distant future welfare, or I believe that any future rejection of my current conception of the

¹⁷ David O. Brink examines a similar conflict between prudence and authenticity, which are simply self-regarding or intrapersonal analogues of the duties we have to others to be beneficent and respect their autonomy. See his "Prudence and Authenticity: Intrapersonal Conflicts of Value," *The Philosophical Review* v. 112, no. 2 (April 2003) 215-45.

¹⁸ Perhaps forcing a heroin addict into rehabilitation is an example of a measure that violates present autonomy while maximizing it in the long run. Continued heroin use surely poses a much graver threat to rational agency than continued tobacco use does. It's not clear whether a ban on smoking would increase the typical smoker's autonomy. Insofar as it prevents him from succumbing to weakness of will and allows him to abide by his own judgment that smoking is bad for him, it enhances his autonomy. But insofar as it does this through others' rather than his own agency, it diminishes it. The net effect of a ban on his *level* of autonomy might be a wash. But none of this affects the claim that a ban on smoking *violates* his autonomy. The autonomy principle is a deontological side-constraint, not a principle that tells us to maximize anything, including autonomy.

good would have to be due to corruption or error—then their interference will be hard paternalism motivated by a subjectivist principle of beneficence.¹⁹

Dan Brock has argued that beneficence correctly understood will seldom clash with autonomy. Brock takes issue with Donald VanDeVeer's claim that autonomy and beneficence collide in an example where a Brother Francis, objecting on the basis of moral conviction to the use of rats in a medical experiment that would kill them, volunteers to serve in the experiment in their stead. (This example seems to raise the same questions as my earlier example of Jack's "altruistic" suicide). Brock claims that we should distinguish Brother Francis's "overall interests" (as determined by his "considered preferences") from his "self interest" and that "the condition that the subject will be made significantly worse off by his action should be interpreted in terms of the subject's overall interests, not merely his self-interests."²⁰ Since his moral convictions about the treatment of animals are among Brother Francis's considered preferences, and thwarting his choice would frustrate those preferences, Brock infers that coercive paternalism towards Brother Francis would both violate his autonomy and leave him worse off.²¹

¹⁹ What's wrong with subjectivism so defined, and what is the best account of prudential value—pure hedonism, "informed pleasures," ordinary "self regarding" preference satisfaction, "informed" self-regarding preference satisfaction, an "objective list" or "perfectionist" account of the good, etc.—is an issue not addressed in this essay. A defense of hard paternalism, I think, can be largely neutral on this question. What needs to be avoided is the all too common assumption that the practice of hard paternalism presupposes an objectivist or perfectionist theory of prudential value.

²⁰ Brock, "Paternalism and Autonomy," p. 555.

²¹ Brock does not deny that conflicts between autonomy and beneficence are possible, or that beneficence sometimes overrides respect for autonomy when they conflict. His example of Adam, an impetuous chooser who affirms his impetuosity as valuable part of his character and has impetuously chosen to put himself in significant danger (pp. 562-5) is supposed to be a case where beneficence overrides autonomy. Brock's thesis is that "autonomy trumping rights views like Feinberg's and VandeVeer's

The problem with Brock's view is that it is Brother Francis's self-interest and not his "overall" interests that the principle of beneficence tells us to attend to. A subjectivist principle of beneficence enjoins us to promote his *prudential* good (as he conceives it). It requires that we promote either his own conception of what is in his self-interest or the subset of his considered preferences that concerns his life (i.e. his "self-regarding" preferences).²²

If this is right, then the definition of soft paternalism based on the conditions of prior or hypothetical consent is incomplete, because it creates the possibility that soft paternalism will not be paternalism at all. To see this, consider the example of Jack's retributive suicide. Suppose that Jack fails to kill himself only because of what he regards as an irrational lack of resolve. Suppose that I agree with Jack that his death would be bad for him but that he deserves it as punishment for infidelity to his wife. Accordingly, I kill him, ignoring his pleas (stemming from that irrational lack of resolve) not to do so. Since the hypothetical consent condition is fulfilled here, this counts as an instance of "soft paternalism" under our definition. But clearly it is not paternalism at all, because my act is motivated by my belief that he deserves to die, not a belief that he is better off dead. To avoid the absurdity that soft paternalism might not be paternalism, the definition of it has to be strengthened. It is necessary not only that a prior or hypothetical consent condition

do not have significantly different implications for paternalism than does a plausible interpretation of the promotion of the good justification." (p. 551)

²² This same objection applies to the "single party welfarist consequentialism" that Arneson appears to defend in "Paternalism, Utility, and Fairness," pp. 102-105. In claiming that subjectivist theories of prudential value should identify a person's good not with all of his preferences but with the subset of them that includes only "self-regarding" preferences, I don't mean to imply that there is a clear, intuitively plausible criterion of whether a preference is self-regarding. There may be no such a criterion.

be satisfied, but also that the target of the paternalism would agree with the agent of it that the interference is justified because it leaves the target better off.

There are three possible views about what should be done when autonomy and beneficence conflict. The first is that of AAHP, namely, that autonomy always trumps beneficence; the autonomy principle is absolute, inviolable. Feinberg tries to defend the absoluteness of the right to self-determination by pressing an analogy between it and a nation's right to sovereignty. He says:

If we take the model of national sovereignty seriously, we cannot make certain kinds of compromises with paternalism. We cannot say, for example, that interference with the relatively trivial self-regarding choices involves only "minor forfeitures" of sovereignty whereas interference with the basic life-choices involves the virtual abandonment of sovereignty, for sovereignty is an all or nothing concept; one is entitled to absolute control of whatever is within one's domain however trivial it may be. In the political model, a nation's sovereignty is equally infringed by a single foreign fishing boat in its territorial waters as by a squadron of jet fighters flying over its capital city.... Only a nation's own sovereignty (in the guise, say, of "self defense") may ever be placed on the scales and weighed against another nation's acknowledged sovereignty, for sovereignty decisively outweighs every other kind of reason for intervention.²³

One problem with the analogy has been noted by Richard Arneson: if taken seriously it supports a strict policy of no paternalism whatsoever, not Feinberg's own position, which strictly forbids only hard paternalism. After all, if the government of a country enacts tax policies that will achieve exactly the opposite results that it seeks, because it holds beliefs about fiscal matters that are demonstrably false, another country may not intercede to force it to adopt tax policies that will better serve its own goals.²⁴

The second view is that since the right to self-determination has value only as a *means* to promoting one's own good, beneficence always overrides autonomy when the

²³ Feinberg, *Harm to Self*, p. 55.

two conflict. This is a possible hard paternalist position, but not the only one. The third view is that the principle of respect for autonomy is as basic or nonderivative as the principle of beneficence, and sometimes beneficence will override autonomy while other times autonomy will override beneficence. As noted earlier, this is the view that is implicit in commonsense morality.²⁵ According to this view, both principles can be likened to Rossian “prima facie” duties. It is the other form that the hard paternalist view can take, and it is the one I shall adopt.

The Libertarian Objection to Soft Paternalist Laws

Is soft paternalism ever wrong? We should distinguish cases where a friend engages in soft paternalism towards me from cases where the state does so via a (coercive) law that applies to a large class of which I am one member. We’ll consider the person to person case later. For now we should take note of why the libertarian is opposed to nearly all coercive soft paternalist legislation.

The Missouri law barring people who put themselves on a list from entering any of its casinos is an example of soft paternalism. Aside from worries (surely unwarranted) that it will lead *via* a slippery slope to tyranny, there is no reason why the libertarian should condemn it. After all, that law applies to a class of persons for whom the prior authorization condition is satisfied in every single case. If a law had to satisfy a unanimity requirement to count as soft paternalism (i.e. the prior authorization or hypothetical

²⁴ See Arneson’s “Paternalism, Utility, and Fairness,” p. 110.

²⁵ It is the “fourth” position described by Feinberg in *Harm to Self*, p. 61.

consent condition had to be satisfied for every single person to whom the law applies), then very few laws would ever qualify. Mandatory seat belts laws should count as soft paternalism on the plausible assumption that *most* drivers would regard them as taking out a “social insurance policy” (to use Gerald Dworkin’s felicitous phrase) against their own proclivities to laziness, forgetfulness, impatience, etc.²⁶

Consider another example. Suppose that the majority fear that they are prone to excessive alcohol consumption, and so, to forestall that possibility, they support a ban on alcohol sales. Such a ban should qualify as soft paternalism even though it would have to apply to those in the minority who want the opportunity to imbibe spirits and are confident that they can do so in moderation. Why would it have to apply to them too? The answer is that it simply is not feasible to enforce a ban on alcohol consumption that applies only to “likely reckless drinkers” and not “responsible, moderate drinkers” as well. Such a law would be an administrative nightmare, requiring investigations into the life history of each individual drinker to determine whether he is in violation of it. The purpose of a ban that applies to everyone would be to restrict the liberty of the intemperate drinkers only; the liberty of the temperate drinkers is restricted only as an unintended but unavoidable side-effect of having enforceable soft paternalist restrictions for the good of the intemperate ones. To those temperate drinkers whose liberty is also restricted, we say (as Feinberg

²⁶ Gerald Dworkin, “Paternalism,” p. 75. Again, though, the fact that they suffer from weakness of will is not enough to make the law soft paternalism. The reason why the hypothetical consent condition is satisfied in their case is that they do not regard the threat of a fine for not buckling up in the same way that most smokers would regard a ban on smoking—that is, as very onerous.

does in connection with Social Security) that “the compulsion is for their sakes, not yours.”²⁷

Of course, this may not be much consolation to them. They may well resent having *their* liberty restricted because *others* cannot exercise *their* liberty responsibly. The libertarian holds that such resentment is justified. She insists that it is unjust to restrict the liberty of wise choosers as part of the means necessary to protect foolish choosers from themselves.²⁸

A possible reply to the libertarian is that the sort of justification for coercion to which she objects in the paternalistic case seems perfectly acceptable when it is used in nonpaternalistic cases. Consider, for example, speeding laws, which I shall assume exist primarily to prevent harm to people other than the speeding driver. Suppose that the average driver with average reflexes poses an unacceptably high risk of harm to others when he exceeds 55 mph on a certain stretch of road, while a skilled driver with quick reflexes does not pose a comparable risk until he exceeds 75 mph on the same road. The speed limit for that road should be set with the average driver in mind at 55 mph—to the detriment of the gifted driver. If it’s permissible to enact laws that “penalize” the more gifted in the nonpaternalistic case, why shouldn’t it be permissible in the paternalistic case as well?

²⁷ Feinberg, *Harm to Self*, p. 18.

²⁸ Arneson notes that if we follow the libertarian and reject such restrictions, then it’s likely that we will exacerbate inequality, since wise choosers are likely to be already better off than foolish ones. He argues that that will be objectionable on the plausible assumption that people are wise or foolish choosers largely owing to genetic and environmental (early upbringing and socialization) factors for which they are not responsible. (Though he doesn’t mention them, Walter Mischel’s “marshmallow” studies on 4 year olds might be evidence that the assumption is true). See “Paternalism, Utility,

Another Objection to Some Soft Paternalism

With person-to-person soft paternalism the libertarian objection does not apply, because the only person affected by it gives prior or concurrent hypothetical consent to it. Assume that my friend knows my character and values intimately and engages in an act of soft paternalism towards me knowing that the hypothetical consent condition is satisfied. His act might be thought wrong because: i) I *deserve* to suffer the consequences of my imprudence, which is a moral vice, or ii) contrary to what my friend and I think, his act really harms me. i) implies that a moral vice that is self-regarding (harming no one other than the agent) may nonetheless merit punishment; soft paternalism is sometimes wrong because it insulates a morally defective character from its just deserts. I'm skeptical of such an argument and shall not pursue it. But there are some conceivable acts of person-to-person soft paternalism that are wrong for the reason stated in ii).

Suppose that I'm a very impetuous person who frequently regrets his actions, and I authorize my guardian angel to step in and thwart my choices whenever he foresees that I will later regret them. He protects me not just from my extremely bad choices, but from my minor mistakes as well. This keeps him very busy, since I make several such mistakes every day. I'm frequently angry with him for blocking these choices, but knowing that he helps and never hinders my pursuit of my own deepest values and preferences, I never rescind his power to interfere. Isn't the long term effect of his guardianship on my character likely to be negative? Why should I ever bother to deliberate before choosing if I know that my guardian angel will veto any choice I would later regret? Certainly he

and Fairness," pp. 86-88.

eliminates an incentive I might have to cultivate such traits as self-control, patience, discriminating judgment, and the like. Of course with a guardian angel by my side it might be claimed that I have no need for such traits. But the reply to that claim is that the traits have intrinsic and not merely instrumental value, their possession and exercise being necessary to lead a flourishing life. Certain perfectionist theories of prudential value support that judgment; so too might the “qualitative” hedonism defended by J.S. Mill in *Utilitarianism* (supposing that possession of the traits is one of the things that distinguish the dissatisfied Socrates from the satisfied fool). A principle of beneficence based on such theories implies that any paternalism likely to discourage the cultivation of these traits is wrong, even if it is soft.

The Refutation of AAHP

We can *imagine* a case in which beneficence clearly trumps autonomy, thereby demonstrating that AAHP is false. Suppose that the state could know with near certainty that X is the true religion. X is theologically intolerant, implying that only those who accept it are saved; everyone else is eternally damned. Salvation for each of us consists in living an eternally long life each stage of which is excellent; eternal damnation consists in a life of unending, excruciating torture. Finally, suppose that there are mildly coercive measures that will increase the likelihood of citizens rejecting false faiths and embracing X. If AAHP were true, then it would be wrong for the state to institute these measures no matter how mild they may be for its citizens’ own good. This, I claim, is absurd and demonstrates the absurdity of AAHP. Of course we should oppose any effort by the state

to limit our religious freedom for our own good. The suggestion that the state could ever know with near certainty that some theologically intolerant religion is true is really quite absurd. But to say this is to doubt that *beneficence* supports the measures.²⁹ What AAHP implies is that because autonomy *always* trumps beneficence, the state may not threaten me with small fines *even if* it knew that such action would probably provide me infinite benefits and/or protect me from infinite evil. That's not proper solicitude for the right to self-determination. It's "fanaticism."³⁰

In the face of counterexamples like this one, opponents of coercive hard paternalism have to abandon a dogmatic, uncompromising AAHP based on an absolutist principle of autonomy in favor of something more pragmatic. The claim will have to be that while beneficence might override autonomy in some *conceivable* cases, it never does "in the actual world" or under "realistic" assumptions. Even if, strictly speaking, AAHP is false, much less moral mischief will be done if we act as though it were true than if we stand ready and willing to restrict liberty on hard paternalist grounds. Call this "pragmatic AHP."

²⁹ In "Neutrality, Skepticism, and the Fanatic" [*Social Theory and Practice* (January 2006): 35-60], I argue that to rebut this argument for intolerance one has to assume either religious skepticism or a dogmatism that claims that theologically and politically intolerant religions are false. Since neither of those doctrines is "neutral," I infer that the liberal position on freedom of religious worship cannot be neutrally justified

³⁰ Such is Richard J. Arneson's verdict on Feinberg's "absolutist antipaternalism" in his "Joel Feinberg and the Justification of Hard Paternalism," forthcoming in *Legal Theory*.

The Limits of Soft Paternalism: Legalized PAS and a Ban on Prostitution

Feinberg's primary strategy for opposing hard paternalism as a "liberty limiting principle" that legitimates criminal legislation was to argue that several examples of acceptable laws that appear to require a hard paternalist justification in fact do not. It turns out, for example, that there is either a soft paternalist or nonpaternalist rationale for a ban on dueling, anti-usury laws, Social Security, and the criminal law's refusal to accept the consent of the victim as a justification for homicide. I agree with all this and would add to the list the law's refusal to enforce "slavery contracts" voluntarily entered into by both parties. Though given a paternalistic justification by J.S. Mill in *On Liberty*, it probably has a sound nonpaternalistic rationale. Such contracts, like employment contracts that forbid departing employees to work for a rival company for the rest of their lives, would make labor markets less fluid and competitive, thus reducing economic efficiency. Society has an interest in maintaining the inalienability of the right to one's labor. Finally, there's an obvious soft paternalist rationale for not permitting suicide assistance by anyone to anyone who wants it. The overwhelming majority of suicides in the real world, especially by those who aren't terminally ill, are due to mental illness or extreme emotional agitation and are not rational, voluntary, or autonomous. Such people, if rational and well-informed, would probably consent to a law forbidding others to assist their suicide attempts.³¹

³¹ Feinberg denies that the majority of suicide attempts are nonvoluntary. He says "In respect to suicide and other controversial self-regarding acts, there is commonly no reliable presumption at all [about their voluntariness] of a statistical sort. Where interference, detention, and inquiry into voluntariness are justified, they are so because of

But the issue of suicide assistance, I think, creates a problem for pragmatic AHP. Conservatives oppose the legalization of physician assisted suicide (PAS) altogether (on “sanctity of life” grounds and/or because of worries that there is no feasible way to prevent rampant abuse), while libertarians insist that suicide assistance ought to be available to any competent adult who wants it. Many people, myself included, would like to support the standard (nonlibertarian) liberal view on PAS, namely, that it ought to be legally available *only to the terminally ill*. All competent adults who are not terminally ill—and they include people who are perfectly healthy and able, at one extreme, and people with severe disabilities (like ALS), at the other— should be ineligible for PAS. To take an extreme case, consider Frank, who suffered a shoulder injury a year ago that will prevent him from ever again playing competitive golf, his life’s passion. He is no longer depressed about his situation, but he is certain that he has nothing to live for and would be better off dead. Thus, he seeks help to end his life. The nonlibertarian liberal position is that people like Frank should be ineligible for suicide assistance. I do not see any way to defend this position by means of soft paternalism. It requires a hard paternalist defense: the reason why those who are not terminally ill should be ineligible for PAS is that the

the magnitude of the risks and the gravity or irrevocability of the harms, not because of the initial probability of nonvoluntariness.” (*Harm to Self*, p. 127) Is this consistent with a soft paternalist justification for a policy of “detention and inquiry” for suicide attempts? I think it is. Suppose that only 10% of all suicide attempts were nonvoluntary. Feinberg could say that interference with all attempts is justified “for the sake of” that 10%. A “detention and inquiry” policy promotes the autonomy and welfare of that 10% greatly, while only delaying the other 90% a little in their pursuit of their goals. I see no trace of hard paternalism in this defense of the policy, but am puzzled by Feinberg’s unwillingness to grant that the majority of suicide attempts are nonvoluntary.

majority of them are better off alive than dead; if they think otherwise, then they are simply mistaken.³²

Won't the same soft paternalist rationale for not permitting suicide assistance by anyone to anyone support the limit on PAS eligibility that the nonlibertarian liberal favors? The answer is "no," for the following reason. The soft paternalist has to oppose a law that allows anyone to kill anyone else who request suicide assistance, because it's predictable that the hypothetical consent condition will not be satisfied in the majority of such cases. But the way to guarantee that all who seek PAS satisfy that condition is not to limit eligibility to the terminally ill. Rather, it is to require a screening process that filters out the clinically depressed, emotionally agitated, and so forth. Whether we limit PAS to the terminally ill or not, we would have to require such screenings in order to ensure that those who are supposed to be ineligible for PAS (because they are not terminally ill and/or not making a free and informed decision to opt for it) are not provided it. Such a screening process might be expensive, but surely it is administratively feasible in a way that restricting alcohol consumption only to responsible, moderate drinkers would not be. Granting eligibility for PAS to those who are not terminally ill would increase the work

³² To admit that it is hard paternalism is to admit that there is a violation of autonomy if people like Frank are prohibited from soliciting and accepting suicide assistance from others willing to provide it. I agree with Daniel Callahan ("Self-Determination Run Amok" *Hastings Center Report*, March-April 1992, 52-55) that *if* the only argument for legalizing PAS were one based *solely* on the autonomy principle, *then* there would be no justification for limiting eligibility to the terminally ill. But I claim that a hard paternalist defense of the limit is possible. Felicia Ackerman [in "Assisted Suicide, Terminal Illness, Severe Disability, and the Double Standard" in *Physician-Assisted Suicide: Expanding the Debate*, ed. by M. Pabst Battin (New York: Routledge, 1998)] has argued that a quality of life ethic is no more able to justify limiting PAS to the terminally ill than the autonomy principle can. I attempt to answer her claim that such a limit can only reflect anti-disability bias in my paper, "The Elitism Objection to a Quality of Life

load of the review boards charged with monitoring the entire process, but perhaps not by very much. Hence, soft paternalism implies that people like Frank should be eligible for PAS after their decisions to end their lives have been vetted by these boards to ensure that they are substantially voluntary and in accord with their deepest values.³³ But that is not the nonlibertarian liberal's view. Her view is that people like Frank should be ineligible for PAS, period.

Hard paternalism is not the only way to defend a denial of suicide eligibility to people like Frank. J. David Velleman has argued that even if, in cases like Frank's, one is better off dead, suicide and assisted suicide remain wrong because they deny the "interest-independent" value ("dignity") that persons have.³⁴ Dignity is an unconditional value, the very condition of one's prudential good (a merely conditional value) mattering at all. If hard paternalism holds that beneficence overrides autonomy in cases like this one, and beneficence involves concern for Frank's prudential good, then the Kantian objection to Frank's assisted suicide is not based on hard paternalism, because it holds that it is Frank's dignity rather than his prudential good that trumps his autonomy.³⁵

Velleman's Kantianism, unlike the sanctity of life ethic, does not oppose suicide assistance for the terminally ill. As a terminal illness progresses, it is usually

Ethic" (in progress).

³³ Surely this is the policy on assisted suicide that is in the same spirit as the "detention and inquiry" policy for unassisted suicide attempts that Feinberg favors.

³⁴ J. David Velleman, "A Right to Self-Termination?" *Ethics* v. 109 no. 3 (April 1999): 606-28.

³⁵ The distinction we drew earlier between the honest and dishonest hard paternalist is applicable here. The honest Kantian admits that a legal prohibition on suicide assistance to people like Frank violates the autonomy principle. The dishonest Kantian claims that since Frank's suicide is contrary to the Categorical Imperative, and Frank's true, rational will is to obey the Categorical Imperative, that prohibition doesn't "really"

accompanied by a deterioration of those faculties of free will and practical reasonableness by virtue of which one is a “person” with “dignity.” Velleman says that “dignity can require not only the preservation of what possesses it but also the destruction of what is losing it, if the loss would be irretrievable.”³⁶ Velleman’s Kantianism implies that PAS for the terminally ill is permissible—not because they are better off dead, but because they are approaching the condition of “undignified life.”

I do not find the Kantian rationale for limiting PAS eligibility to the terminally ill more persuasive than the hard paternalist one, for several reasons. First, if personhood were indeed the condition necessary for one’s prudential good to matter at all, then the prudential good of dogs and pigs would not matter at all; but it does. Second, it seems to me inconsistent to say, on the one hand, that my prudential good matters (because I am a person), but on the other hand that there are cases where my suicide leaves me better off and no one else is worse off yet it remains wrong. In *those* cases, whatever they are, my prudential good does *not* matter. Third, if the capacities by virtue of which one has dignity really did have a value “beyond all price,” then it would be wrong to impair or abandon them even *temporarily* for the sake of the greatest prudential good. That doesn’t seem right.³⁷ Finally, Velleman claims that dignity *requires* the destruction of what is losing it. Consider a PVS patient whose advance directive said that he wants “ordinary”

violate his autonomy.

³⁶ Velleman, “A Right to Self Termination?” p. 617.

³⁷ Imagine someone who roughly once a month gets very drunk—to the point that he babbles incoherently and cannot think straight. With drunkenness he temporarily anesthetizes his capacity for free will/practical reasonableness, and he does so because it makes him feel good, providing him a short reprieve from the pressures of his job. It seems excessively puritanical to condemn this person’s occasional “bender.” Another example, this one nonhedonistic, is “Schelling’s reply to armed robbery” as described by

care continued should he ever become PVS, because of his belief in the sanctity of life doctrine as the Roman Catholic Church understands and defends it. This doctrine implies that it is mere biological human life, not Kantian personhood, that is the ground of our “dignity.” Since PVS patients have irretrievably lost their capacities for personhood, Velleman’s Kantianism apparently requires that we disregard this patient’s advance directive. I find this implication of his view deeply implausible. I conclude that a hard paternalist justification for limiting PAS eligibility to the terminally ill is superior to the Kantian one offered by Velleman.

Another issue on which soft paternalism probably fails to support the position held by many liberals who are not libertarians is prostitution. Many such liberals, myself included, are sympathetic to its criminalization for paternalistic reasons: it is a “lifestyle” that is not good for the women who practice it. The dangers that a ban seeks to protect these women from are not just physical ones to life and health but “ethical” ones to spirit and character. Nonlibertarian liberals are usually loath to acknowledge that protecting someone from “corruption” might be a good reason for restricting his liberty. Antipathy to hard paternalism leads most to try to defend a ban on soft paternalist grounds: the choice of prostitution as a trade is substantially nonvoluntary, necessitated, for example, by the need to support a drug habit. In a recent paper on paternalism, Peter de Marneffe has argued that this picture does not fit the facts. There is abundant evidence that most women in “sex trades” choose their line of work voluntarily.³⁸ If that’s correct, then

Derek Parfit, *Reasons and Persons* (New York: Oxford University Press, 1984), p. 12.

³⁸ Peter de Marneffe, “Avoiding Paternalism,” *Philosophy and Public Affairs* v. 32, no. 5 (2005) 68-94. de Marneffe’s references to the literature on prostitution are on p. 93.

nonlibertarian liberals either have to admit that a ban is justified for reasons of hard paternalism or agree with libertarians that it is not justified at all.

Limiting Hard Paternalism

If people like Frank are to be denied PAS eligibility because their views about whether their lives are worth living are “simply mistaken,” and commercial sex is to be banned because it is bad for the people who engage in it, where do we stop? Why not sometimes override the decisions of competent adults to reject unwanted medical care? Surely there are cases in which people reject life saving treatment for bad reasons (e.g. a refusal to consent to an appendectomy because one believes that wearing a crystal of the right color will cure one’s appendicitis).³⁹ Should we override the adult Jehovah’s Witness’s refusal to consent to a life saving blood transfusion? Why not a ban on “fornication” aimed at individuals who seek nothing more than a “one night stand”? And so on. Pragmatic AHP alleges that trying to determine where to draw the line between permissible and impermissible hard paternalism is probably hopeless and certainly dangerous. It’s safer to suppose that in the real world beneficence *never* overrides the duty to respect autonomy or personal sovereignty.

Of course if what I’ve argued so far is correct, then accepting pragmatic AHP itself has costs: we have to either reject both the “legalization with limits” of PAS and a ban on

³⁹ Assume that the belief is held in a manner that makes it immune to possible falsification (e.g. in conjunction with the *ad hoc* auxiliary assumption that the crystal will be efficacious only when worn by someone with the proper, spiritually attuned mindset). The belief is part of a New Age metaphysics. However irrational it may be, a choice

prostitution, or continue to support them with feeble and disingenuous soft paternalist arguments. If both of those alternatives are unacceptable, then we are stuck with hard paternalism and the slippery slope worry about how and where to limit it. I wish to suggest a couple of strategies that the defender of hard paternalism might employ in an effort to allay that worry.

We should be clear about which cases are and which are not on the slope. Our question is how to determine whether beneficence trumps autonomy or *vice versa* in cases where *the two are in conflict*. Suppose that a PVS patient has left behind an advance directive stipulating that should he ever become PVS, his wish is that “ordinary care” (feeding tubes, etc.) be discontinued. His caregivers disregard those wishes on the grounds that he is better off with mere biological life than he would be dead. This would be an act of hard paternalism in connection with which the question does *not* arise, because it’s not a case plausibly regarded as presenting a conflict between beneficence and autonomy. The caregivers’ notions of what’s in their patient’s best interests seem mistaken. Incidentally, though there probably are people who oppose pulling the plug on PVS patients on the basis of beneficence, those who oppose it on “sanctity of life” grounds need not be among them. “Sanctity of life” is best construed as a deontological side-constraint on the permissibility of killing, not as a prudential value judgment about mere biological life.

The first thing that the hard paternalist might say about cases in which the two values *are* in conflict is this. It has to be admitted that there are cases in which hard paternalism would protect someone from serious harm, yet it remains clearly wrong.

based on it remains fully voluntary.

Forcing a life saving blood transfusion on an adult Jehovah's Witness is such a case.⁴⁰

The hard paternalist can say that the reason why it is clearly wrong in these cases is not solely or even primarily due to its being a violation of autonomy; it is due to its violating some *other* right. These other rights include the rights to free speech, to religious practice, to form one's own personal relationships, to be free from physical assaults on one's body, and not to be killed without one's consent. The idea here is that these more specific and substantive rights are *not* derived from or dependent on a more fundamental right to self-determination, but instead are to be justified in some other way. (If anything, the right to self-determination is obtained *via* generalization from these specific rights). The reason why hard paternalism is clearly wrong in the adult Jehovah's Witness case is not that the right to self-determination trumps beneficence, but that the rights to freedom of religion and to be free from unauthorized physical assaults on one's body trump it. By contrast, denying PAS eligibility to people like Frank does not violate any other right besides the right to self-determination.⁴¹ The hard paternalist needn't deny that a right to self-determination exists and covers all of our self-regarding choices, including the ones not

⁴⁰ Some have claimed that beneficence properly understood really does not call for interference here, on the grounds that while a forced transfusion might save the JW's life, it would do so at the cost of leaving him miserable because convinced that he is eternally damned. Jehovah's Witness teachings, however, do not support that claim. JW's insist that blood transfusions are unsafe and almost always medically unnecessary, but they seem unanimous that anyone who is forced to undergo one is faultless and suffers no spiritually adverse consequences. The Jehovah's Witness who "strenuously" objects to a blood transfusion but it given one anyway is no more guilty of sin than a woman who "strenuously" objects to being raped is guilty of fornication (*Watchtower*, 1991, 6/15 p. 17).

⁴¹ Surely the "right to die" is an offspring of the right to self-determination and entirely dependent on it for its justification. A "right to die" should not be confused with a right not have to suffer physical pain, mental anguish, and slow deterioration of mental faculties before succumbing to death caused by terminal illness.

covered by any of these other, more specific rights. But she can say that if it does exist it is less weighty than these other rights and by itself not weighty enough to trump the dictates of beneficence in cases where paternalism would benefit the subject greatly.

In the case of prostitution, the only plausible candidate I can see for another right (besides the “right to sexual autonomy”) that a ban would violate is the right to control one’s own labor power. But if OSHA regulations are acceptable in spite of the fact that they limit that right, it’s hard to see why a ban on prostitution should be unacceptable for that reason. A right to enter into the personal relationships of one’s own choosing is vague, but I assume that there is a way to explicate the notion of “personal relationships” so that it does not cover sex-for-money exchanges or any other commercial transactions. If “fornication” is understood as sexual relations between unmarried adults, then a ban on fornication would violate the right to choose one’s own personal relationships. Admittedly, “one night stands” in which each person regards the other as little more than a means to sexual gratification probably shouldn’t qualify as “personal relationships” any more than sex-for-money exchanges should. But the hard paternalist can plausibly argue that whatever harm people who pursue them do to themselves is likely to be much smaller than the harms suffered by prostitutes, and furthermore, there is no feasible way to prohibit mere “one night stands” without also interfering with acts that *are* fully protected by the right to form one’s own personal relationships.⁴²

⁴² Interpreting the Constitution’s “right to privacy” so that it protects “personal relationships” rather than “sexual autonomy” allows us to judge anti-sodomy laws unconstitutional without having to do the same for bans on prostitution. See Michael J. Sandel, “Moral Argument and Liberal Toleration: Abortion and Homosexuality,” 77 *California Law Review* 77 (1989). For a defense of a right to privacy that protects prostitution, see David A.J. Richards, *Toleration and the Constitution* (New York: Oxford

I've suggested that hard paternalism will be easier to justify in cases where the *only* right that it violates is an abstract "right to autonomy" and not another more specific, substantive right too. Fully developing that suggestion would require that I furnish a developed account of what our rights are, which I don't have. My second suggestion is that even in cases where self-determination is the only moral right at stake, some violations of it are more severe than others, and hard paternalism will be harder to justify the more severe is the violation of autonomy that it requires. Everyone will admit, I think, that forcing someone to forsake his religion for his own good is a more serious violation of his autonomy than fining him for not having a life vest on his sailboat. It is a more severe violation simply because he is likely to value practicing his religion more than sailing without life vests. What I want to suggest is that we can plausibly judge a violation of autonomy to be less serious when it's likely that there would be retrospective authorization after the violation and/or regret that no violation occurred if one didn't. Suppose, for example, that the majority of smokers have the values described earlier, such that forcing them to quit for their own good is hard paternalism. If it's likely that i) most of them in twenty or so years would thank us if we forced them to quit, and/or ii) most of them in twenty or so years would wish that we had ignored their earlier wishes and forced them to quit, if we hadn't, then that would make a ban on smoking a less serious violation of their autonomy. For the purpose of determining whether the autonomy principle has been violated, we should still attend only to the person's present and past values. (I'm not suggesting that if there's subsequent, retroactive approval of the intervention, then there was *no* violation of autonomy). But for the purpose of determining *how grave* the

violation is, we should adopt a more temporally neutral perspective that takes into account likely future values as well. When “you’ll thank me later” turns out to be true, it diminishes the severity of the violation of autonomy.

If this suggestion is right, it provides another reason why the hard paternalist can disapprove of coerced blood transfusions for Jehovah’s Witnesses, and thus, avoid too precipitous a slide down the slope. Not only would forced transfusions violate rights to freedom of religion and freedom from unwanted bodily intrusions, but they would also be severe violations of autonomy. The violations are likely to be severe, because it is unlikely that Jehovah’s Witnesses would later come to approve of them if their wishes were ignored. The case of Frank is probably different in this regard. People who’ve suffered disabilities that shatter their life’s dreams but are no obstacle to engaging in a wide range of other worthwhile and rewarding life activities usually adapt to their new circumstances with time. If he is patient, Frank is likely to find new challenges that he wants to tackle, in which case he will be grateful that his request for suicide assistance was denied.⁴³ It might be objected that if this is true, then a rational Frank who knew this about himself would consent to being denied PAS eligibility, giving us a soft paternalist rationale for the PAS limit and making a hard paternalist one otiose. But the problem is

⁴³ The Philosopher’s Brief submitted to the Supreme Court in the case of *Washington v. Glucksberg* (1997) makes the same point. See Ronald Dworkin, Thomas Nagel, Robert Nozick, John Rawls, Thomas Scanlon, and Judith J. Thomson, “Assisted Suicide: The Philosopher’s Brief,” *New York Review of Books*, March 27 1997, pp. 41-47. What’s not so clear is whether the authors of the brief are assuming that there is no violation of autonomy at all whenever there is retrospective authorization, or whether they are assuming that such authorization instead allows beneficence to override autonomy. Both alternatives seem inconsistent with the position on paternalism that Ronald Dworkin defends in *Sovereign Virtue*, Chapter Six.

that a Frank who can foresee such a change in his future self might be repulsed rather than gratified it. He might regard adapting to a life without golf as corruption or “selling out.”

If most smokers in their later years wished that they had been forced to quit when younger, that would make the violation of autonomy from a smoking ban less egregious than coerced blood transfusions for Jehovah’s Witnesses. But that does not mean that we should support a hard paternalist ban on smoking. A ban would create all of the problems that Prohibition did (e.g. criminal syndicates that thrive by selling the contraband on the black market), and the social costs of trying to enforce it would probably outweigh whatever good it does for smokers. But on the issue of whether or not a ban would be wrong *even if* it were perfectly and costlessly enforceable, I confess to having unclear intuitions. It’s not clear to me whether we should think that there is another more substantial right at stake here in addition to the thin, abstract right to self-determination.

Conclusion

Of course some opponents of a criminal ban on prostitution may wish to argue that difficulties in enforcement and negative side-effects make that ban unwise too. A ban doesn’t prevent prostitution; it merely pushes it underground and forces prostitutes to rely on exploitative and often violent pimps for protection from abusive customers. Legalization would not only contribute to the physical well-being of prostitutes but also abate those of its morally corrupting effects that are due to its being part of the criminal underworld.

I find this argument unpersuasive because I'm skeptical of its empirical assumptions. But if those assumptions are correct and the argument is sound, then what follows is that an enlightened beneficence will not support a ban on prostitution. It does not follow that a ban would be wrong because autonomy trumps beneficence. I do not claim that we should ban prostitution no matter what the consequences of doing so may be on welfare. That view is more likely to be held by the "pure legal moralist" than the paternalist.⁴⁴ My claim is rather that we should not think that its violating anyone's right to self-determination is an especially strong reason not to ban it. A similar point applies to the "legalization with limits" of PAS. If legalization harmed many more people than it benefited, either because conservative worries about the inevitability of rampant abuse turn out to be justified, or because (as some disability advocates contend) it would lead to reduced hospice and medical care for the terminally ill with no wish to opt for PAS, then I agree that it should not be legal for anyone. My claim is merely that it's not a very weighty reason against limiting eligibility to the terminally ill that doing so would violate the autonomy of people like Frank. Hard paternalism may be ill-advised in both the prostitution and PAS cases for other reasons, but we should not oppose it on the grounds that autonomy always trumps beneficence.

⁴⁴ See Joel Feinberg, *Harmless Wrongdoing* (New York: Oxford University Press, 1990).