Axes of Paternalism

Richard H. Dees*, Department of Philosophy, University of Rochester, Rochester, NY, USA

*Corresponding author. Richard H. Dees, Department of Philosophy, University of Rochester, P.O. Box 270078, Rochester, NY 14527-0078, USA. Tel.: 585 275 4105; Email: richard.dees@rochester.edu

Paternalistic actions interfere with the autonomy of an agent, so they can only be justified by strong moral considerations. Yet how seriously an action infringes on autonomy depends on how it affects autonomy. Autonomous actions require at least three independent elements: agents must act with an understanding of the risks of their actions, without the coercive influence of others and in accord with their most important values. Each of these elements exists along a spectrum, and by mapping measures in the resulting three-dimensional space, we get a better understanding of what is at stake in any debate about them, and we can understand when paternalistic measures might be justified and when they are not.

John Stuart Mill famously argues that paternalism is always wrong, that

> the sole end for which mankind are warranted individually or collectively, in interfering with the liberty of action of any of their number is self-protection. That only purpose…is to prevent harm to others. His own good, either physical or moral, is not a sufficient warrant (Mill, 1859: 1.9).

Despite this apparently straightforward statement, what exactly Mill means and what exactly he finds objectionable are not so clear. In his trenchant analysis, Gerald Dworkin focuses on the ways paternalism interferes with people's freedom:

> *X acts paternalistically towards Y by doing (or omitting) Z:*
> 1. *Z* (or its omission) interferes with the liberty or autonomy of *Y*.
> 2. *X* does so without the consent of *Y*.
> 3. *X* does so because *X* believes *Z* will improve the welfare of *Y*…or in some way promotes the interests, values, or good of *Y* (Dworkin, 2020).

The problem with paternalism, then, is that it restricts freedom. Some, however, point to other aspects of Mill's view as the deeper problem with paternalism. On this view, autonomy is the key value, and we should always seek to allow people as much room as possible to express their own individuality, to engage in what Mill calls "experiments in living" (Mill, 1859: 3.1), as long as they do no harm to others. To interfere with those choices, he thinks, treats adults as if they were children (Mill, 1859: 5.10). For that reason, Seana Shiffrin and others (Shiffrin, 2000; Cholbi, 2017) have argued that the problem with paternalism is not simply that it restricts freedom, but that it substitutes the will of one person for that of another, even when the former is better positioned to make a good judgment. The problem with paternalism, then, is that it restricts a person's rational will.

Yet both the freedom and the rational will accounts of paternalism allow the possibility that paternalism might be justified. While Mill, famously, seems to conclude that paternalism is never warranted, even he hedges.[1] Most obviously, the prohibition of paternalism applies only to competent adults (Mill, 1859: 1.10). After all, it is only "paternalism" if we are acting on behalf of people as if they were children; compelling children to do what is good for them against their will is (usually) just good parenting. And in other respects, Mill is not quite as liberal as we might think: he includes among those adults who are not competent those who live in "backward states of society in which the race itself may be considered in its nonage" (Mill, 1859: 1.10). Yet leaving aside these obvious cases, Mill still hedges on a complete prohibition against paternalism. We can, for example, stop the person who aims to cross a rickety bridge, by force if necessary, at least until we can explain to him the danger of the bridge, "for liberty consists in doing what one desires and he does not desire to fall into the river" (Mill, 1859: 5.5). In addition,

https://doi.org/10.1093/phe/phag008
Online publication date: 15 May 2026

© The Author(s) 2026. Published by Oxford University Press. All rights reserved. For commercial re-use, please contact reprints@oup.com for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site—for further information please contact journals.permissions@oup.com.

Downloaded from https://academic.oup.com/phe/article/19/2/phag008/8679220 by University of Rochester user on 21 May 2026

we can forbid someone from selling himself into slavery: "It is not freedom to be allowed to alienate his freedom" (Mill, 1859: 5.11).

The key questions, then, are about the circumstances in which paternalism might be justified. The presumption is that since paternalism usually curtails our freedom or our rational wills, it could only be justified by very strong moral considerations. My project in this paper is to bring clarity to this discussion by reframing the issues. By thinking about three aspects of autonomy, we can understand better when autonomy is seriously infringed and when it is not and thereby understand better the circumstances in which paternalism might be justified.

Setting the Stage

Much of the discussion about when paternalism can be justified focuses on the difference between "soft" and "hard" paternalism, in which "soft" paternalism is frequently justifiable, and "hard" determinism never is. Unfortunately, the term "soft paternalism" is used in at least two different ways (Hanna, 2018), and even those different ways do not reflect the complexity of the moral landscape. Sarah Conly, for example, uses the term to refer to measures that do not use much force: "Soft paternalists try to influence individuals not to make that self-destructive choice, but leave them the freedom to do the self-destructive thing if they nonetheless want to do that" (Conly, 2013b: 241), citing education, warning labels and laws that limit soda portions in restaurants as examples (Conly, 2013b; see more generally, Conly, 2013a). More often, "soft paternalism" refers to measures that prevent people from making decisions that are not voluntary. As Joel Feinberg puts it, "Soft paternalism holds that the state has the right to prevent self-regarding harmful conduct…*when but only when* that conduct is substantially nonvoluntary, or when temporary intervention is necessary to establish whether it is voluntary or not" (Feinberg, 1986, 12). We have reason to doubt that a person is acting on their own volition if their "choice stems from ignorance, coercion, derangement, drugs, or other voluntariness-vitiating factors" (Feinberg, 1986, 12). Most often, such measures are efforts to ensure that people do not make decisions without understanding the risks they are undertaking. So, we typically do not allow people to skydive, to ride go-carts or to undergo surgery unless they demonstrate at least a minimal understanding of the risks that might be involved. We also disallow contracts that are made under grave threats or under the influence of alcohol. Invoking soft paternalism allows us to explain why we are permitted to intervene in Mill's bridge case: the man does not understand that the bridge is unsafe, and so can stop him—by force, if necessary—until we can be sure that he knows what he is doing.

In addition, some people think that paternalism might be justified as long as it does not violate what Thomas Nys calls "broad autonomy," which "means that one is allowed to make important life choices, that one exercises control over one's life, that one is the author of one's own story" (Nys, 2008: 67). Unless a measure affects significant portions of our life, he suggests, it need not concern us. In many respects, Mill is most interested in broad autonomy. His emphasis on "experiments in living" and on giving "free scope…to varieties of characters" (Mill, 1859: 3.1) speaks to the importance of having broad control over one's life. But he realizes that broad autonomy requires an ability to make lots of little decisions every day, most of which are not essential to living a particular kind of life, but all of which are key components to creating a life a person can call their own. So, we will want to preserve as much of all these forms of autonomy as possible. Yet, as Stephen Holland (2009: 286) points out, only the most intrusive and comprehensive measures would actually interfere with our ability to direct our own lives, so if we are only interested in preserving broad autonomy, very few measures would be ruled out. Nevertheless, Nys has a point: Other things equal, paternalistic measures that affect only small and unimportant aspects of our lives are less of an intrusion and more justifiable for that reason.

A related and perhaps more powerful argument for when paternalism might be justified is when measures protect what Andrew Sneddon calls "deep autonomy" (Sneddon, 2001: 108; see Sneddon, 2006: 392 and Nys, 2008: 67). We have "shallow autonomy" when we do what we want to do (Sneddon, 2001: 108) in a robust sense: when we act on our desires that are supported by our second-order desires about what we should desire. Deep autonomy goes further: It

> consists in reflection on one's desires and the structure of one's life. It concerns, minimally, deciding whether it is worth doing what one currently wants to want to do. More extensively, deep autonomy is exercised in attending what it would be to lead a life worth living (Sneddon, 2001: 108).

Deep autonomy is about having the ability to reflect on one's whole life, "the ability," Nys notes, "not only to

Downloaded from https://academic.oup.com/phe/article/19/2/phag008/8679220 by University of Rochester user on 21 May 2026

decide and act upon deep-rooted values, but also the capacity to reflect upon those founding values themselves" (Nys, 2008: 67). The prohibition on selling oneself into slavery, Sneddon argues, prevents a person from alienating his ability to act on his own values and thereby violate his deep autonomy.[2]

The general kinds of reasons that have been offered to justify paternalism do so in an effort to protect some important aspect of autonomy. Soft paternalism allows us to ensure that we make decisions voluntarily and measures to protect both broad and deep autonomy safeguard autonomy itself. In general, then, the case against paternalism is that it interferes with autonomy, and we allow exceptions when autonomy is better served by some interference. How seriously a paternalistic measure interferes with autonomy depends on exactly how it affects autonomy.

What, then, does autonomy require? For my purposes, we do not need to wade through the extensive literature on the true value of autonomy or even on its uses in bioethics. Instead, I want to focus on three important elements: (i) autonomous decisions are made with an understanding of the decision and its likely benefits and costs, (ii) they are made freely, in the sense that no one is coercing the person to make a decision, one way or the other and (iii) they are made in ways that reflect the values of the person. We often treat each of these elements simply as a necessary condition for autonomous actions that are either met or not. But a little reflection shows that each can be met more or less, and so the autonomy displayed is a matter of degrees. A little more reflection suggests that each of these three elements is independent of the other two. Each of these elements constitutes a separate axis of paternalistic considerations.

Axis 1: Information

Consider, first, information. Stopping the pedestrian, even forcibly, who does not realize that the bridge ahead is out is, as I have already noted, Mill's paradigm case of justified paternalism; stopping the pedestrian protects his liberty by allowing him to understand what risks he is facing. Notice that we are sometimes justified in using considerable force to ensure that he understands: To the man at the bridge, Mill says, we "may *seize him* and turn him back" (Mill, 1859: 5.5, emphasis added). Notice also that we are making assumptions about what values are important to him—though the assumption in this case is only that he does not want to fall into a river. But the latter is provisional. Once warned, "no one but the person himself can judge of the sufficiency of the motive which may prompt him to incur the risk" (Mill, 1859: 5.5). Such measures are deemed a form of "soft paternalism," because they do not ultimately prevent any actions; people must simply demonstrate that they understand what they face.

How much information we need to have is, of course, a matter of degrees: A decision will be more autonomous, the better the person understands their situation. But, generally, we think interference is only justified if a person can harm themselves significantly; only then does the information matter enough to enforce the information requirement. Indeed, the greater the risk, the more we think some intervention is necessary. The dangers of sugar and fat are real, but not immediate, so we can require labels on foods so people can easily access the information if they want it. The dangers of canoeing on a lake are real and immediate, but the risks are generally pretty small, so we require people to sign a form on which they acknowledge the risk (even if they rarely actually read it). However, the dangers of surgery under anesthetics are real and immediate, and the risks are significant, so we require people to be counseled on the risks and then to sign a more extensive form. If the surgery carries particularly significant risks and the patient has time to contemplate it, we may even require them to take home materials, so they have the chance to think about their decision before we allow them to sign the form. We might even require them to pass some sort of test to demonstrate that they fully understand their situation.

In principle, then, requiring information is not a breach of autonomy. However, governments can try to make the knowledge requirements sufficiently onerous that they are meant to deter the action. For example, some states require a woman seeking an abortion to undergo an ultrasound, ostensibly to ensure that the woman understands that she is killing a fetus and that she may regret her decision. But since women almost always understand what they are doing when they seek an abortion, the purpose is both to discourage the procedure and to impose a hurdle to having it done. So, in truth, these measures are a moderate form of coercion.

However, we think about it, since it is temporary, the information requirement is not a deep form of paternalism. While it may cause considerable inconvenience, it is not ultimately a barrier to the person doing what they want.

Downloaded from https://academic.oup.com/phe/article/19/2/phag008/8679220 by University of Rochester user on 21 May 2026

Axis 2: Coercion

Autonomous decisions are made freely, without coercion, understanding—with Mill—that coercion can come in many forms, some more subtle than others, as we have just seen in the ultrasound case. In laying out his argument against paternalism, Mill is concerned with any attempt "in interfering with the liberty of action" of an individual, which he then equates with "compelling him or visiting him with any evil in case he do otherwise," which he contrasts with "remonstrating with him, or reasoning with him, or persuading him, or entreating him" (Mill, 1859: 1.9). For Mill, any attempt to use any means that do not appeal to the person's reason and interests broadly construed is illegitimate.[3]

As Mill shows, interference comes in many forms, and we can think of them on a spectrum. On one end of the scale, we can simply prohibit an action by making it impossible. Even if someone tries to sell themselves into slavery, we do not allow them to do so. Someone could sign such a contract, but since the contract is illegal, they can get out of it anytime they choose. But a slave who can leave their slavery anytime they choose is not a slave at all. Slightly less coercive are laws that place severe penalties on an activity. Using heroin is illegal in the United States. People can and do still use it, but they face serious time in prison if they are caught. The government thus applies considerable coercive force against the choice—however well-informed—to use heroin.

On the other end of the spectrum are measures that employ what Richard Thaler and Cass Sunstein call "libertarian paternalism" (Thaler and Sunstein, 2003). As they use the term, libertarian paternalism involves no coercion at all. Choices are structured one way or another, but the way they are structured will influence the decisions people make. Food in a cafeteria has to be organized one way or another; if healthy foods are prominently displayed, people are more likely to choose them. By placing the healthy foods in plain sight, we can make it more likely they will eat better without removing any choices from them. Or, employees will either be automatically enrolled in a retirement plan or they will have to opt into it; if they are automatically enrolled, they are more likely to contribute to their retirement plan earlier, even if they can easily opt out. If we force them to opt in, they are more likely to forget or to neglect the paperwork, even if they wanted to do so. By structuring choices in ways that are likely to benefit people, we influence their behavior for their own good, and if the other choices are readily available, then the "better" option is not being imposed on them in any significant way. Conceptually, the question here is whether constructing such choice architectures counts as paternalism at all since they fail to meet the first of Dworkin's three conditions. But I take it as a limiting case: We are trying to influence people's behavior for their own good by relying on their psychological flaws rather than on rational persuasion.

In between making a practice illegal and encouraging it through noncoercive means, there are a range of options, from those that simply try to encourage one behavior over others to those that make it difficult without banning it. The range can be seen in tax policies. Mill declares that taxing a product is morally equivalent to banning it:

> To tax stimulants for the sole purpose of making them more difficult to obtain is a measure differing only in degree from their entire prohibition, and would be justifiable only if that were justifiable. Every increase in cost is a prohibition to those whose means do not come up to the augmented price; and to those who do, it is a penalty laid on them for gratifying a particular taste (Mill, 1859: 5.9).

While few taxes are so burdensome that they literally keep some people from choosing them, Mill has a point. Cigarette taxes in the USA are onerous, and young people, in particular, are often discouraged from smoking by the costs of cigarettes. Yet people (even poor people) can and do choose to keep smoking (often because they are addicted), but it does significantly affect the other choices that they can make. In other words, the taxes can significantly impede their general freedom of action. They can thus be quite coercive. They direct a certain kind of choice without actually requiring it.

Other taxes are less burdensome: A one-peso-per-liter soda tax in Mexico is a modest tax that discourages soda consumption without making its purchase impossible or particularly burdensome. And, it does seem to have a modest effect on soda consumption (Fernandez and Raine, 2019). Since such taxes encourage people to make choices that would not otherwise make, they are still coercive, if only mildly so. We are attempting to channel people into certain behaviors, and we are using some form of compulsion to do so.

The use of coercion is harder to justify than the requirement that people have sufficient information to

Downloaded from https://academic.oup.com/phe/article/19/2/phag008/8679220 by University of Rochester user on 21 May 2026

make the decision. But even Mill thinks some coercion may be justified, as the slavery case shows. But many think coercion is justified in other cases as well. Seat belt laws, for example, impose significant fines on people for not wearing their seat belts. Such laws are not designed to protect other people: Relatively few injuries in car accidents are caused by unseat-belted passengers being flung into others. But many injuries are caused by unbelted passengers being thrown into steering columns, against dashboards, and through windshields. Seat belt laws, then, force a choice on people for their own protection. They also save almost 15,000 lives each year in the US (NHTSA 2025). But another factor makes a difference here: Very little is at stake in forcing people to wear a seat belt. No one's life plan is upset by wearing a seat belt, and the inconvenience is trivial.

Contrast this claim with motorcycle helmet laws. Those laws too are designed to protect the rider, not others. About 500 lives a year would be saved in the US if all motorcyclists were required to wear helmets (CDC 2026). Some argue that we can justify the laws by claiming that bad accidents create health care costs that the rest of us must bear, but we could—less coercively—simply require riders to have enough insurance to cover the costs of their accidents. The real reason we require helmets is that we want to protect riders from their own recklessness. But strapping on a helmet is much more inconvenient than wearing a seat belt. And for many riders, part of the appeal of the motorcycle is being close to the road and close to nature, feeling the elements as they ride. A helmet interferes with those feelings. For that reason, such laws are harder to justify than seat belt laws—even if we think they are still justified by the number of deaths and serious brain injuries they prevent.

Axis 3: Values

The contrast between seat belt laws and motorcycle helmet laws brings us to the third axis of paternalism: values. Autonomy is served when we act on values that are important and that we self-consciously endorse. Being in control of the central decisions in our life is supremely important. Protecting broad autonomy is, then, essential. So paternalistic actions are harder to justify the more the values compromised are closer to the core values that someone has. Different measures will, of course, impact the values of different people in different ways. With public health measures, we need to be cognizant of the areas in which a restriction is likely to impinge upon important values. We will have to weigh the good that the measures might achieve against how likely it is to impinge on important values for a lot of people. Since, for example, religion is a key to many people's self-identity and to the values that are often entangled in that identity, measures that impinge on religion are particularly fraught. For that reason, we may want to allow religious exemptions from particular paternalistic public health measures. So, if we were contemplating a mandate for, say, the covid vaccine now—when it does little to prevent the spread of the virus from one person to another and therefore does little to prevent harm to others—then we may want to allow an exemption for those with religious objections. When we can, out of respect for important values, allow exceptions to a general rule without undermining the general good the rule is supposed to achieve, we should do so.

On the other hand, sometimes we can protect broad autonomy best if we are prevented from taking actions that undermine our important goals. If I want to limit my intake of empty calories from drinking sodas, a measure that makes me buy a 12 oz. soda rather than a 20 oz. one helps me achieve my goals. Of course, even if I am on a health kick, this goal is relatively trivial. Yet knowing my own cognitive biases, measures that help me combat those flaws will help me live by my own values. Indeed, as Conly argues, occasionally

> What we need is to make decisions in optimal situations that we will later be forced to abide by, even when we are so tempted that we do not want to. Typically, the only way to truly bind ourselves is to allow legislation that prevents us from taking the tempting, but self-destructive, step. The control needs to go, not to superior people, but to all of us when we are doing our superior thinking (Conly, 2013b: 242).

Promoting my own broad autonomy may require some external help.

Beyond promoting broad autonomy, preserving deep autonomy may justify interventions.

Maintaining my ability to make choices for myself in the future is almost certainly an important value. A prohibition on selling myself into slavery would, then, help me promote my own autonomy. Nys thinks that deep autonomy is not only important, but inalienable, so he thinks promoting deep autonomy is always required. He wants to tie health to deep autonomy as a

Downloaded from https://academic.oup.com/phe/article/19/2/phag008/8679220 by University of Rochester user on 21 May 2026

justification for many public health measures. By preventing serious health problems that will undermine a person's ability to continue making choices for herself, he thinks, public health protects deep autonomy. However, he overstates the connection: While health is necessary for most life plans, it has a subordinate role in others. Playing football, skydiving and free solo climbing all pose risks to health, but each might be an important part of a person's life plan. As Holland notes (2009: 286–88; see Wilkinson, 2025: ch. 3), Nys assumes (i) that health is always essential to deep autonomy and (ii) that deep autonomy is always more important than the simple ability to choose as one likes. Holland points to the romantic idea that some people want to live fast and die young, but we also might point to lives of scientific dedication and self-experimentation that carry significant risks, or those of ballet dancers who ruin their bodies through their art. Nevertheless, even if Nys's view is too strong, deep autonomy should not be violated except in the most extreme circumstances.

The deeper point here is that paternalistic measures can affect the realization of our values. We thus have another spectrum between actions that promote important values and those that impede them, with a large swath of actions in the middle where the values involved are not particularly important, one way or the other. They either do little to promote or impede key values or they promote or impede values that are relatively unimportant. Wearing a seat belt is at worst a trivial annoyance; autonomy is not greatly impacted by the requirement. Since it saves a great number of lives, it can be justified on those grounds. The same cannot easily be said of motorcycle helmet laws. While I do not personally share the thrill of riding a Harley without a helmet, I can certainly understand it. It is probably not, however, central to anyone's identity, even if riding a motorcycle is. Still, given the importance of that experience to some people, it will impinge on their autonomy, and we must weigh the importance of that interference against public health gains of mandating helmets.

Spectra

We now have three spectra of considerations when we think about autonomy. Since they are independent, they create a three-dimensional space that we can use to frame discussions about paternalism. Since the information spectrum is mostly temporary, I will focus on the two-dimensional space of the coercion spectrum and the values spectrum. Consider Figure 1. Acts at the bottom left involve decisions that promote values and use little force, so they do not infringe much upon autonomy, while those in the upper right do. The diagonal line represents a spectrum that reflects the degree to which a measure encroaches on autonomy. If life were simple, everything in the upper right would be prohibited, everything in the lower left would be permitted, and the only debates would be about the balance between coercion and values that fall in the other two quadrants.

Alas, life is not so simple, for two reasons. First, paternalism in a public health context only makes sense when there is a health benefit to the person. So, for our purposes, the health benefit functions as a separate axis.[4] How much health benefit people will receive makes a difference to how much we should be permitted to interfere with a person's autonomy. In general, then, a paternalistic action that would be placed in the lower left corner of the figure need only produce modest health benefits to be justified, while no health benefit could ever justify those that fall in the upper right. For example, a seat belt law produces significant benefits, and it neither impedes nor promotes important values, so it is justified even though it is quite coercive. The placement of food in a cafeteria involves food, so it is not completely trivial, but virtually no coercion is involved. So, it is justified. Modest soda taxes are mildly coercive, and they do not involve a practice that is deeply valued, so they are justified—but only if they work. The evidence suggests that they reduce soda consumption somewhat, but there is no convincing evidence that such measures actually reduce obesity. A ban on smoking, on the other hand, is extremely coercive, and it concerns a practice that is—for better or worse—a part of some people's identity and so they value their connection to it. For that reason, then, it is probably unjustified. However, the fact that smoking accounts for roughly 480,000 deaths per year in the USA (CDC 2024) is a strong reason to think that such a ban should not be considered out of the question. Heavy taxes are significantly less coercive, so given those deaths, they are, I think, justified.

Public Health and Pluralism

The second reason the picture is complicated is that, for the purposes of illustration, I have made assumptions about the importance of various values, and I have

Downloaded from https://academic.oup.com/phe/article/19/2/phag008/8679220 by University of Rochester user on 21 May 2026

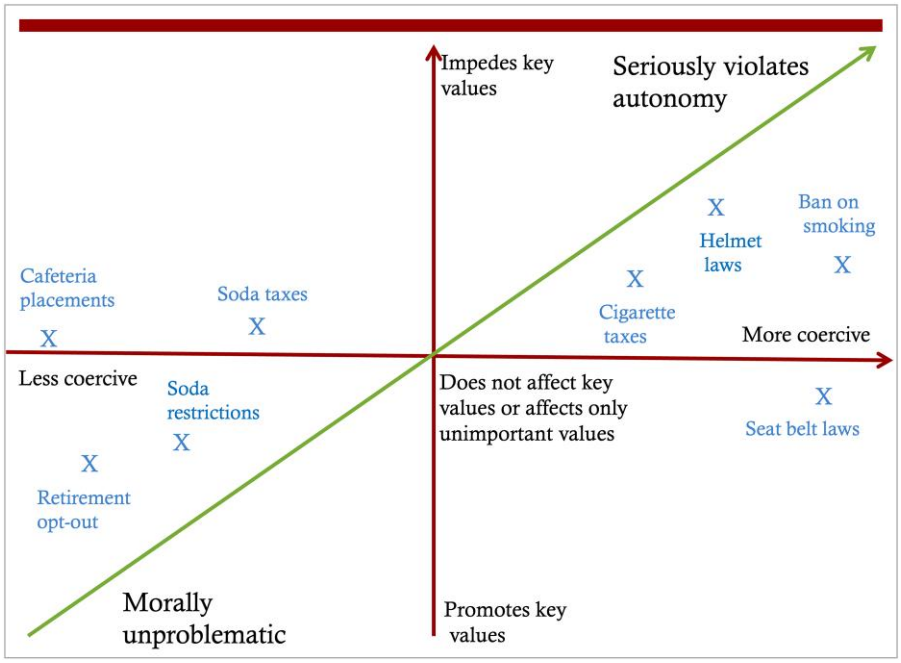

Figure 1. Spectrum of coercion mapped against spectrum of value.

placed various actions on the grid in Figure 1 based on those assumptions. But those assessments are simply my judgments about where most people would place those actions. I fully realize, however, that different people will place different importance on different values—including the values of health and autonomy themselves—so different people will judge how serious an infringement on autonomy any particular action is in different ways. We can imagine someone who thinks that riding in a car without a seatbelt is not only pleasurable, but an important joy in their life (though I think in fact we should have some worries about such a person). But many motorcyclists value helmet-free rides and value them a lot. If everyone had the same values as I do, then no one's autonomy would be seriously compromised by a law requiring helmets, so such a law might be justified. But since they do not, it is not. Indeed, we might argue, because people have different values, every public health intervention will impose some value scheme on people and thereby restrict someone's autonomy significantly. If so, then no measure can be placed definitively on the values spectrum, and we can never measure the trade-off between autonomy and health that would be necessary to justify a law. Even in this scenario, thinking of autonomy in a complex three-dimensional space is useful because it helps us situate what is at stake in each decision. But it will do little to frame decisions about public health laws.

This skeptical argument, however, rests on the assumption that pluralism about values implies that no justified claims can be made about what is good for people in any general way. T.M. Wilkinson, for example, bases his broadly skeptical arguments about public health paternalism by assuming that "ultimate preferences measure well-being" (Wilkinson, 2025: 43). For that reason, "the case for public health interventions must give reasons to think that, in making people healthier, they also make people better off as judged by their own preferences" (Wilkinson, 2025: 53). Such a subjectivist view of well-being is itself controversial, and Wilkinson only accepts it as a kind of default position, the least controversial alternative, compatible with many more objectivist views (Wilkinson, 2025: 42–49). But the same problem would arise if different people simply valued in different ways things that are objectively valuable. But even Wilkinson rejects an extreme form of the problem. Sometimes, he thinks, ultimate preferences can be mistaken: the anorexic who is willing

Downloaded from https://academic.oup.com/phe/article/19/2/phag008/8679220 by University of Rochester user on 21 May 2026

to die rather than eat, he concedes, has made a mistake about what is really valuable for her (Wilkinson, 2025: 52–53). For my purposes, this concession is significant. If we accept that people can be mistaken about the value of their ultimate preferences, then those preferences do not decide all questions about what is valuable for them and what will therefore best promote their autonomy. In so, there will be some cases in which we can place an intervention at a definite place on the values spectrum, and we can then ask whether the health gains are worth whatever intrusions on autonomy it entails. We will then have debates about particular cases and whether the values involved are ones that should be respected just because some people choose them. But in that case, thinking about the spectra of autonomy is useful.

Yet we can make a further claim. The placement of a measure on the values axis as measured by people's preferences and values is not usually scattered randomly. Even if values impacted by an intervention cannot be narrowed to single point, they usually exist along a relatively small range. If so, then the interference with autonomy will exist in a more or less broad area of the coercion-value plane, and if the health benefits are great enough, then the measure would still be justified. So, measures that help many people achieve their goals or which do not interfere with those goals as measured by their own preferences may justify a paternalistic law—especially if the coercion involved is not serious. As Wilkinson notes, even many smokers endorse heavy taxes on cigarettes as a means of discouraging smoking (Wilkinson, 2025: 87–89). In this case, even though the taxes are fairly coercive, the non-health values are not so important that the health benefits of the taxes are outweighed. This kind of logic then may permit broader measures that promote the health goals of many people, perhaps even the kinds of laws in which we bind ourselves collectively, as Conly has suggested. Even if, like Wilkinson (2025: 129), we think such cases are rare, placing these interventions on the plane of autonomy helps us to see better what is at stake.

Conclusion

My purpose here is not to determine exactly when public health laws might be justified. Instead, I hope to reframe debates about when paternalism in public health is appropriate. We should reject the simplistic view that paternalism is never justified and think seriously about when exactly intrusions on liberty and autonomy are small enough and the benefits great enough to justify an intervention. By looking at the nuances of the values that are at stake, we can enrich the debate. By trying to see where in the three-dimensional space a particular action lies helps us to see which elements of autonomy might be compromised—and how much they will be. In doing so, we can make better judgments about the propriety of public health measures.

Notes

1 Notably, Joel Feinberg (1986) and Richard Arneson (1980) defend the view that paternalism is never justified, mostly by explaining away the cases that I cite here.
2 Nys offers two other grounds to justify paternalism: (i) that we can effectively consent through democratic processes to actions that limit our freedoms (Nys, 2008, 68–69) and (ii) that actions that benefit us may be justified on grounds of justice or solidarity (Nys, 2008, 69–70). However, democratic majorities can be just as paternalistic as authoritarian powers, and actions done for the sake of others are not paternalistic, even if some people chafe against the loss of liberty.
3 Such a view is much broader than that considered by the philosophical literature on coercion that arose from Nozick, 1969. For an overview, see Anderson, 2023.
4 Health, of course, is itself a value, and one of the issues at stake here is how important a value it is. We separate it out here because we are interested in thinking about a set of measures that are intended to promote health. Presumably, we could perform a similar exercise for other values, like truth or beauty.

Acknowledgements

The author thanks Nicholas Novk and the students of the University of Rochester for discussions of these issues and the anonymous referees for Public Health Ethics for trenchant comments that helped clarify the issues at stake.

Manuscript received: June 2025

References

Anderson, S. (2023). Coercion. In Zalta, E. (ed.), *Stanford Encyclopedia of Philosophy*, available

Downloaded from https://academic.oup.com/phe/article/19/2/phag008/8679220 by University of Rochester user on 21 May 2026

from: https://plato.stanford.edu/entries/coercion/ [accessed 6 May 2026].
Arneson, R. J. (1980). Mill Versus Paternalism. *Ethics*, **90**, 470–489.
CDC (2026). *Motorcycle Safety*, available from: https://www.cdc.gov/pedestrian-bike-safety/about/motorcycle-safety.html?CDC_AAref_Val=https://www.cdc.gov/transportationsafety/mc/ [accessed 11 April 2026].
CDC (Centers for Disease Control and Prevention). (2024). *Cigarette Smoking*, available from: https://www.cdc.gov/tobacco/about/index.html [accessed 11 April 2026].
Cholbi, M. (2017). Paternalism and our Rational Powers. *Mind*, **126**, 123–153.
Conly, S. (2013a). *Against Autonomy: Justifying Coercive Paternalism*. Cambridge: Cambridge University Press. 2013.
Conly, S. (2013b). Coercive Paternalism in Health Care: Against Freedom of Choice. *Public Health Ethics*, **6**, 241–245.
Dworkin, G. (2020). Paternalism. In Zalta, E. (ed.), *Stanford Encyclopedia of Philosophy*, available from: https://plato.stanford.edu/entries/paternalism/ [accessed 6 May 2026].
Feinberg, J. (1986). *Harm to Others*. Oxford: Oxford University Press.
Fernandez, M. and Raine, K. (2019). Insights on the Influence of Sugar Taxes on Obesity Prevention Efforts. *Current Nutrition Reports*, **8**, 333–339.
Hanna, J. (2018). Hard and Soft Paternalism. In Grill K. and Hanna J. (eds), *The Routledge Handbook of the Philosophy of Paternalism*. London: Routledge, pp. 24–34.
Holland, S. (2009). Public Health Paternalism—A Response to Nys. *Public Health Ethics*, **2**, 285–293.
Mill, J. S. (1859). *On Liberty*. Indianapolis: Hackett Publishing. 1978. Referenced by chapter and paragraph number.
NHTSA (National Highway Traffic Administration). (2025). *Seat Belts*, available from: https://www.nhtsa.gov/vehicle-safety/seat-belts#:~:text=Being%20buckled%20up%20during%20a,vehicle%20is%20almost%20always%20deadly.&text=If%20you%20don't%20wear,injure%20or%20even%20kill%20you [accessed 11 April 2026].
Nozick, R. (1969). Coercion. In Morgenbesser, S., Suppes, P. and White, M. (eds), *Philosophy, Science, and Method: Essays in Honor of Ernest Nagel*. New York: St. Martin's Press, pp. 440–472.
Nys, T. (2008). Paternalism in Public Health. *Public Health Ethics*, **1**, 64–72.
Shiffrin, S. V. (2000). Paternalism, Unconscionability Doctrine, and Accommodation. *Philosophy & Public Affairs*, **29**, 205–250.
Sneddon, A. (2001). What's Wrong with Selling Yourself into Slavery?: Paternalism and Deep Autonomy. *Critica: Revista Hispanoamericana De Filosofia*, **33**, 97–121.
Sneddon, A. (2006). Equality, Justice, and Paternalism: Recentreing Debate About Physician-Assisted Suicide. *Journal of Applied Philosophy*, **23**, 387–404.
Thaler, R. and Sunstein, C. (2003). Libertarian Paternalism. *American Economic Review*, **93**, 175–179.
Wilkinson, T. M. (2025). *The Ethics of Public Health Paternalism*. Oxford: Oxford University Press.

Downloaded from https://academic.oup.com/phe/article/19/2/phag008/8679220 by University of Rochester user on 21 May 2026