

Making Sense of Medical Paternalism

By Robin Hanson, Sep.7, 2007

Why do we regulate the substances we can ingest, the advisors we can hear, and the products we can buy far more than similarly-important non-health choices? I review many possible arguments for such paternalistic policies, as well many possible holes in such arguments. I argue we should either be clearer about what justifies our paternalism, or we should back off and be less paternalistic.

Paternalistic Regulation

Government regulations often limit our choices, arguably for our own good. For example, securities regulations limit the ventures in which we can invest, education laws limit the schools we can attend, and legal rules limit the contracts to which we can agree. Such limits are especially common in health and medicine regulation. For example, food and drug rules limit the substances we can ingest, professional licensing limits who can advise us, product safety rules limit the products we can buy, vehicle safety rules limit the vehicles we can use, and speech limits restrict the advertisements we can hear about all these choices.

Such regulations are often called “paternalistic” because they seem like parents limiting the actions of children. That is, they seem like parents who tell their children “you may not play in the street,” “you must eat this food,” or “you must go to bed now.” Parents who say such things tend to honestly believe children would be worse off without such enforced limits.

But eventually, most children reach a point where they think themselves better off without such limits. Such people think it sufficient to be able to seek parental advice when they think such advice might be useful. And in fact this is the situation most of us find ourselves in regarding most of our choices. For example, we usually enjoy broad freedoms regarding the non-health features of our food, associates, products, and travel plans.

Are health choices different?

This raises a fundamental puzzle: Why do we feel we should regulate health and medical choices so differently from other choices? Why do we think people with all the usual human frailties are better off with the freedom to choose their own sex life, roommates, careers, places to live, and children, but not their own food, drugs, cars, and medical advisors? Sure some health choices are hard to reverse and have big consequences, but so are many far less regulated choices such as changing nations or having babies. And we remain paternalistic about even small reversible health choices.

Now we have lots of data showing the public does in fact support paternalistic policy. For example, U.S. citizens were in 2006 asked who “should be primarily responsible for setting the rules for food handling, production and packaging in order to ensure the safety of these products?” In response, 72% chose a government organization, while only 17% chose a private organization [1]. And in 2005 when asked whether there is too much or not enough regulation “making sure prescription drugs are safe for people to use,” 50% said not enough and only 8% said too much [2].

But what we lack is data about the exact reasons people support such policies. And what makes it hard to even collect such data is that most people are not even aware of what the possible reasons might be. Therefore for the remainder of this essay we will content ourselves with merely trying to be clear about possible reasons one might offer for medical paternalism, and about the problems with such reasons.

Attitudes to “cliff-walkers” – an example of health choices and paternalism

Since all situations of paternalism have basically the same general structure, let us focus on a particular example for concreteness. Therefore let us now consider “cliff-walkers” as prototypical health choosers.

Imagine finding yourself near someone about to walk off a cliff. If he seems distracted enough to not notice a crucial bend in the cliff-edge, you might feel quite justified in grabbing his arm, to stop him from falling. You might even expect his gratitude.

But what if he seems well aware of the cliff before him? Well, if he seems crazy, either permanently insane or temporality drugged, you might still grab him. You might also grab him if you knew his family would miss him terribly. In such cases you might at least expect gratitude from his family, his caretaker, or his future sober self. And if you were morally outraged enough by the very idea of walking off a cliff, you might grab him no matter who was grateful or offended.

But what if, aside from the whole cliff thing, he seems no crazier or immoral than most? What if his action mainly affected only him? What if the cliff was only five feet tall, or twenty feet tall over deep water, or if he walked near the cliff at what he considered a close but safe distance? You might still think of grabbing his arm, if you thought you understood something important that he did not. Perhaps you know the wind is unusually gusty, or the ground is unusually slippery. Perhaps there is no time to explain, or he doesn’t understand your language.

But what if he does understand you, and there is time enough to say “Watch out! That cliff is dangerous.” If he dismisses your concern and does not back away, would that justify your intervention? Well we can’t very well allow anyone to intervene in anyone else’s life anytime they feel like it. So if you persist in grabbing we might let him sue you for assault.

But what if you were not alone? What if a great many of you also thought him careless? What if you lived in a democracy and could get enough voters to pass a law banning cliff-walking? Perhaps your law requires tall fences, or threatens to jail those who approach cliffs. Are you justified now?

Even in this situation, you are arrogant if you do not at least consider the possibility the cliff-walker knows what he is doing. After all, you must admit he might understand something about this situation you do not. And surely he knows better than you just how much he cares about, for example, fun and adventure versus safety and security.

Data, consensus and choice

Now we could imagine someone collecting enough solid data to clearly show that humans just do not respect cliffs enough, relative to their long-term self-interest, even after they have been warned. And that would seem to settle the case. But what if you have no such clear data, and this is more a matter of judgment?

Perhaps you think that those who agree with you simply have better judgment than cliff walkers and their supporters. If so, you should ask yourself: what is your basis for this conclusion? Is it because there are more of you, and pretty much all decisions should be made by majority rule? Is it because you think your side belongs to a superior age, gender, class, or ethnicity? If so, why do you also regulate people on this superior side?

Perhaps you can point to objective features of people on your side which suggest they are better informed. For example, maybe most professors of geology, who understand in great detail how cliffs are formed, support your view. But the relevance of geology expertise is not exactly overwhelming here. Doctors who care for cliff-fallers might be more relevant, but even this expertise is not obviously enough to outweigh cliff-walkers' superior expertise in their own skills and values.

Even if your side does clearly have more relevant expertise overall on cliff-walking decisions, this advantage can be completely negated if the cliff-walker is simply a good *listener*. That is, if he hears your warning, realizes you might know more than he, and fully takes this into account when considering his choice, then he would mainly reject your advice only when doing so is on average his best choice. If so, you might be justified in stopping him to benefit someone else, but not for "his own good," as he understands it.

Now perhaps you think people on your side are better listeners. That is, you think you guys carefully listen as much as is appropriate to opinions of cliff-walkers and their defenders, but they in contrast are too proud and sloppy to listen well to your warnings. Well in this case you are arrogant, i.e, biased to presume your own superiority, if you do not at least consider what basis you might have for such a lopsided conclusion. Is it, again, your superior age, gender, class, or ethnicity, or the fact that you are in the majority?

Perhaps both sides are equally bad listeners, and so both sides are too proud and sloppy to give enough weight to valid expertise on the other side. Well in this case the relative quality of your judgments would depend more on your relative expertise. But ask yourself: if you are so proud and biased that you don't listen to the other side's valid expertise, how sure can you be that that similar pride and bias hasn't distorted your judgment of which side has more relevant expertise overall?

Can advice be trusted?

Consider another possibility. Perhaps cliff-walkers aren't so much bad listeners as suspicious ones. That is, what if the reason they don't give full weight to your advice is that they do not trust you to be completely honest and helpful. Maybe in the past people like you gave bad advice to people like them, advice distorted by other agendas. Perhaps such advisors were selling something, or too easily assumed everyone wanted to be like them, or loved the sound of their own voice, or enjoyed controlling the lives of others.

Let us focus on this distrust possibility, as it allows us to analyze the situation without arrogantly presuming your own superiority. So let us set aside concerns about who is how proud, biased, or sloppy, and assume for the purpose of argument that both sides are reasonable. We'll assume your side has far more relevant expertise, but that the other side is completely capable of listening to your advice, *if* they trust you. And assume they find it hard to distinguish you from people who would offer distorted advice.

Imagine you are absolutely sure your advice serves only their interest as they understand it, without distortion from other agendas. Since they won't listen enough because they suspect otherwise, doesn't this justify you forcing them to follow your advice? Well this wouldn't justify a *general* policy of letting people who look like you force their advice on people like them, as this general policy would also let others force distorted advice. And cliff-walkers reasonably judge this to be bad for them.

What if you can choose to intervene without creating a precedent allowing distorted others to also intervene? Does this make intervention all right? Maybe, but shouldn't you admit that you can't be absolutely sure you have *no* other agendas? What if you and they both expect you to have small distortions? If you might have a lot of information that could prevent a lot of harm, shouldn't that outweigh the possibility of a small distortion?

Hopefully you can see that this situation is pretty complicated, too complicated for you to easily trust your own intuition about such things. So let me tell you about the results of an analysis I published a few years ago about this sort of game [3].

The game has two rational agents, a "doer" (e.g., cliff-walker) who chooses an activity level (e.g., how close to the cliff edge to walk), and an "advisor" (e.g., regulator) who knows how risky is that activity. In one scenario both sides know the advisor can only give advice, while in another scenario both sides know the advisor can choose either to

give advice or to “ban” the activity. A ban limits the activity level to some low level; for example, a fence might keep everyone at least twenty feet from a cliff.

Assume that for most risk levels, the activity level the advisor prefers is close to, but not quite the same as, the doer’s. For example, if an all-knowing cliff walker’s ideal cliff distance were ten feet, an all-knowing advisor might prefer eleven feet instead. So if the advisor would be exactly believed by an ignorant doer, that advisor would be tempted to give slightly distorted advice.

Are bans better than advice?

But the doer anticipates such distortions, and so will not believe everything the advisor says. Because of this, the advisor can only communicate a small number of risk distinctions, such as that the risk is either “low,” “medium,” or “high.” The more distortion, the fewer distinctions are possible. The advisor may know more detail, but the doer will not believe more.

As a result, an advisor can regret being only able to give advice. For example, when the advisor knows risk is high, the doer may only believe a moderate warning, and so a ban could produce a better outcome. On the other hand, an advisor who can ban can regret having that power. For example, if the risk is moderate, the doer may not believe a moderate warning, because he expects bans in such cases. So the advisor may have to choose between a mild warning and an extreme ban.

Averaging over many risk levels, is it better if advisors can only advise, or if they can also ban? The answer depends on whether an all-knowing advisor would prefer *more* activity or *less* than an all-knowing doer. When the answer is *less*, the game turns out to be too complex to say much in general. When the answer is *more*, however, then the power to ban on average makes both sides worse off!

For example, a teacher who wants his rational but ignorant and cliff-averse students to study the view from a cliff should on average regret his ability to ban them from the cliff. Similarly, regulators who think that a reasonable public who understood the true quality of doctors and drugs would still choose too little medicine should regret being able to ban via professional licensing and drug regulation. And those who think that the public saves too little of their income may regret being able to ban investments via securities regulations.

What if, after all this analysis, you shake your head and say “I don’t need all this; I just know he is reckless and must be stopped. After all I’m a rational, well-educated, and rich, all my friends agree with me, and he’s just an irrational nobody.” Then I will shake my head at your inexcusable arrogance.

In conclusion

Do you support imposing limits on the food and drugs people can buy, or the medical advisors they can choose?

If you want to convince yourself and the rest of us that your support for such paternalism is based on more than a simple arrogant presumption that people like you can run other people's lives better than they can, you should make some effort to explain to yourself and the rest of us exactly why you think your paternalism is justified.

References

[1] Harris Interactive, online survey within the United States between December 12 and 14, 2006,

<http://www.harrisinteractive.com/news/allnewsbydate.asp?NewsID=1140>

[2] Kaiser Health Poll Report Survey February 2005,

http://www.kff.org/healthpollreport/feb_2005/16.cfm

[3] Robin Hanson, [Warning Labels as Cheap Talk: Why Regulators Ban Drugs.](#) *Journal of Public Economics* 87(9-10):2013-2029, September 2003.