

Vote On Values, Outsource Beliefs

Posted on May 8, 2014 by Scott Alexander



I

Today I learned about [social impact bonds](#). They are a thing that exists. I would expect them to be in an adequate civilization like [Raikoth](#) or [dath ilan](#). But they are a thing that exists on Earth.

The basic idea is: government could save a lot of money if some problem got fixed. For example, if people stopped committing crime, they could spend less money on prisons. So they make a deal with a corporation. The corporation agrees to spend a certain amount of money to prevent crime for five years. And if crime goes down and the government saves on prisons, the corporation gets half the savings (or a third, or whatever).

Zero taxpayer money gets risked. It is entirely up to the corporation to fund the problem-solving effort. If they fail, then it's their own loss. If they succeed, then the government pays them money, but less than the government made, so the taxpayers still get a profit.

(The main exception I can think of is if by coincidence, crime was about to drop by 50% anyway right when the program started, and

the government ends up giving half of its prison savings to the corporation for no reason. But presumably you hire a couple of mediocre economists and they are able to price out this risk. Also, a lot of the social impact bonds use a slightly different method of assessment, where they compare crime among the people the corporation has helped to crime among a control population to be sure it was the intervention that did it.)

The particular article I read about this today was [How Goldman Sachs Can Get Paid To Keep People Out Of Jail](#). It was the name “Goldman Sachs” that got me excited. They’re an investment bank. Their job is predicting risk. I don’t know if they’re any good at it or not. But they’re the sort of organization that potentially could be. So we have people who understand risk trying to figure out what social policies will produce which results, with money riding on the decision.

This is looking *impressively* close to prediction markets. Futarchy says [“vote on values, bet on beliefs”](#). Asking a corporation to invest money in crime-solving is a form of betting on belief – they are betting on what anti-crime programs will decrease crime most and win them the most reward. You still have the elected government deciding what bonds to place – voting on values – but you’re outsourcing your beliefs to the corporation involved and giving them an incentive to get it right.

Think of all the possibilities.

Right now we have a system where we don't really help people in need, unless the need becomes desperate, in which case we would feel bad about not helping, so we do, but then the cost of helping has gone up by an order of magnitude. This is exactly the sort of stupid thing that a market should be able to profit from solving.

We could have a health insurance company giving free preventative care to the poor, and the government paying them out of decreased emergency room visits.

A psychiatry clinic giving therapy to at-risk patients, and the government paying them out of decreased involuntary commitments.

A university accepting students without tuition, and the government paying them out of the increased tax revenue when they take higher-paying jobs.

Planned Parenthood offering free IUDs for women who need them, and the government paying them the money it saves from not having to put the kids through school.

Trade schools offering free classes to people on welfare, and the government paying them back from not having to give them welfare checks once they get good jobs.

I'm not sure what it means that we're not doing those sorts of things already. But if we can't figure out a way to solve those prob-

lems without bringing in a corporation to profit off of our incompetence, I say bring in the corporations.

II

I think many people are against government social programs for a lot of the same reason that *The Last Psychiatrist* is [against maintenance of certification exams](#) (a position I [totally called](#)). There's too much temptation to use it as a signal that you are Doing Something while in fact funding [programs like DARE](#) which look virtuous, but do nothing or even actively make the problem worse.

If you lean this way – and I think I do – then it is not solely out of stupidity that we wait until problems have become dire before doing anything about them. Yes, it would be great to give free job training to people on welfare and save money when they come off welfare more quickly. But actual job training programs for welfare recipients are abysmal and have been denounced as a “charade” from both [the left](#) and [the right](#). They may be a lost cause, but I would like to see someone who has an incentive to succeed try first before writing them off – or at least get the evidence that would be provided by no such person being willing to try.

III

For a while I was confused by the old libertarian talking point that “greed is good”. I think I could phrase it a little better now. Greed

isn't *good*, per se. It is *honest*. You know where you stand with greed. You never wonder if greed has an ulterior motive, because it's already the most ulterior motive there is. Greed feels no temptation to corruption, because the thing it would do if it were corrupt is precisely what it's doing anyway. Greed is like the Harlot in one of Khayyam's rubaiyat:

A Sheikh beheld a Harlot, and said he:
"You seem a slave to drink and lechery"
Replied the Harlot: "What I seem ... I am!
Are you, O Sheikh, all that **you** seem to be?"

As I see it, capitalism isn't about worshipping greed, but about figuring out how to make greed work for good ends. So far, it has mostly tried to apply greed to get us cheap and attractive consumer products. And the amount of cheap and attractive consumer products is, like, the one thing that everyone can unambiguously agree our civilization hasn't dropped the ball on. If we all die tomorrow and aliens discover Earth ten thousand years from now, their anthropologists will publish books saying "They sure were screwed up, but *man* did they have a lot of cheap and attractive consumer products."

And I think some of the most exciting proposals for the future involve finding ways to use this privileged incorruptible perfectly-incentivized status of greed to do other things. Prediction markets are promising because they use greed to fix epistemology. Neocameralism is promising because it uses greed to fix governance.

And social impact bonds are promising because they use greed to fix social problems.

...which isn't to say it's going to be easy. Ozy's first response is that Goldman Sachs should use their \$10 million to give ten thousand people in the control group a \$1000 bribe each to commit a small crime; this will be more than enough to demonstrate a *vastly* reduced probability of criminality by being in the intervention group and earn Goldman \$20 million.

I told Ozy zir plan is unnecessarily complex. Look at [the numbers](#). Two hundred potential criminals. And they need a 50% decrease in jail time to meet their target and earn \$20 million.

So go to the potential criminals and tell them "I'll give you \$50,000 to not commit any crimes in the next few years. \$25,000 now, in order to help you solve whatever problems turned you to criminality. And \$25,000 at the end, after you've successfully avoided jail, as a reward." If half of them stick to it, then boom, you get \$20 million and you've made a \$10 million profit. And incentivized the next generation of criminals, but you've already *got* your profit, that's the next generation's problem.

The fact that this would *work* probably says a lot about the inefficiency of prison compared to any other conceivable way of dealing with crime. And about the profits Goldman Sachs or anyone else willing to face the inefficiency head on could make.

I don't know if it's exactly a *good* idea to bring in the people who caused the financial crash to help the people who came up with the prison system. But since all we've got is incompetent institutions, maybe sticking *different* incompetent institutions in different roles might at least shake things up a little.