Bursting the selfishness bubble

Shalvi, S.

Bursting the selfishness bubble

Shaul Shalvi, University of Amsterdam

Madam Rector Magnificus, Mister Dean, Mister former Dean, Ladies and gentlemen, so good to see you all here.

In 2015, I joined the School of Economic, here at the University of Amsterdam. Trained as a psychologist, I was very happy that our management team invited me to co-teach a new first-year course called Principles of Economics and Business. Since then, together with my Business School colleague Nick Brunsveld, we have been teaching the section on 'the moral limits of markets'.

Preparing this course was a great opportunity for me to read literature I was unfamiliar with, and dive further into the world of economics. It also made it very clear to me why economists and psychologists like to work together. After all, economic decisions are often taken by people.

One paper I read was by Milton Friedman[1], probably the most influential economist of the 20th century. The title of the paper is "the social responsibility of business is to increase its profits". The idea is that business executives should only try to meet the goals set by their stakeholders. If stakeholders set the company's goal to maximizing profit, an executive – serving as their agent – must only attempt to achieve this goal. Jeopardizing the company's bottom line by producing products that are for example, less polluting, would accordingly be morally wrong.

Friedman suggests that executives, as private people, can choose to donate part of their salary to reducing pollution. But if they want their company to produce environmentally friendly products, they must convince stakeholders that it is either profitable, or a worthy cause to narrow the profit margin for. Friedman writes: "there is one and only one social responsibility of business —to use its resources and engage in activities designed to increase its profits so long as it stays within the rules of the game, which is to say, engage in open and free competition without deception or fraud."[2]

Reading Milton Friedman is great. The arguments are crystal clear and logical. It is difficult to disagree with even a single word. However, While Friedman's argument is accurate, the message it send is fundamentally wrong.

---

[1] Friedman, M. (2007). The social responsibility of business is to increase its profits. In Corporate ethics and corporate governance (pp. 173-178). Springer, Berlin, Heidelberg.

[2] Friedman, M. (2002). Capitalism and freedom: Fortieth anniversary, p. 133.

When Friedman writes profit, he mean long term profit. But when you study people's psychology you learn that when people think profit, they often think short term. Executives have to present achievements to stakeholders. The company's stock value should consistently increase. Sales must be on the rise.

Today, I propose that when people strive for profit, they start thinking in a bubble. The selfishness bubble. Thinking in a selfishness bubble, is dangerous. It can lead to shortsighted decisions and poor-quality products. Such products risk not only the company's long-term success, but also human lives.

Let's consider the evidence for selfishness bubble thinking, and how to burst this bubble.

Take the aviation company Boeing, for example. Ten years ago, pressured by the success of Airbus, Boeing considered two alternatives: develop a new aircraft, or redesign the 737 series. Choosing the cheaper option of the two, in 2017 Boeing introduced the fuel efficient fourth generation 737-Max.

When a company re-designs an existing model it must report to the regulator, if major changes were made. In such case, all certified pilots must be re-trained. Re-training takes time and is expensive. According to an excellent Netflix documentary on the case[3], and please watch it, if you have not done so yet, Boeing executives, were eager to turn a quick profit. To do so, they guided the engineers and lower-level management to re-design but avoid re-training. The aircraft was eventually engineered in a way that did require re-training. Certification papers were filed, however, suggesting re-training was not required.

This short-term profit driven decision proved deadly. In October 2018, an Indonesian 737-max crashed 13 minutes after take-off, killing all 189 passengers and crew. Blaming the Indonesian crew, Boeing suggested the crash was due to human error. The documentary convincingly shows, that Boeing knew what the cause of the problem was. They also knew that a similar crash is rather likely, estimating that it will occur within the next two years. Instead of grounding the aircraft immediately until they solve the problem, or retraining pilots to work around the problem, flights continued. Boeing were trying to the fix the problem while keeping the aircrafts in the air.

They hoped all would go well. It did not. Less than a year later, in March 2019, an Ethiopian 737-max crashed 6 minutes after take-off, killing all 157 passengers and crew. Again, Boeing refused to take responsibility, suggesting human error caused the crash. This time regulatory bodies around the world grounded the aircraft until the problem was fixed.

Boeing suffered heavy financial consequences. The 737 Max was the company's main source of profit. Orders for more than a thousand aircrafts were cancelled. Being

---

[3] 'Downfall: The Case Against Boeing' on Netflix.

charged with fraud, Boeing settled by paying more than 2.5 billion US dollars. The CEO was fired. Short term greedy decisions led to many lost lives. They also severely damaged the reputation and financial stability of a leading 100-year-old company. This is not what stakeholders wish for. How did the company's executives, managers, and employees fail in such way?

My answer to this question is that people to think in a selfishness-bubble. The idea is that people's interests – such as making a short-term profit – shape the way they both perceive and choose information. In turn, people use this information to make their decisions.

Our field of behavioral ethics aims to understand how incentives shape the way we think and behave. Let me share some of the things we have learned in recent years about the selfishness bubble. Let's begin with considering how incentives shape the way we perceive information, and the decisions that follow.

Here is a great exercise from our Harvard Business School colleagues Max Bazerman and Ovul Sezer.[4] Imagine you are an investment advisor. Your client has a long-term perspective and a moderate tolerance for risk. You are considering one of four funds to recommend the client. Figure 1 shows the returns for each of the funds over the last four years. The light blue line at the bottom represents the S&P 500 index tracing market performance; Here are the four funds to choose from: The light green line, is Tobacco Trade Investments; The red line, is alpha investments; The purple line is fortitude investments; And the orange line at the top is power trade investments. Please take some time to choose the fund you would recommend.
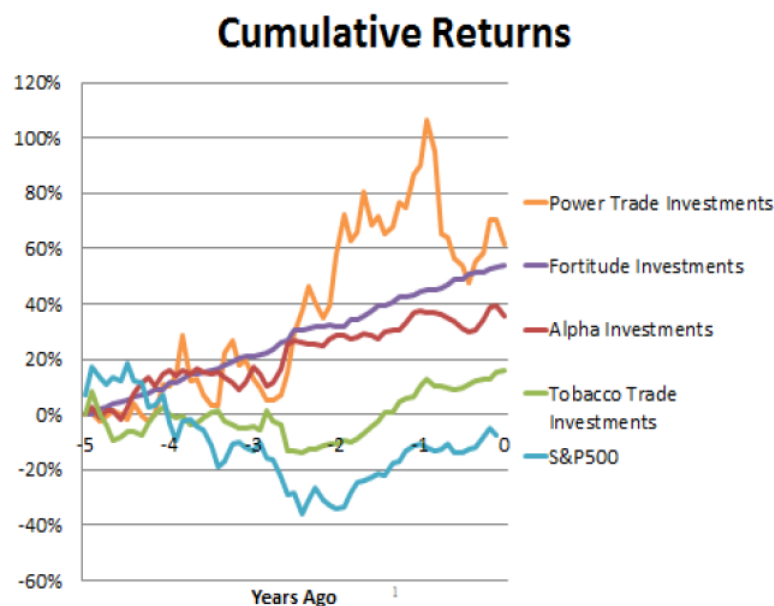
[4] Bazerman, M. H., & Sezer, O. (2016). Bounded awareness: Implications for ethical decision making. *Organizational Behavior and Human Decision Processes*, 136, 95-105.

Remember, the client has a long-term perspective and a moderate tolerance for risk. And you want this client to stay with you and recommend friends to hire you as well. So you really want a happy client. Let's vote.

Who would recommend Tobacco trade, the green line? Also our students do not choose this option. Who would recommend Alpha investments? Not so many. Who would recommend Fortitude? Many. Who would recommend Power trade? some.

The results today are very similar to those we get in class with economics students. People recommend Fortitude. It delivers a steady profit year in year out. You and our students clearly know how to serve the interests of your clients. Or do you?

Fortitude's performance here is the performance of Berny Madoff's fund, which was based on a Ponzi scheme. He had very rich people invest with him, and falsely reported to them at the end of every financial year about the profits they allegedly made. Happy investors brought their friends to invest, providing Madoff with a lot of money, which he spent on his luxurious life style. When investors asked for part of their investment so that they could go on a holiday for example, Madoff could easily give it to them as new money kept rolling in. When people started to doubt however, and demanded their money back, it quickly became clear that there was no money. The 65 billion dollars invested in Madoff's fund were all gone.

The lesson here is that it is difficult for us to spot information that is too good to be true. It is not possible to have such a steady stock performance in a noisy market environment. We know that. But this knowledge does not help. We fail to notice those things because we want happy clients that will bring us more business. A clear case of selfishness-bubble thinking.

Here is a task we used at our lab in a paper with former lab members Andrea Pittarello, Margarita Leib and Tom Gordon.[5] Participants in our experiment were seated behind a computer screen and asked to report the number nearest to the 'X' that was about to appear on their screen. In our experiment, each participant had to do the task for about 200 times, with different combinations appearing on the screen.

We had two groups of students doing this task. Both groups were instructed the same: identify the number closest to the X. But we paid the two groups according to different rules. The first group of students was paid according to their accuracy. Giving an accurate answer meant earning some money. The second group of students was paid according to the number they reported observing as closets to the 'X'. If they reported that '6' was the closest, they got 6 euros. If they reported '1' was closest they got 1 euro, and so

---

[5] Pittarello, A., Leib, M., Gordon-Hecker, T., & Shalvi, S. (2015). Justifications shape ethical blind spots. *Psychological science*, 26(6), 794-804.

on. Reporting a high number, meant earning more money. Even if the number was not really closest to the X.

In Figure 2 you can see the example from before. The Y-axis indicates the proportion of participants making a mistake by reporting the second closest value to the X. That would be '6' in the example above. Here are the results of the group paid for accuracy. They found the task easy. On the left side you see the results of trials in three difficulty levels. From low difficulty, on the left, meaning that the X was clearly closer to the target number than to the number next to it, to high difficulty, on the right, meaning that the X was closer, but not by much, to the target number than to the number next to it.
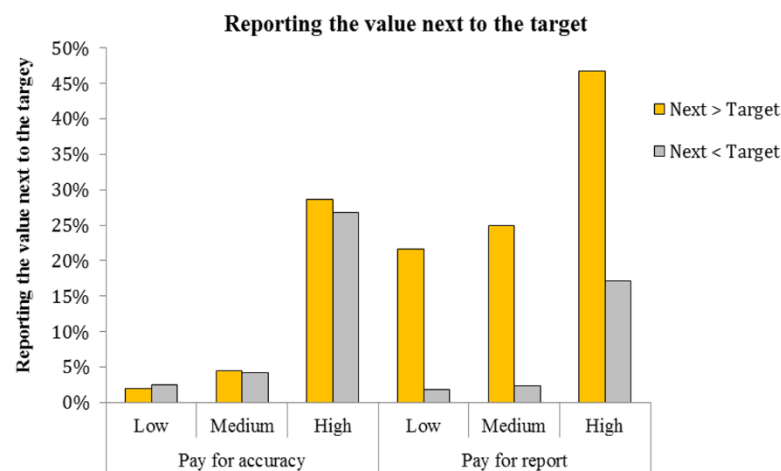
Pittarello, Leib, Gordon, & Shalvi (2015)

You can see that when people are paid for accuracy, and the task is easy, they make very few mistakes. When the task is more difficult, they make some more mistakes. But still, they made about 30% mistakes, which is better than choosing between the two numbers at random. You can further see that the value appearing next to the target, does not change the results. The orange bars represent trials in which the value next to the target

was higher than the target, whereas the gray bars represent trials in which the value next to the target was lower than the target. This made no difference to the participants who were trying to be accurate. Here are the results of the group paid for reporting high numbers. This group made many more mistakes. Not just any mistakes, self-serving mistakes. As you can see on the right side of the figure, there were more mistakes, particularly when the number second closest to the 'X' was higher than the target, the orange bars. We found much fewer mistakes when the number second closest to the X was lower than the target, the gray bars.

This study, as well as studies by Guy Hochman and others,[6] shows that when accuracy pays, mistakes are rare. When reporting high numbers pays, self-serving mistakes are much more common. We see what we want to see. A clear case of selfishness bubble thinking.

Another example. At universities, mice are often used for medical research. After the study is done, the mice are killed. Keeping them alive is too expensive. Participants at a study by Armin Falk and Nora Szech[7] were asked to make the following choice. Receive 10 euros or give up the 10 euros to save the life of one such mouse. Real money. Real mouse. What would you do?

I will not ask you to vote, but in class I do. Results are often very similar to what the researchers found, between 40% and 50% of our students choose the money over the mouse. This is already interesting, if we think about the monetary value of life.

But results had an important twist. Falk and Szech had additional market treatments. For example, in one of those market treatments participants had to negotiate with another person how to split 20 euros between the two of them. If they agree on a split, say 10 euros for each, they get the money and the mouse dies. If they do not agree on a split, or agree to give up the money, they get nothing and the mouse lives.

In the baseline, individual treatment, almost 46% of participants chose the money killing the mouse. Here is what happened in the market treatments. Now more than 70% of people took the money and killed the mouse. The increase from 46% to more than 70% of people taking the money, and killing the mouse, demonstrates how adopting a market mindset, in which financial decisions are made, makes us prioritize money over life.

Here is another study we ran at our lab to assess the extent to which people trade different norms with one another. In this case, cooperation and honesty. Arguably, people care about both things: Being cooperative and being honest. But what happens when you cannot achieve both? That is, if being cooperative requires lying. Would people be honest and less cooperative? Or would they cooperate by bending the rules, and lie?

Together with Ori Weisel, we designed task to answer this question. Participants entered the lab in groups of 20. They were seated behind individual computers. Each person was coupled with another person. But participants did not know with whom they were coupled. They only knew it was one of the other participants in the room. One participant of the two, was randomly assigned to be player A. The other participant in the pair was assigned to be player B. Player A was asked to roll a dice and type in the number he or she saw. Player B then learned what A typed in. After observing what A has reported, player B

[6] Hochman, G. U. Y., Glöckner, A., Fiedler, S., & Ayal, S. (2016). "I can see it in your eyes": Biased processing and increased arousal in dishonest responses. *Journal of Behavioral Decision Making, 29*(2-3), 322-335.
[7] Falk, A., & Szech, N. (2013). Morals and markets. *Science, 340*(6133), 707-711.

was asked to roll the die and type in the number he saw as well. If both reported the same number, they received this number in euros. If both reported different numbers, they both got nothing. After player B reported the number, Player A learned what B reported, and played again. They did that for 20 rounds. Since only participants knew the number they have rolled, they could have typed in any number. So if participants want to make money, they can cooperate and report high numbers, like 6-6. Even if in fact they saw other numbers.

Figure 3 shows what we found. The figure shows the distribution of outcomes you would expect, if people are reporting honestly. The Y-axis shows the reports of player B, the X-axis shows the reports of player A. Here you see the results of a simulation of rolling two dice and honestly reporting their outcome. As you can see, the dots are equally spread around the board. The likelihood that A saw a 1 and B a 3, a result that would lead them to earn nothing, is the same as the likelihood that A saw 6 and B saw 6, a result that would lead them to earn 6 euros each. The likelihood of both players to observe the same number, that is to roll a double, is 1 in 6. Player A reports a number, which B needs to hit. B has a 1 in 6 chance to roll this number. What is the likelihood to roll a double 6? Player A has a 1 in 6 chance to roll 6. Player B as well. Multiple the two probabilities and you get the likelihood: 1 in 36. Less than 3%.


Leib, Köbis, Soraperra, Weisel, & Shalvi (2021)

Now here is what happens when you let 20 pairs of students, play the game for 20 times. Statistics are not needed here. It is rather clear that participants are lying. Most reports are on the diagonal, and especially, on the most profitable 6-6 option. One of the things I like most about experiments, is the fact that they tell us something fundamental about how people think and decide.

A common critique of experiments is however, how do we know that the observed effect is robust? Would it replicate if we try again? We recently provided some answers to those questions in a paper[8] with Margarita Leib, Nils Kobis, Ivan Soraperra and Ori Weisel. We conducted a meta-analysis. This means we aggregated the results of all studies we could find who use the task I just described. We have managed to obtain data of more than 2,700 participants, making more than 50,000 decisions.

But how did we find these data? One thing we did was to systematically search within published literature. This is a good starting point, but focusing only on published work is tricky. The reason is that people who run a study but do not find the effect may consider the study 'unsuccessful' and not even try to publish it. That is called the file-drawer problem.

To verify we provide the best overview, we reached out to the scientific community in calls we have sent via various mailing lists. People sent us their unpublished work including dissertations and master's thesis work. In our final sample, 35% of the studies are unpublished. And the results of the published and unpublished work are similar. What did we find? Rather similar results to the original study. People report doubles more often that you would expect, especially the profitable 6-6 option.

In a way, our students behave rather similar to the Boeing employees from our earlier example. Both violated ethical rules, to increase profit. Participants, by reporting higher numbers than they actually saw; Boeing employees, by suggesting re-training pilots was not needed when in fact it was.

The selfishness bubble mind set is not only about how we perceive information, but also about the information we choose to expose ourselves to. Do we really want to know where and how the clothes we wear were made? Do we want to know the conditions that farm animals producing milk, cheese, and eggs live in? Often we prefer to remain willfully ignorant to such information.

Figure 4 shows the task, our colleague Jason Dana, has used together with others to study willful ignorance.[9] In their baseline treatment, Jason and his colleagues asked participants to choose one of two options: A or B. As you can see in the picture, the choice influenced their own outcome but also the outcome of another anonymous person. If participants chose option A they received 6 American Dollars, while the other received 1 dollar only. If they chose B, both they and the other received 5 Dollars. So participants had to decide whether to give up 1 dollar, going from 6 to 5 Dollars, so that the other person, will get 4 dollars more.

[8] Leib, M., Köbis, N., Soraperra, I., Weisel, O., & Shalvi, S. (2021). Collaborative dishonesty: A meta-analytic review. *Psychological Bulletin, 147*(12), 1241.

[9] Dana, J., Weber, R. A., & Kuang, J. X. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33(1), 67-80.

Try to think what you would do. You do not know the other person, you will never know who she or he is, they will only receive the amount you chose to give them.

In this situation, 74% of people chose the altruistic option, B.
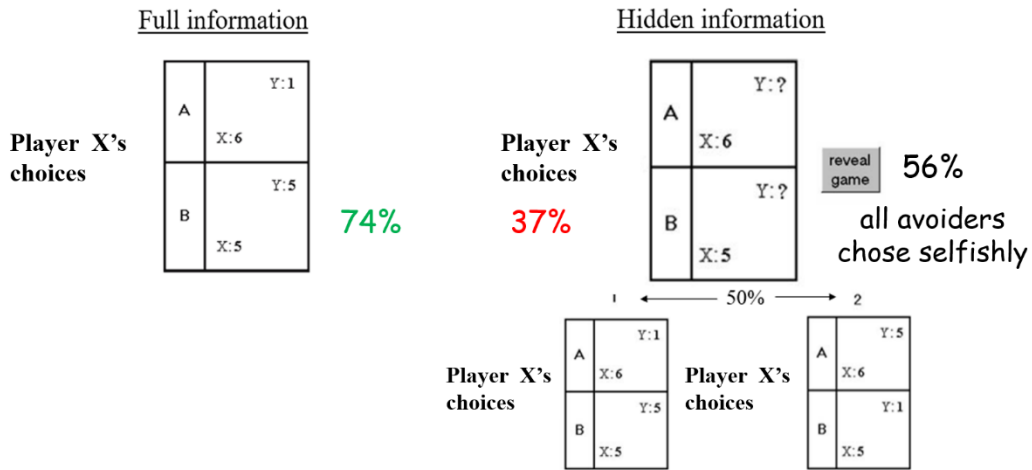
Dana, Weber, & Kuang (2007)

There was a twist, of course. In another setting, the experimental treatment, participants were asked again, to choose between two options: A and B. The amounts of money they would receive from both options were identical to those in the control treatment: $6 if they choose A and $5 if they choose B. But here, participants did not know what the other would receive. With 50% probability the other would receive the same amounts as before, $1 if participants chose A and $5 if participants chose B. With 50% probability, however, the other would receive $5 if participants chose A, and $1 if participants chose B.

Importantly, participants could learn the outcomes of the other person. All they had to do was click the 'reveal' button. And this is rather useful information. If we are in a situation in which choosing A leads both the participant and the other person to receive more money, it is a no brainer. You choose A. Self-interest and being nice go hand in hand.

But what if revealing the information leads you to learn that choosing A gives you one more dollar but reduces the other's outcome by $4? Now it does not feel so good perhaps to choose A. In this setting, people can engage in willful ignorance. They can choose not to reveal the consequences of their actions. Not knowing allows them to feel good, while being selfish.

Dana and colleagues found that when information about the other's outcomes had to be revealed, only 37% of participants chose the altruistic option, B. The reason for the drop from 74% of people choosing altruistically to only 37%, was that more than half of participants chose not to reveal information. And all of those who chose not to know, chose the selfish option. Those people did not know what the consequences of their choices were, so they chose selfishly, without feeling bad about it.

You may again ask yourself, how do we know that the observed effect is robust? Would it replicate if we try again? And what happens if we try to run the study in a different country?

We recently provided some answers to those questions in a paper[10] with Linh Vu, Joel van der Weele, and Ivan Soraperra all from our group here at CREED – the Center for Research in Experimental Economics and Decision Making. We conducted a meta-analysis. This time we aggregated the results of all studies we could find who use the willful ignorance task.

The list on the left shows all the studies we found. In this list, 36% of the studies were unpublished. And again, the results of the published and unpublished work are similar. What did we find? In yellow, you can see the original study results. the green dot represents the percentage of altruistic choices in the baseline treatment, 74% as you may remember.

The red dot represent the percentage of altruistic choices in the experimental treatment, 37%.

As you can see, across many studies, with more than 6,500 participants, the effect is robust and clear. When people can justify their selfishness by ignorance, they do. We found that close to 40% of participants avoid information. As a result, altruism drops by 16% on average. You can see this average difference at the very bottom of the figure. We further find that there is not much of a difference when considering the country in which the study was conducted.

But to be fair, our data comes from only six countries. And all are western, educated, industrialized, rich, and democratic. Also known as WEIRD countries.[11] This is why together with Catherine Molho, Ivan Soraperra, and Jonathan Schulz, we now study willful ignorance around the globe.[12] Including countries like Egypt, Indonesia, Nigeria, Peru, and more.

I like this paper also because we use a rather new academic publishing format, the registered report. In this type of publishing, we submit our plan to be evaluated by our

[10] Vu, L., Soraperra, I., Leib, M., van der Weele, J., & Shalvi, S. (2022). Willful Ignorance: A Meta Analytic Review, a working paper.

[11] Henrich, J., Heine, S. J., & Norenzayan, A. (2010). Most people are not WEIRD. *Nature*, 466(7302), 29-29.

[12] Molho, C., Soraperra, I., Schulz, J., & Shalvi, S., (2022). Guilt- and Shame-Driven Prosociality Across Societies, a working paper.

colleagues before running the study and analyzing the data. Once the editor gives us the 'good to go' signal, we collect the data and report the results. The paper is then published regardless of emerging pattern of results. The format nicely removes a lot of the inherent biases academic publishing may have. Such as having to report a sensational finding to get into a top journal.

The studies I talked about today, demonstrate how the selfishness bubble works. The selfishness bubble leads people to view the world in a self-serving way. To make mistakes they benefit from. To value life differently. To ignore inconvenient information. The question however, remines, what should we do about it?

My answer is: think long term and fix the system, not the people. Let me explain. 15 years ago, a bestseller book called Nudge was published. [13] Even if you have not heard of the book, you may still have heard of the term: nudging.

The authors, Richard Thaler and Cas Sunstein, made an intriguing proposal; to change people's behavior, all you need to do is redesign the settings in which they make their choices. Instead of placing fried food and soft drinks on eye levels shelves, place vegetables and water. Because people buy more things that are on their eye level, re-organizing will lead them to buy healthier food. It even gets better, people still have the choice of buying the unhealthy food, if they really want to. Paternalistic set-up coupled with free choice.

In most cases, nudge solutions are cheap and do not upset people. Policy makers love this idea. Social scientists in turn, received funding to figure out the best ways to cheaply nudge people to eat healthier food, get vaccinated, save electricity, and so on.

There are voices however, suggesting that by focusing on ways to nudge people, we liberate governments to do little in terms of making long-term changes and fixing the system. In a recent paper, Nick Chater and George Lowenstein remind[14] us, that nudges and nudge-based short-term campaigns, should not replace long-term regulation change.

Consider for example, trying to reduce the consumption of sugary drinks. Here in the Netherlands, the government encourages eating healthy by providing eating recommendation on its website and regulating the information on food packages. Not bad measures, but relatively soft ones. Nudging like that is easy for the public to accept. Imposing a Sugar tax in contrast is unpopular, but more likely to reduce consumption.

Another example, reducing the use of short-distance flights. Recently, KLM – the Dutch aviation company – issued a campaign suggesting people should 'fly responsibly' and

[13] Thaler, R. H., & Sunstein, C. R. (2009). Nudge: Improving decisions about health, wealth, and happiness. Penguin.
[14] Chater, N., & Loewenstein, G. (2022). The i-frame and the s-frame: How focusing on the individual-level solutions has led behavioral public policy astray. Available at SSRN 4046264.

consider replacing flying with zoom meetings.[15] This is a soft nudge, that may have some impact. But I doubt KLM are taking a big financial risk by nudging people to fly less. Compare this soft nudge with a flight-ban on short-distance flights, or increasing aviation fuel tax. Unpopular, but more likely to have impact.

Unlike short-term nudge campaigns, some nudges do lead to changing regulation. Those are the ones, I believe we should be focusing our efforts on. The best example, for a nudge that informed a major, long-term policy change, is organ donations.

In 2003 Johnson and Goldstein[16] discovered that in countries where people have to sign up to become organ donors after they die, donations are rather rare. In the countries they looked at, less than 30% of the people were listed as organs donors. In contrast, in countries where people become donors automatically after they die, unless they opt-out, donations were more common. In those countries, more than 85% of the people were listed as organ donors. The reason for the gap is that people do not like to think about what will happen to them after they die. So they take no action and passively chose the default option.

Based on these finding in 2020, the Netherlands changed its default from opt-in to opt-out. Whereas in the past people had to sign up to become an organ donor, now they are donors unless they choose not to be. One year after the regulation was introduced, the Netherlands, a country with a population of 17,5 million people, saw an increase of 4,6 million people now registered as organ donors.[17]

When basic science about human behavior is used to informs policy, impact can be achieved.

What ties the Boeing and organ donation examples together, is that for success, people need to adopt a long-term perspective. If Boeing executives were not so focused on quick returns and instead focused on the long-term perspective, they would have likely developed better aircrafts, increasing future revenues. When it comes to organ donations, the Dutch government used scientific insights to change regulation, rather than nudging individual people or trying to educate them.

The government's long term solution, achieved real, lasting impact.

And this brings me back to Milton Friedman and the social responsibility of business.

But before Friedman, let me thank some people.

---

[15] https://flyresponsibly.klm.com/gb_en#home
[16] Johnson, E. J., & Goldstein, D. (2003). Do defaults save lives?. Science, 302(5649), 1338-1339.
[17] https://www.donorregister.nl/actueel/nieuws/2021/07/07/door-nieuwe-donorwet-staan-bijna-11-miljoen-mensen-met-een-eigen-keuze-in-het-register

I thank the Executive Board of the University of Amsterdam for the trust placed in me and the decision to establish the new chair in Behavioral Ethics.

I would also like to thank the current faculty Dean, Roel Beetsma, and the former faculty Dean Han van Dissel, for nominating me for the position.

Thank you all my mentors, Carsten, Michel, Ilana, Jason, Maurice. The extended CREEDers in past and present, and the behavioral ethics lab members, for the pleasure of working together.

Thanks to our amazing administrative team, led by Wilma. You are allowing us to do our work, that is priceless.

Of course, I would not be standing here today without the support of my family. My kids, Daniël, Ruben, and Ayelet and their wonderful mother, Marieke. Love you! And of course my parents Judith and Micha and sisters, Netta and Tamar, who supported me in every step, including moving to the Netherlands. Thank you!

Thank you all for joining us here today as well as those watching from home. Those include my grandmother, Alice 95 years old, watching live from Jerusalem. Love you savta!

Now, let me close where we started. Milton Friedman wrote that "the social responsibility of business is to increase its profits". But when people hear profit, they often think of short-term profit.

So I propose that the social responsibility of business is to encourage a long-term perspective. This is not just semantics. Words shape mind-sets. Mind-sets shape behavior.

By encouraging a long-term perspective, executives will focus on building the company's reputation not only on their current balance sheet. Our companies can go back to producing products they can be proud of. And make money.

But to do so, they must first, burst the selfishness bubble.

Ik heb gezegd