

# Harry Potter and the welfare of the willfully blinded\*

Felix Bierbrauer<sup>†</sup>

January 14, 2019

## Abstract

This paper presents a selective review of welfare economics. It is first argued that welfare analysts need to turn a blind eye to various aspects of individual preferences. Otherwise applications of welfare economics yield repugnant conclusions. The problem is first illustrated with characters from Hogwarts and then related to the theory of optimal taxation. We also look at individual decisions to ignore relevant information and discuss some of its welfare implications. Finally, we discuss the suppression of information that would otherwise affect the behavior of others. Such acts may be in conflict with liberal values. In the presence of behavioral biases, they may still come with positive effects on welfare. The logic is akin to the theory of the second best due to [Lipsey and Lancaster \(1956\)](#).

*Keywords:* Welfare economics; Deliberate ignorance; Optimal taxation.

---

\*Background paper prepared for the Ernst Strüngmann forum on deliberate ignorance. I benefitted from conversations with Carina Bierbrauer, Martin Hellwig and Francois Maniquet.

<sup>†</sup>CMR - Center for Macroeconomic Research, University of Cologne, Albert-Magnus Platz, 50923 Köln, Germany. E-mail: [bierbrauer@wiso.uni-koeln.de](mailto:bierbrauer@wiso.uni-koeln.de)

# 1 Introduction

We consider a decision that affects the well-being of several individuals. Decision-relevant information is available. Before dealing with the decision problem of interest, there is another decision to be taken, namely whether to use this information or to take the decision under ignorance.<sup>1</sup>

To fix ideas, consider the following problem: A cake of given size has to be divided between Harry and Draco. Albus, a benevolent planner chooses the division. He contemplates an application of utilitarian principles. If both Harry and Draco were selfish, representing their preferences by the same concave utility function and maximizing a sum of utilities gives rise to an equal split of the cake. At first glance, this seems to be an appropriate outcome. Harry, however, has altruistic feelings and derives utility also from every piece received by Draco. Taking these feelings into account implies that Albus should assign a larger share to Draco. Deviating from a fifty-fifty split in this way has an opportunity cost, Harry's forgone utility as he is eating less, and a welfare gain, Draco's extra utility from eating more plus Harry's extra utility from Draco's extra utility. The latter implies that the welfare gain dominates. The conclusion that Draco should receive a larger share is, moreover, reinforced by Draco's sensations of envy which imply that every piece assigned to Harry reduces Draco's utility by more than just his forgone consumption utility. Thus, a consequence of utilitarianism seems to be that Harry is punished for his altruism and Draco is rewarded for his envy. Albus thinks twice. What information about preferences and utilities should be taken into account? What information should be ignored?

The example illustrates the possibility that taking account of information on preferences makes it possible to achieve higher welfare levels, while the consequences seem repugnant. More generally the question is what types of information welfare analysis should be responsive to and to what types of information welfare analysis should turn a blind eye. In the first part of this note we will discuss this problem at a broad conceptual level. The second part deals more specifically with the welfare implications of information acquisition and information avoidance by individuals. To what extent are individual choices in this regard aligned with social welfare? To what

---

<sup>1</sup>To be clear, in our discussion of this problem, we are not interested in a trade-off of the following sort: The information, if available, would improve the collective decision. Acquiring it is costly, however. A cost-benefit analysis therefore has to strike a balance between improved decision quality and the costs of information acquisition. Therefore, we will assume throughout that information is available at no cost.

extent would a welfare maximizer want to interfere with individual choices?

## 2 Blinding the welfare analyst

**What should be the domain of welfare analysis?** Let us begin with an example, due to [Coase \(1960\)](#), in which an application of welfare economics seems uncontroversial. A fishery and a chemical plant reside along the same river. The chemical plant resides upstream and pollutes the water of the river. This has a negative impact on the fish population and reduces the return to the fishery. This classical example is discussed in textbook treatments of the market failures that arise in the presence of externalities. Under *laissez faire*, the chemical plant does not take into account that its activities have negative consequences elsewhere. The scale of its operation is then too large from a welfare perspective: A reduced scale in combination with a monetary compensation for the forgone profits would make both the chemical plant and the fishery better off.

An alternative example by [Sen \(1970\)](#) involves a person's decision whether or not to read "Lady Chatterley", a novel with explicit accounts of sexual actions. Another person is a prude and feels that no one should read such a book. Thus, if the first person reads the book, this has adverse consequences for the second. Does such a negative externality warrant the same treatment as the one in the relation between the chemical plant and the fishery? The logic of the latter suggests that censorship by the second person, in combination with a compensation for being censored to the first person, would make both better off.

Sen uses this example to illustrate a conflict between the principles of welfare economics and liberal values that arises as soon as individuals have preferences over the private choices of others, i.e. what books to read, how to dress, whom to meet, what opinions to express etc. Liberal principles require that such choices are respected. A stubborn application of welfare economics, by contrast, suggests that such choices should be corrected or moderated, in return for compensation.

[Goodin \(1995\)](#) argues that the preferences which are used in welfare economics should go through a process of "laundering". Goodin is concerned with perverse or sadistic preferences such as Draco's sensations of envy in the introductory example. Welfare economics would be misguided if it took such preferences seriously. Draco would then be rewarded for his envy by receiving a bigger chunk of the cake than

Harry. In Goodin's view, the censoring of the preferences that are fed into welfare analysis does not require a paternalistic justification. It can often be justified by distinguishing what people really want, their *true preferences*, and the preferences that seem to be revealed by their choices, their *revealed preferences*. Revealed preferences may be shaped by temptations, short-term desires or other sensations. Individuals might not want that value judgements reflect these sensations. In this case, there is a discrepancy between the normatively relevant true preferences and the revealed preferences.

Note that the laundering of preferences can also be applied to deal with Sen's conflict of liberal values and welfare economics. The laundering in this case would have to remove preferences over the private choices of others. A liberal prude would admit that he feels annoyed by a fellow's reading of "Lady Chatterley" while not wanting that this sensation is used as a justification for the interference with private choices. An illiberal prude might disagree. But there is no conflict between liberal values and welfare economics provided that, for the purposes of welfare analysis, all liberals agree that their preferences should be laundered from attitudes towards the private choices of others.

The preceding discussion illustrates the question what types of preferences one should allow for in welfare analysis. This is a normative question that cannot be answered by an application welfare analysis itself. Laundering preferences from sensations of altruism or envy may be a way to avoid repugnant conclusions, such as punishing Harry for his altruism and rewarding Draco for his envy. Dismissing preferences over the private choices of others is a way to avoid illiberal conclusions, such as the censoring of books.

From the perspective of applied welfare economics, the *deliberate ignorance* of perverse or illiberal preferences may come with a cost. The welfare measure used in applied analysis would possibly take a higher value if dirty preferences were allowed for. The following examples from the welfare analysis of tax policy illustrate this point.

**Welfare economics of taxation.** In applied welfare analysis, there is a set of outcomes and individuals have preferences over these outcomes. The outcomes and the preferences are the primitives of the problem. The problem then is to find the "right" outcome. Often this is taken to be the outcome that maximizes a utilitarian welfare measure. A more cautious approach – that avoids interpersonal comparisons

of utility – identifies a whole set of “right” outcomes, typically a set of Pareto optima.<sup>2</sup> A field that makes heavy use of this framework is the analysis of tax policy in public finance. To give a feel for the relevance of the preceding discussion in applied work, we discuss various examples from this line of work.

The theory of optimal taxation applies welfare economics to the study of tax policy. The basic ingredients are a government that uses tax policy to generate revenues and consumers who choose how much to consume, how much labor to supply, or how much to save. The choices of consumers are affected by the tax policy. A labor income tax affects the incentives to supply labor, a tax on capital income affects the return on savings. An optimal policy maximizes a utilitarian welfare objective taking these behavioral responses of consumers into account.<sup>3</sup> A well-known result is that taxes should follow an inverse elasticities rule: taxes should be high when behavioral responses, usually measured by the price-elasticity of supply and demand, are low and vice versa.

It is an intuitive finding. If capital income is shifted abroad in response to taxation but labor income is not, then the tax on capital income should not be as high as the tax on labor income. If the demand for necessities such as bread or gas is less price sensitive than the demand for luxuries, then the tax on bread should be higher than the tax on champagne. These examples raise distributive questions, but even if those are taken into account, the logic of the inverse elasticity rule remains intact, albeit in a *ceteris paribus* sense: When two goods are consumed equally by the rich and the poor, the one with the lower elasticity of demand should be taxed at a higher rate, see e.g. [Diamond \(1975\)](#).

In any case, such taxes interfere with the private choices of individuals. If the demand for books such as “Lady Chatterley” or the “Satanic Verses” was less price sensitive than the demand for more respectable types of literature such as Hamlet or the Koran, the logic of optimal tax theory would suggest to have higher taxes on the former and lower taxes on the latter. Hence, another round of laundering may be needed to avoid repugnant or illiberal conclusions from the application of optimal tax theory. The consequence of such laundering, however, is that the resulting tax system is not optimal from the perspective of a welfare measure that is based on the

---

<sup>2</sup>The defining property of a Pareto optimal outcome is that moving away from it necessarily makes some people worse off.

<sup>3</sup>The modern analysis of this problem dates back to [Ramsey \(1927\)](#). A rich a body of literature has refined this approach in various ways. Seminal contributions are [Mirrlees \(1971\)](#), [Atkinson and Stiglitz \(1976\)](#), [Diamond \(1998\)](#) and [Saez \(2001\)](#).

preferences that individuals reveal through their market behavior. Like an insurance company that loses profit when non-discrimination requirements remove a possibility to have different premia for men and women, a welfare maximizer then has to live with the fact that the laundering prevents her from reaching welfare levels that would otherwise be attainable.

From the perspective of practical tax policy, having different tax rates for different types of books is a contrived example. There are, however, more plausible implications of optimal tax theory that raise similar issues. There is a rich literature on the optimal taxation of couples, see [Boskin and Sheshinski \(1983\)](#) or [Kleven et al. \(2009\)](#) for seminal references. In these studies, it is shown that a differential taxation of the primary and the secondary earner in a couple is desirable from a welfare perspective. This finding combines the logic of the inverse elasticities rule with the empirical observation that the labour supply of females, who are more often in the role of the secondary earner, is more tax sensitive than the labour supply of males. Thus, an optimal tax system should apply different tax rates to the primary and the secondary earner in a couple. In particular, income due to the secondary earner should be taxed at a lower rate.

Obviously, such differential taxation interferes with a private choice. It affects the assignment of roles in a couple, in particular the decision who should contribute how much to family income. With a progressive income tax system, income splitting<sup>4</sup> is the only possibility to have a couple's tax burden only depend on its overall income, irrespectively of who contributed how much. Hence, an attempt not to interfere with private choices implies that the inverse elasticities logic is not applied, with the consequence that potential welfare gains remain on the table, and with the consequence, that female labour supply is discouraged more than it would otherwise be.

The treatment of altruism plays a prominent role in the theory of capital and inheritance taxation. Results on the desirability of capital taxes crucially depend on assumptions about the altruism of parents towards their children. In the analysis of [Farhi and Werning \(2010\)](#) altruism implies that a bequest is a source of utility both for parents and children. A bequest subsidy is warranted to make sure that the positive externalities from leaving a bequest are taken into account. In [Piketty and Saez \(2013\)](#), by contrast, the degree of altruism varies from generation to generation

---

<sup>4</sup>Let  $y^p$  be the income of the primary earner and  $y^s$  the income of the secondary earner, under income splitting the tax burden of a couple is  $2 T\left(\frac{y^p+y^s}{2}\right)$ , where  $T$  gives the tax burden as a function of income.

in an unpredictable way. This is shown to imply that a redistributive taxation of bequests is desirable. The redistribution from lucky children with high bequests to unlucky ones is shown to be part of a welfare-maximizing policy. Needless to say that such analyses only make sense on the assumption that altruism is a legitimate ingredient of welfare analysis. The conclusions on the desirability of bequest taxes and subsidies would not survive a laundering of preferences from altruism.

These examples demonstrate the difficulty of the question what preferences to feed into welfare analysis. A naive use of revealed preferences may give rise to repugnant or illiberal conclusions. Forcing the welfare-maximizer to turn a blind eye to the dirty or the private aspects of individual preferences may be an appropriate remedy. As we have seen in our discussion of tax policy, the extent to which this is done can have drastic implications for the policy implications of “applied work.”

**Internalities.** The Coasian example that involves the chemical plant and the fishery is one of *externalities*. One firm pursues its economic interests at the expense of another one. Welfare economics stipulates that such externalities must be taken into account, and moreover, that doing so in an appropriate way, would make both firms better off.

This logic has been extended to deal with the internal conflicts that individuals may have. Self control problems are a prominent example. An individual may have the long run goal to lead a healthy life. In the short run, the individual is confronted with temptations such as drinking alcohol, smoking cigarettes or eating unhealthy food. Giving in to such temptations damages the individual’s long-run goals. The literature often refers to such self-damaging behavior as creating *internalities*. Applied work in optimal tax policy has discussed corrective taxes that deal with such internalities. For instance, [O’Donoghue and Rabin \(2006\)](#) characterize “optimal sin taxes” that mitigate self-damaging behavior.

Public policy that addresses internalities interferes with the private choices of individuals. This raises the question whether it provokes the type of conflict between liberal values and welfare economics that is illustrated by Sen’s “Lady Chatterley”-example. In that example, a problem arises as one person has preferences over the private choices of another person. Here, the public policy maker has preferences over the life-style of individuals such as their drinking and smoking habits. Isn’t this the same kind of problem?

This question has spurred controversies, see [Gul and Pesendorfer \(2008\)](#) and

Loewenstein and Haisley (2008). For the proponents of such policies, the answer is clearly “no”, provided that the policy maker does not pursue an own agenda, but has preferences that are aligned with the individuals’ long term goals. The agenda on “soft” or “liberal” paternalism, see Thaler and Sunstein (2008) for the best known example, focusses on situations where individuals can be enabled to behave in accordance with their long-run goals, without harming others who do not suffer from the same kind of self-control problem.

In line with this program, O’Donoghue and Rabin (2006) look at a population of smokers who differ in the intensity of their self control problems. Some are heavy smokers and have pronounced self-control problems, others are occasional smokers and have self-control problems that are not as severe. The authors show that a sin tax may, nevertheless, make all smokers better off. The heavy smokers like the tax as it helps with the self-control problem. The revenue that is generated can be used to compensate the light smokers who would otherwise be harmed by the sin tax.

The analysis also points to the limits of liberal paternalism. A Pareto-improving sin tax is possible only if there is a one-to-one relation between the number of cigarettes smoked and the intensity of the self-control problem. If one introduces heavy smokers with no self control problems to the system, the possibility of a Pareto improving sin tax is gone. In this case, one has to make rational smokers worse off when one attempts to help the smokers with self-control problems. From the perspective of the rational smoker, this is akin to an illiberal interference with a private choice.

### 3 Blinding oneself

We now turn from the question what information a welfare maximizer *should* ignore, to the information that individuals *do* ignore. In the subsequent section, we will then turn to the deliberate suppression of information that is relevant for the decisions of others.

A strand of behavioral research has investigated circumstances under which individuals take decisions while deliberately ignoring decision-relevant information. This literature is surveyed by Golman et al. (2017) and Hertwig and Engel (2016). Our focus will be on the welfare implications of such information avoidance. We will go through some prominent examples and discuss the criteria developed in the previous



section:

1. When a person engages in information avoidance, does this give rise to externalities (positive or negative effects on others) or internalities (positive or negative effects on her- or himself)?
2. If so, are mechanisms in place that ensure that these are taken into account?
3. If not, is information avoidance a private choice that should not be the subject of welfare analysis?
4. If not, are “dirty” preferences at play that should be removed from welfare analysis?

**Climate change denial.** Climate change is one of the most pressing problems of mankind. It is also a prime example of an externalities problem. The CO<sub>2</sub> emissions of past and current generations have drastic consequences for younger generations. Classical welfare economics stipulates to address such problems using corrective taxation or quantity controls. Yet, taking measures which are really effective at mitigating climate change is politically controversial. Some opponents of such policies even deny the existence of a problem that needs to be solved: There is deliberate ignorance of the overwhelming evidence on climate change.

The Coasian example with the fishery and the chemical plant rests on the premise that all accept the description of the problem. The calculus of welfare economics is applicable only if both parties agree that the emissions of the chemical plant are harmful to the fishery. If, say, the fishery denies this, there is no point in having a cost-benefit analysis determine the optimal level of emissions reduction, accompanied by compensation payments that ensure that both parties are better off. Thus, a deliberate ignorance of the harm that is caused by emissions is equivalent to denying that policies to address this problem can be justified with an appeal to welfare.<sup>5</sup>

This raises the question whether welfare economics can be applied in the face of such deliberate ignorance. From the perspective of those who deny climate change, public policies that seek to address it are unjustified and paternalistic. Should their

---

<sup>5</sup>According to the impossibility result by [Myerson and Satterthwaite \(1983\)](#), efficient Coasian bargaining is not possible if the intensity of the externalities problem cannot be objectively verified. If only the fishery knows how much harm is caused by emissions and only the chemical plant knows how costly it is to avoid them, efficiency is out of reach.

welfare still be taken into account when such policies are evaluated? If so, which preferences should enter the cost-benefit calculation, the preferences that are articulated in the political process, or a laundered version that no longer contains traces of deliberate ignorance?

**Ignorance of performance evaluations.** Climate change denial has a political motivation. As another example, consider a teacher who does not want to look at teaching evaluations, for fear of a negative outcome. Here, the motivation is more personal, the desire to keep a positive self-image. Still, there are externalities. A teacher's reluctance to look into ways of improving his teaching is harmful for his students.

[Golman et al. \(2017\)](#) argue that the hedonic consequences of information avoidance should be taken into account. This concrete example raises the question whether the teacher's hedonic utility from keeping a positive self-image should be weighed against the students' benefits from improved teaching. The alternative perspective is that a desire to keep an unjustified positive self image is a "dirty" preference that should be removed from a cost-benefit analysis of additional training.

**Reluctance to test for diseases.** [Hertwig and Engel \(2016\)](#) report the case of James Watson who had his genome sequenced but chose to remain ignorant about his predisposition for Alzheimer's disease. This is an example of private choice, i.e. of a choice that affects the welfare of only one person, or at least should be treated as such. Remember the lesson from Sen's "Lady Chatterley" example: Preferences over the private choices of others have to be removed from welfare analysis. Otherwise, liberal principles are in conflict with welfare analysis. If James Watson's decision is not treated as a private choice, what else should be? If the disease was a infectious and if the risk of infecting others could be reduced, e.g. by a vaccination, the conclusion would of course be different. In this case, externalities would enter the picture. A welfare analysis that weighs the personal costs of acquiring unpleasant information against the health risks of others might appear quite reasonable.

## 4 Blinding others

The previous examples involve individuals who blind themselves, possibly with negative consequences for others, e.g. for those who want to fight climate change or the

student's who would benefit from better teaching. We know turn to the deliberate blinding of others.

**Motivated beliefs.** Bénabou and Tirole (2006) present a model of motivated beliefs in which individuals suppress unfavorable information to handle cognitive dissonances. Specifically, individuals have a desire to believe that the world is just, that those who work hard or invest in their human capital can reap the rewards and become financially better off. At the same time individuals are confronted with the evidence that social mobility is imperfect, that economic inequality tends to persist over generations and that hard work doesn't necessarily pay off. There is evidence that individuals bias their perceptions of social mobility against this evidence and instead remain overly optimistic. They stick to the American dream despite the facts that point to the contrary, see also Alesina et al. (2018).

In the model of Benabou and Tirole the suppression of this unfavorable evidence has a benefit. It keeps individuals going. They invest more in human capital than they would otherwise. The positive effect is due to the assumption that individuals also suffer from a present bias. Educational effort therefore tends to be inefficiently low. Individuals give too much weight to the immediate costs of acquiring human capital and too little weight to the higher future income that results from the investment. A suppression of unfavorable information on the returns to education can therefore mitigate the individuals' tendency to procrastinate.

In their preferred interpretation of the model, Benabou and Tirole take an intergenerational perspective. Parents tell their children about the returns to effort. The children in turn choose how much effort to exert when going to school. Thus, the parents shield their kids from unfavorable information on the returns to effort in an attempt to overcome their laziness. What are the welfare implications of these choices? Are the parents doing harm to their kids? The answer would be "yes" if there was no present bias. In this case, children who become victims of their parents' propaganda would invest more than is in their own interest. With the present bias, however, the parents' indoctrination may be regarded as second-best alternative so that the children are better off with it than without.

The example illustrates a more general lesson from what is known as the *theory of the second-best*, see Lipsey and Lancaster (1956). With distortions already in place, adding another distortion may have beneficial effects for welfare. A welfare analysis of deliberate ignorance might therefore be misguided if it focusses on only one type

of deliberate ignorance in isolation. Discovering that deliberate ignorance serves a useful purpose may require to evaluate it against the background of the whole menu of individual biases.

**Manipulating the salience of taxes.** Positive welfare effects of blinding others have also been documented in the context of tax policy. [Chetty et al. \(2009\)](#) report on a field experiment that involves a manipulation of the price tags in US supermarkets. The standard are price tags that do not include sales taxes. The manipulated price tags highlighted, in addition, the tax inclusive price. It was found that the manipulation triggers a behavioral response: fewer items are sold. The authors document that consumers are well informed about sales taxes. The price tag therefore did not provide new information, it only made available information more visible. It was shown that this visibility has consequences for demand. Consumers buy more if the information on taxes is suppressed.

The conventional perspective in public finance is that any sales tax has an efficiency cost. Such a tax drives a wedge between the prices paid by consumers and the prices received by producers. As a consequence, gains from trade are not exhausted. A consumer who is willing to pay ten but faces a tax inclusive price of eleven will not buy. If producers are willing to sell for nine, there are gains from trade between the producer and the consumer. Those gains would be realized if there was no tax, but not with the tax. The forgone benefits of such transactions constitute the efficiency costs of taxes.

How is this logic affected by the behavioral responses to the salience of taxes? [Chetty et al. \(2009\)](#) assume that the demand with tax inclusive prices reflects true preferences. Thus, individuals overconsume when taxes are not salient. This overconsumption in turn helps to mitigate the efficiency costs of taxation. This is another instance of a second-best logic, one that combines a behavioral bias with an inefficiency that also prevails with rational agents, the distortionary effects of taxation.<sup>6</sup>

---

<sup>6</sup>The analysis is, however, sensitive to the assumption that individuals overconsume when taxes are not salient. Consider the alternative assumption that true preferences correspond to the demand that is observed when consumers see the price tags they are used to. In this case, making taxes more salient will aggravate the tax distortions.

## 5 Concluding remarks

In its first part, this paper has revisited classical welfare economics. It highlighted that any application of it requires a principled decision of what type of preferences to feed into the analysis. Taking account of sensations of envy or altruism may give rise to repugnant conclusions. An incorporation of preferences over the private choices of others will clash with liberal values. Thus, to be relevant, welfare analysis needs to turn a blind eye to certain aspects of individual preferences.

The second part discussed decisions of individuals to ignore information that is readily available, or to suppress information that would otherwise be available to others. It was emphasized that a welfare analysis of such choices runs into the difficulty of delineating the proper domain of welfare economics: Should welfare analysis take the preferences of those who deny climate change into account or ignore them? Should genetic tests for health risks be treated as a private affaire that is not subjected to welfare analysis?

An interesting line of recent research looks at related questions from an empirical perspective. It tries to elicit the preferences that individuals want to be fed into welfare analysis. For instance, [Weinzierl \(2017\)](#) reports that individuals demand a laundering of preferences from sensations of envy.<sup>7</sup> This avenue might prove useful for further research on the welfare implications of deliberate ignorance.

## References

**Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso**, “Intergenerational Mobility and Preferences for Redistribution,” *American Economic Review*, February 2018, *108* (2), 521–54.

**Atkinson, Anthony and Joseph Stiglitz**, “The Design of Tax Structure: Direct versus Indirect Taxation,” *Journal of Public Economics*, 1976, *1*, 55–75.

**Bénabou, Roland and Jean Tirole**, “Belief in a just world and redistributive politics,” *Quarterly Journal of Economics*, 2006, *121* (2), 699–746.

---

<sup>7</sup>[Charité et al. \(2015\)](#) analyze whether individuals support the use of welfare measures that respect the reference-point dependence of preferences. [Weinzierl \(2014\)](#) investigates the support for alternatives to utilitarian welfare maximization.

- Boskin, Michael J. and Eytan Sheshinski**, “Optimal tax treatment of the family: Married couples,” *Journal of Public Economics*, 1983, *20* (3), 281 – 297.
- Charité, Jimmy, Raymond Fisman, and Ilyana Kuziemko**, “Reference Points and Redistributive Preferences: Experimental Evidence,” Working Paper 21009, National Bureau of Economic Research March 2015.
- Chetty, Raj, Adam Looney, and Kory Kroft**, “Salience and Taxation: Theory and Evidence,” *American Economic Review*, September 2009, *99* (4), 1145–77.
- Coase, Ronald H.**, “The Problem of Social Cost,” *The Journal of Law and Economics*, 1960, *3*, 1–44.
- Diamond, Peter A.**, “A many-person Ramsey tax rule,” *Journal of Public Economics*, 1975, *4* (4), 335 – 342.
- , “Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates,” *American Economic Review*, 1998, *88*, 83–95.
- Farhi, Emmanuel and Iván Werning**, “Progressive Estate Taxation,” *The Quarterly Journal of Economics*, 2010, *125* (2), 635–673.
- Golman, Russell, David Hagmann, and George Loewenstein**, “Information Avoidance,” *Journal of Economic Literature*, March 2017, *55* (1), 96–135.
- Goodin, Robert E.**, *Laundering preferences* Cambridge Studies in Philosophy and Public Policy, Cambridge University Press,
- Gul, Faruk and Wolfgang Pesendorfer**, *The Case for Mindless Economics*, Oxford University Press, 2008.
- Hertwig, Ralph and Christoph Engel**, “Homo Ignorans: Deliberately Choosing Not to Know,” *Perspectives on Psychological Science*, 2016, *11* (3), 359–372. PMID: 27217249.
- Kleven, Henrik Jacobsen, Claus Thustrup Kreiner, and Emmanuel Saez**, “The Optimal Income Taxation of Couples,” *Econometrica*, 2009, *77* (2), 537–560.
- Lipsey, R. G. and Kelvin Lancaster**, “The General Theory of Second Best 1,” *The Review of Economic Studies*, 1956, *24* (1), 11–32.

- Loewenstein, George and Emily Haisley**, *The Economist as Therapist: Methodological Ramifications of "Light" Paternalism*, Oxford University Press, 2008.
- Mirrlees, James**, "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 1971, 38, 175–208.
- Myerson, Roger B and Mark A Satterthwaite**, "Efficient mechanisms for bilateral trading," *Journal of Economic Theory*, 1983, 29 (2), 265 – 281.
- O'Donoghue, Ted and Matthew Rabin**, "Optimal sin taxes," *Journal of Public Economics*, 2006, 90 (10), 1825 – 1849.
- Piketty, Thomas and Emmanuel Saez**, "A Theory of Optimal Inheritance Taxation," *Econometrica*, 2013, 81 (5), 1851–1886.
- Ramsey, Frank P.**, "A Contribution to the Theory of Taxation," *The Economic Journal*, 1927, 37 (145), 47–61.
- Saez, Emanuel**, "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, 2001, 68, 205–229.
- Sen, Amartya**, "The Impossibility of a Paretian Liberal," *Journal of Political Economy*, 1970, 78 (1), 152–157.
- Thaler, Richard H. and Cass R. Sunstein**, *Nudge - Improving Decisions About Health, Wealth, and Happiness* 2008.
- Weinzierl, Matthew**, "The promise of positive optimal taxation: normative diversity and a role for equal sacrifice," *Journal of Public Economics*, 2014, 118, 128 – 142.
- , "A Welfarist Role for Nonwelfarist Rules: An example with envy," Working Paper 21009, Harvard Business School July 2017.