Economic Influences on Moral Values

Robert Östling[*]

[*]Stockholm University, robert.ostling@iies.su.se

Economic Influences on Moral Values[*]

Robert Östling

Abstract

This paper extends standard consumer theory to account for endogenous moral motivation. Building on cognitive dissonance theory, I show how moral values are affected by changes in prices and income. The key insight is that changes in prices and income that lead to higher consumption of an immoral good also affect the moral values held by the consumer so that the good is considered less immoral. A preliminary empirical analysis based on the World Values Survey is consistent with the model's predictions with respect to income.

KEYWORDS: consumer theory, moral values, endogenous preferences, cognitive dissonance, self-serving bias

1   Introduction

> It is easier for a camel to go through the eye of a needle than for a rich man to enter the kingdom of God. (Matthew 19:24)

People hold moral values that influence their behavior. Economists have long recognized this and the observation has been used to explain a wide range of economic phenomena. Many have realized that the reverse is also true – economic factors influence moral values (e.g., Lindbeck, 1995, 1997, Bowles, 1998 and the papers cited below). In contrast to most previous economic research, this paper focuses on the individual determinants of internalized moral values. I show how standard consumer theory can be extended with a simple and yet plausible psychological mechanism to study changes in moral motivation. More specifically, the model shows how consumers that view a certain good or activity as immoral may self-servingly change their moral values as income and prices change.

The idea that moral motivation is affected by changes in prices and income has several important implications. The prevalence of moral motivation in consumer markets is demonstrated by the demand for environmentally friendly products, ethical investments, organic foods and fair trade labeled goods. Policy makers might be interested in increasing demand for such products and therefore need to be aware of the impact of economic policies on the moral motivation of consumers. For example, the expansion of low-cost airlines might increase consumers' moral tolerance of carbon emissions, which may counteract measures taken to combat climate change. Relatedly, higher incomes may increase consumption of "immoral goods", for example air travel, which is likely to affect moral attitudes. The framework for endogenous moral values laid out here can help to explain moral attitudes regarding "immoral" consumer goods, and it can also explain attitudes toward tax evasion and benefit fraud.

In the model, prices and income affect the incentives for immoral behavior, entailing a conflict between narrow self-interest and moral values. This conflict gives rise to cognitive dissonance, which the consumer can reduce by exerting effort in order to modify her moral values. The main prediction of the model is that higher consumption of a good implies that the consumer will view that good more favorably from a moral point of view. For normal goods, higher incomes therefore lead to higher moral acceptability, whereas the opposite is true for inferior goods. An empirical analysis using data from the World Values Survey supports the prediction regarding the effects of income

on moral motivation. Unfortunately, however, the empirical analysis is not entirely conclusive due to data limitations.

The model aims to capture an idea that Elliot Aronson, a leading social psychologist, presents as hypothetical advice from a modern Machiavelli: "If you want people to soften their moral attitudes toward some misdeed, tempt them so that they perform that deed" (Aronson, 2003, p. 162). This type of self-justification of moral attitudes is discussed by Aronson (2003, chapter 5), but the underlying psychological idea is based on the cognitive dissonance theory developed by Festinger (1957).

Cognitive dissonance was first introduced in economics by Hirschman (1965) and Akerlof and Dickens (1982) and has since been used in several economic applications.[1] This paper's combination of cognitive dissonance theory and standard consumer choice theory is most closely related to Rabin (1994) and Konow (2000). My model extends their framework in two respects. First, rather than having only one consumption good, I allow the consumer to choose from a finite number of goods. Second, the consumer faces a budget constraint and there are prices attached to all goods. This links their approach to standard consumer theory and allows comparative statics in terms of easily observable variables such as prices and income. In particular, having more than one consumption good in the model is required to distinguish between normal and inferior goods, which is critical for the empirical identification of the model's main predictions.

Apart from Rabin (1994) and Konow (2000), there have also been some other attempts to model moral motivation endogenously. For example, Brekke et al. (2003) study moral motivation in a public goods provision model where a commonly shared moral norm can be affected by policy. Specifically, they assume that people are utilitarian and apply Kantian reasoning, i.e., the moral norm is determined by the action that maximizes the sum of all players' utility given that everybody takes the same action. This paper instead takes moral values as given and focuses on a psychological mechanism that may change moral motivation irrespective of what moral philosophical principles that underlie consumers' moral values. Frey (1997) makes a distinction between intrinsic and extrinsic motivation and uses it to explain the crowding out of blood supply suggested by Titmuss (1970). Frey (1997) discusses how

---

[1]For example, economics of crime (Dickens, 1986), moral behavior and social change (Rabin, 1994), biased fairness norms in the dictator game (Konow, 2000), mobility-reducing norms of low-productivity farmers (Haagsma and Koning, 2002), environmentally friendly goods (van de Ven, 2003), formation of underclass attitudes (Oxoby, 2003, 2004), redistributive politics (Bénabou and Tirole, 2006a) and changes in political attitudes after elections (Beasley and Joslyn, 2001, and Mullainathan and Washington, 2009).

extrinsic motivation may affect intrinsic motivation, but he provides no formal model of this interaction.[2] In my model, there is a tension between extrinsic (prices) and intrinsic (moral) motivation, but the effects of extrinsic motivation cannot be reversed through crowding out of moral motivation.

2   Model

Consider the familiar utility maximization problem of a consumer with a fixed endowment w that she spends on N consumption goods. The consumed non-negative quantities are denoted by the vector $\boldsymbol{x} = (x_1, x_2, ..., x_N)$. Consumption of these goods gives the consumer material utility $u(\boldsymbol{x})$, which is a standard utility function that is twice continuously differentiable, strictly concave and increasing in all goods.

The consumer not only cares about material utility, but she also has original moral values $\boldsymbol{\mu} = (\mu_1, \mu_2, ..., \mu_N)$ that measures how immoral the consumer considers consumption of the different goods to be. Moral values are non-negative and good i is called *immoral* if $\mu_i > 0$ and *amoral* if $\mu_i = 0$. The original moral values are exogenous, but the consumer can choose to change moral values when making her consumption choice. The chosen moral values are denoted $\boldsymbol{m} = (m_1, m_2, ..., m_N)$.

Consuming a good may result in cognitive dissonance, which is measured by non-negative functions $d_i(m_i, x_i)$. The dissonance functions are twice continuously differentiable, strictly convex and increasing in m_i and x_i. There is no dissonance if a good is not consumed or if the consumer considers a good to be amoral, i.e., we assume $d_i(0, x_i) = 0$ and $d_i(m_i, 0) = 0$. Furthermore, the marginal dissonance with respect to consumption of an immoral good is increasing in the moral value, i.e., the cross-derivatives of the dissonance functions are positive, $\partial^2 d_i(m_i, x_i)/\partial m_i \partial x_i = \partial^2 d_i(m_i, x_i)/\partial x_i \partial m_i > 0$.

Deviating from original moral values, however, comes at a utility cost. This cost of self-deception is increasing and strictly convex in $|\mu_i - m_i|$. For simplicity, we assume that this cost is $\delta_i(\mu_i - m_i)^2$ with $\delta_i > 0$. Overall utility is additively separable in its $2N + 1$ components:

---

[2]By now there are some economic models that potentially can explain this and other related phenomena. Bénabou and Tirole (2003) show how motivation might change if a person has imperfect self-knowledge and therefore may be sensitive to signals from a more informed player. Bénabou and Tirole (2006b) and Ellingsen and Johannesson (2008) instead assume that players are concerned about signaling their character traits. This paper differs from these theoretical models since I do not assume any form of interaction with other players – the focus is on a psychological mechanism on the individual level.

$$U\left(\boldsymbol{x}, \boldsymbol{m}\right) \equiv u(\boldsymbol{x}) - \sum_{i=1}^{N} \left[d_i(m_i, x_i) + \delta_i(\mu_i - m_i)^2 \right].$$

The consumer maximizes utility by simultaneously choosing consumption and moral values. It may appear more intuitive that consumers first choose consumption quantities and then rationalize their consumption decisions by adapting moral values. However, Lieberman et al. (2001) provide suggestive evidence that this intuition is likely to be false and that attitude change is a highly automated process that is hard to temporally separate from the behavioral decision. The utility maximization problem is therefore

$$\max_{\boldsymbol{x}, \boldsymbol{m}} \quad u(\boldsymbol{x}) - \sum_{i=1}^{N} \left[d_i(m_i, x_i) + \delta_i(\mu_i - m_i)^2 \right],$$

subject to

$$\boldsymbol{p} \cdot \boldsymbol{x} \leq w,$$
$$\boldsymbol{x} \geq 0,$$
$$\boldsymbol{m} \geq 0,$$

where $\boldsymbol{p} = (p_1, p_2, ..., p_N)$ are the positive goods prices.

For the sake of simplicity, we assume throughout that there is at least one amoral good and that $\lim_{x_i \to 0} U(x_i, \cdot) = -\infty$ for all goods. These assumptions imply that the budget constraint is binding and that positive quantities of all goods are consumed. To guarantee that the solution to the maximization problem is unique and given by the Kuhn-Tucker conditions, overall utility is assumed to be strictly quasiconcave (see Theorem M.K.4 in Mas-Colell et al., 1995). Proposition 1 and 2 hold also if there are multiple solutions, but uniqueness is required to derive the result in Proposition 3. The solution to the utility maximization problem is denoted $\boldsymbol{x}^* = (x_1^*, x_2^*, ..., x_N^*)$ and $\boldsymbol{m}^* = (m_1^*, m_2^*, ..., m_N^*)$. Note that it is always optimal to choose $m_i^* = 0$ when $\mu_i = 0$, so there is no dissonance or self-deception associated with the consumption of amoral goods.

The model predicts that chosen moral values, $\boldsymbol{m}^*$, differ from original values, $\boldsymbol{\mu}$, for all immoral goods. It follows directly from the first-order conditions with respect to moral values (found in the Appendix) that chosen moral values are given by

$$m_i^* = \max \left\{ 0, \mu_i - \frac{\partial d_i(m_i^*, x_i^*)}{\partial m_i} \frac{1}{2\delta_i} \right\}. \tag{1}$$

Since the dissonance function is increasing in m_i, and δ_i is positive, condition (1) implies that $m_i^* \leq \mu_i$. The fact that the consumer deceives herself when consuming immoral goods is an example of self-servingly biased moral values (see Babcock and Loewenstein, 1997, for an introduction to self-serving biases in economic problems).

Using the optimality condition (1), it is straightforward to establish comparative statics with respect to income and prices. The results below are stated globally, but they all apply locally around the optimum. For example, goods do not have to be normal or inferior at all prices and income levels – it is sufficient that they are normal or inferior around the price and income levels under consideration.

The first proposition states how chosen moral values are affected by income changes.

Proposition 1 *Chosen moral values, $\boldsymbol{m}^*$, are weakly decreasing in income for normal goods and weakly increasing in income for inferior goods.*

Proof. If m_i^* is positive, the result follows from differentiating the second (positive) part of (1) with respect to income:

$$\frac{\partial m_i^*}{\partial w} = -\left[\frac{\partial^2 d_i(m_i, x_i)/\partial m_i \partial x_i}{2\delta_i + \partial^2 d_i(m_i, x_i)/\partial^2 m_i}\right]\frac{\partial x_i^*}{\partial w}.$$

The result immediately follows from this expression once it is noted that all terms within brackets are positive by assumption. In the corner solution, when m_i^* is zero, moral values are unaffected by infinitesimal income changes (unless the second part of (1) exactly equals zero). ■

The intuition for the result in Proposition 1 is straightforward. When an immoral good is normal, higher income leads to higher consumption. Higher consumption of an immoral good creates cognitive dissonance, which can be reduced (at the margin) by changing moral values so that the immoral good is believed to be less immoral than before. In other words, if we consume more of goods that we believe are immoral to consume, then we adjust our values in order to reduce the dissonance that the increased consumption gives rise to.

There are of course other ways in which higher incomes affect moral values. For example, as argued by Shleifer (2004), higher income also provides greater opportunity to behave morally when it is costly to do so. Moreover, when a higher income is observed by others, there could be a change in social pressures to behave morally. Nevertheless, the model points at another rather general mechanism. Higher income leads to higher consumption, which has

consequences for our moral attitudes. Most people cannot stand considering themselves to be immoral persons, and so they need to adjust their moral values to be compatible with their consumption patterns. A historical example of this effect is when the Catholic Church lifted their ban on eating meat on Fridays in the mid-1960s, supposedly because incomes had grown and meat had become relatively cheaper.

The second proposition states how the moral value for one good is affected by changes in the price of another good.

Proposition 2 *(i) If good i is a gross substitute for good j, i.e., $\partial x_i^* / \partial p_j > 0$, then the chosen moral value for good i, m_i^*, is weakly decreasing in the price of the other good, p_j. (ii) If good i is a gross complement for good j, i.e., $\partial x_i^* / \partial p_j < 0$, then the chosen moral value for good i, m_i^*, is weakly increasing in the price of the other good, p_j.*

Proof. Whenever $m_i^* > 0$, the result follows from differentiation of the positive part of (1) with respect to p_j:

$$\frac{\partial m_i^*}{\partial p_j} = - \left[\frac{\partial^2 d_i\left(m_i, x_i\right) / \partial m_i \partial x_i}{2\delta_i + \partial^2 d_i\left(m_i, x_i\right) / \partial^2 m_i} \right] \frac{\partial x_i^*}{\partial p_j}.$$

If $m_i^* = 0$, then that moral value is generally unaffected by infinitesimal changes in prices. ∎

Again, the intuition for the result is clear. If a change in the price of some good leads to higher consumption of an immoral good, then that immoral good is considered less immoral.

Proposition 2 also demonstrates a kind of crowding-out effect. Suppose that an immoral and an amoral good are gross complements. A price decrease of the amoral good will lead both to higher consumption and higher moral acceptance of the immoral good. Hence, extrinsic motivation—a lower price of the morally superior amoral good—can crowd out intrinsic moral motivation. This is a kind of motivational crowding-out effect, although the effect goes via a change in consumption of the immoral good and not directly from extrinsic motivation (lower price) to intrinsic motivation (moral values).

The price effects illustrated by the model are relevant for policy to the extent that policy makers are interested in stimulating particular moral values. Consider the case of organic food subsidies. For many people the choice of organic food is to some extent motivated by moral concerns. The model illustrates that such subsidies will not only affect the consumption of organic

foods, they may also affect the demand for other goods and the moral values attached to these goods. It could be argued that lower prices on organic food will result in higher consumption of other more environmentally harmful goods, which would affect moral attitudes in favor of these other goods.

The above discussion hides one complication. The terms normal and inferior goods have been used as if consumption is determined independently of moral values. However, moral values and the consumption of immoral goods are chosen simultaneously and both affect the level of cognitive dissonance. Despite this, it is straightforward to derive the regular Slutsky equation and show that the standard interpretation of income and substitution effects carries over to this setting.

Combining the Slutsky equation (derived in the Appendix) with condition (1) allows us to characterize price effects more clearly. Proposition 3 is stated in terms of substitutes and complements, which are defined in the usual way. Let $h_i(\boldsymbol{p}, \boldsymbol{\mu}, U)$ denote compensated demand for good i, i.e., $h_i(\boldsymbol{p}, \boldsymbol{\mu}, U)$ is the expenditure-minimizing consumption level of good i at prices $\boldsymbol{p}$, original moral values $\boldsymbol{\mu}$ and utility level U (see the Appendix for details). Two goods i and j are substitutes if $\partial h_i(\boldsymbol{p}, \boldsymbol{\mu}, U)/\partial p_j$ is positive and complements if $\partial h_i(\boldsymbol{p}, \boldsymbol{\mu}, U)/\partial p_j$ is negative.

Proposition 3 *(i) If good i is a normal good, then the chosen moral value, m_i^*, is weakly increasing in the price of that good, p_i. (ii) If good i is an inferior good and a substitute for another good j, then the chosen moral value m_i^* is weakly decreasing in the price of the other good, p_j. (iii) If good i is a normal good and a complement for another good j, then the chosen moral value m_i^* is weakly increasing in the price of the other good, p_j.*

Proof. Whenever $m_i^* > 0$, differentiation of second part of (1) with respect to p_j yields

$$\frac{\partial m_i^*}{\partial p_j} = -\left[\frac{\partial^2 d_i(m_i, x_i)/\partial m_i \partial x_i}{2\delta_i + \partial^2 d_i(m_i, x_i)/\partial^2 m_i}\right] \frac{\partial x_i^*}{\partial p_j}.$$

Substituting $\partial x_i^*/\partial p_j$ from the Slutsky equation (derived in the Appendix) implies

$$\frac{\partial m_i^*}{\partial p_j} = \left[\frac{\partial^2 d_i(m_i, x_i)/\partial m_i \partial x_i}{2\delta_i + \partial^2 d_i(m_i, x_i)/\partial^2 m_i}\right] \left[\frac{\partial x_i^*}{\partial w} x_j^* - \frac{\partial h_i(\boldsymbol{p}, \boldsymbol{\mu}, U)}{\partial p_j}\right].$$

Note that the terms within the first brackets are positive. The first part of the result follows from the law of compensated demand, i.e., $\partial h_i(\boldsymbol{p}, \boldsymbol{\mu}, U)/\partial p_i \leq 0$ (see the Appendix for details). The remaining two parts of the result

comes from the usual definition of substitutes and complements, i.e., that good i and j are substitutes if $\partial h_i\left(\boldsymbol{p},\boldsymbol{\mu},U\right)/\partial p_j > 0$ and complements if $\partial h_i\left(\boldsymbol{p},\boldsymbol{\mu},U\right)/\partial p_j < 0$. Finally, note that when $m_i^* = 0$, marginal changes in prices will generally not affect chosen moral values. ∎

If an immoral good is normal, then the income effect is negative. This implies that an increase in the price of that good leads to lower consumption and an upward adjustment of the moral value, i.e., it is considered more immoral. This result suggests that the failure of policy makers to correct for externalities might be associated with an additional "moral cost". For example, due to international agreements there is currently no tax on air fuel in most countries, which means that the negative environmental externality of air travel is not reflected in the market price. Since air travel is cheaper without the tax, more people are travelling and their attitudes toward pollution are less negative than they would have been if the tax was in place.

The second and third parts of Proposition 3 clarify the effects on moral values for one good of changes in the price of some other good. For example, if a good is normal and a complement for some other good, increasing the price of that other good decreases consumption of the first good and implies that the first good is considered more immoral.

3 Extensions

Before proceeding to the empirical analysis, I consider a number of possible extensions of the basic model.

3.1 Dynamics

The model presented above is static and silent about what happens to the consumer's moral values after the consumption decision. In order to develop a dynamic version of the model there are three issues that need to be resolved.

First, does the self-deception investment in reduction of moral values have a transitory or permanent effect? Recall from the expression for the optimal moral value (1) that chosen moral values are lower than original moral values (for immoral goods). At one extreme, $\boldsymbol{\mu}$ may be constant over time and consumers merely deceive themselves at the time of consumption. This is likely to be a reasonable assumption if moral values reflect clear and well-defined moral principles that are less susceptible to long run self-deception. For example, if moral values are determined by external effects caused by

consumption (as in Brekke et al., 2003) and there is little uncertainty about these external effects, then it might be more difficult to permanently and self-servingly adapt the moral rule. At another extreme, self-deception could be a one-time investment – the chosen moral value in one period is the original moral value in the next consumption decision.

Second, the predictions of a dynamic model would also depend on to what extent people are forward-looking and manage to predict their own changing preferences. In many situations people seem to systematically mispredict changes in future preferences (Loewenstein and Schkade, 1999) and it is therefore unclear if it is reasonable to assume that people correctly anticipate changes in moral values.

Finally, note that in the static model, the consumer does not derive any positive utility from holding particular moral values – the original moral values only result in utility costs in the form of dissonance and self-deception costs. It may be more reasonable, however, to assume that the consumer derives positive utility from a high level of μ, perhaps because of concerns about maintaining a positive self-image. This would not matter in the static model since μ is unaffected by the consumer's choice, but it could play a role in a dynamic model where original moral values depend on past consumption.

If self-deception efforts have long-lasting effects and people either 1) anticipate future changes in moral values but do not care intrinsically about moral values or 2) do not anticipate future preference changes, then the model implies that moral values tend to erode over time. This does not necessarily imply that we should expect all consumers to not hold any moral values. Even in a dynamic model that has the implication that moral values are zero in the long run, certain consumption decisions are taken irregularly and consumers might therefore not have experienced sufficiently many such consumption decisions to completely erode moral values. The next section considers a simple dynamic extension of the model.

3.2 Moral slippery slope

Moral rules are often dichotomous, for example, it is probably more common with a moral value stating that it is always wrong to shoplift rather than a rule that allows for some shoplifting. Although there are many potential reasons for this, a simple dynamic version of the model can show that individuals might be better off with a discrete rather than a continuous choice set. A continuous choice set might tempt the consumer to consume some of the immoral good and to adapt moral values accordingly. If self-deception has a permanent effect and future changes of preferences are not anticipated, then there may be a "moral

slippery slope": consuming a little of the immoral good may lead to more consumption in the future and to an erosion of the consumer's moral values.

To illustrate this idea, consider the following simple dynamic version of the model. Suppose that there is one immoral (x_t) and one amoral good, and that the immoral good is available in quantities between zero and one. Let the price of both goods be 1 and let $w > 1$ so that the consumed quantity of the amoral good is $w - x_t$. (Since both the original and chosen moral values attached to the amoral good are always zero, these are left out of the discussion.) The original moral value at one point in time is given by the chosen moral value in the previous time period, i.e., $\mu_t = m^*_{t-1}$. The consumer myopically maximizes utility in each time period and does not take into account the effect of today's consumption decision on future consumption decisions. Furthermore, assume that the consumer intrinsically derive utility (linearly) from holding a particular moral value. Specifically, let utility in period t be given by

$$U\left(x_t, w - x_t, m_t, \mu_t\right) = u\left(x_t, w - x_t\right) - d\left(m_t, x_t\right) - \delta\left(\mu_t - m_t\right)^2 + \mu_t.$$

For the sake of the argument, let us assume that the consumer prefers consuming some of the immoral good to consuming nothing, but also prefers consuming nothing compared to one unit of the immoral good. Specifically, assume that for all $\mu_t \geq 0$ there is some $x^*_t \in (0,1)$ and $m^*_t \geq 0$ such that

$$U\left(x^*_t, w - x^*_t, m^*_t, \mu_t\right) > U\left(0, w, 0, \mu_t\right).$$

Assume also that for all $\mu_t \geq 0$ the following holds:

$$U\left(0, w, 0, \mu_t\right) > U\left(1, w - 1, m^{**}_t, \mu_t\right),$$

where m^{**}_t denotes the optimally chosen moral value when one unit of the immoral good is consumed and the original moral value is μ_t. Note that this implies that the previous assumption that utility is negatively unbounded for zero consumption is disregarded here.

These assumptions trivially imply that the consumer never consumes the immoral good when the choice set is discrete, i.e., when $x_t \in \{0, 1\}$. However, when the consumer can consume any quantity between zero and one, these assumptions imply that $x^*_1 > 0$, and consequently by (1) that $m^*_1 < \mu^*_1$ for all $\mu_1 > 0$. Since chosen moral values carry over to the next period and the consumer is myopic, chosen moral values will decrease in all time periods until chosen moral values are zero.

Let x_0 denote the optimal consumption when moral values are completely eroded (i.e., $\mu = 0$). After sufficiently many time periods (perhaps only in the

limit) consumer utility will simply be $u\left(x_0, w - x_0\right)$, whereas utility with the discrete choice set is $u\left(0, w\right) + \mu_1$ in every period. Although $u\left(x_0, w - x_0\right) > u\left(0, w\right)$ by the assumption above, it is clear that $u\left(0, w\right) + \mu_1 > u\left(x_0, w - x_0\right)$ for sufficiently high μ_1. Consequently, if the consumer has strong initial moral values, she will end up being better off with the constrained choice set than with the continuous choice set (although she may prefer the latter in the short run).

3.3   Moral multipliers

Up until now it has been assumed that there is one moral value associated with each immoral good. A natural extension of the model is to consider the case when the same moral value is associated with several different goods. For example, a consumer might have moral values not to cause too much carbon emissions, which may give rise to dissonance both when travelling by air and by car. In this case there might be a moral multiplier effect of taxing one of the goods – a higher price of one immoral good might not only decrease consumption of that good, but may also lead to stronger moral values which reduces consumption of some other good.

To illustrate this idea in the simplest possible way, suppose there are two immoral goods, x_1 and x_2, and one amoral good, x_A. Furthermore, suppose that the second immoral good is free $\left(p_2 = 0\right)$ and that material utility is additively separable. The utility function is given by

$$u_1\left(x_1\right) + u_2\left(x_2\right) + u_A\left(x_A\right) - d_1(m_1, x_1) - \delta_1(\mu - m_1)^2 - d_2\left(m_2, x_2\right) - \delta_2(\mu - m_2)^2.$$

Since the second immoral good is free, consumption is simply given by equalizing marginal material utility and marginal dissonance. Due to additive separability of material utility, consumption of the free immoral good is independent of the price of the first immoral good.

Suppose instead that the same moral value, m, causes dissonance when both immoral goods are consumed. One natural way to incorporate this in the consumer's utility function is the following:

$$u_1\left(x_1\right) + u_2\left(x_2\right) + u_A\left(x_A\right) - d_1(m, x_1) - d_2\left(m, x_2\right) - \delta(\mu - m)^2.$$

Since the second immoral good is free, consumption is again given by equalizing marginal material utility with marginal dissonance. However, now marginal dissonance for the free immoral good depends on the commonly chosen moral value. Since the moral value depends on the consumption of both immoral goods, consumption of the free moral good will generally depend on the price

of the first immoral good. To see this more clearly, note that the optimally chosen moral value is now given by

$$m^* = \max\left\{0, \mu - \left(\frac{\partial d_1(m^*, x_1^*)}{\partial m} + \frac{\partial d_2(m^*, x_2^*)}{\partial m}\right)\frac{1}{2\delta_i}\right\}.$$

A price decrease of the first immoral good which leads to higher consumption of that good may lead to lower moral values and an increase in consumption of the free immoral good. With this specification, a price decrease of p_1 which is accompanied by an increase of x_1 may, in principle, also lead to higher m and lower consumption of x_2. The main point here, however, is that the independence between p_1 and x_2 is broken due to the common moral value.

3.4 Other determinants

Although this paper focuses on economic determinants of moral values, it is worth noting that the model can also shed light on other determinants of moral values. Most directly, people that derive more material utility from a certain immoral good will view consumption of that good more favorably from a moral point of view. For example, the French and Hungarians seem to derive more pleasure out of eating foie gras than people in most other countries. Based on the model in this paper, we should consequently expect the French and Hungarians to consider force-feeding of birds for food production morally more acceptable than people in other countries. This idea has been discussed and formalized by Rabin (1994), so I do not elaborate further on it here.

Relatedly, to the extent that personal characteristics affect the availability or cost of certain immoral activities, the model also has clear predictions. For example, if men take the car more often than women, we could expect men to be more tolerant toward overspeeding. People that live in areas with good public transport are likely to find commuting by car less acceptable than people that live in areas where there is no alternative to taking the car.

4 Empirical Analysis

In order to test the predictions of Proposition 1, I use data from the 1999-2004 wave of the World Values Survey (WVS).[3] The 1999-2004 wave of the

---

[3]The data has been obtained from www.worldvaluessurvey.org and the latest wave of the survey has been extracted from the following integrated data file: European Values Study Group and World Values Survey Association. EUROPEAN AND WORLD VALUES SURVEYS FOUR-WAVE INTEGRATED DATA FILE, 1981-2004, v. 20060423, 2006.

WVS contains responses to survey questions from 101,000 individuals in 70 countries. The list of countries included, and a more detailed description of the data, can be found in the Appendix. Respondents are, among other things, asked about their moral attitudes toward certain behaviors. These questions are phrased as follows: "Please tell me for each of the following statements whether you think it can always be justified, never be justified, or something in between" and the respondents were asked to answer on a scale from 1 to 10 where 1 means "always justifiable" and 10 means "never justifiable" for different types of activities. The scale has here been reversed to be consistent with the interpretation of m in the theoretical model.

The effect of income on moral values can be identified by estimating the following regression:

$$m_i = \alpha + \beta y_i + \boldsymbol{X}_i \boldsymbol{\gamma} + \varepsilon_i,$$

where m_i is the stated moral value, y_i the income of the respondent and $\boldsymbol{X}_i$ a vector with country dummies and individual characteristics. The individual characteristics are sex and age in the "short" specification, whereas the "long" specification, in addition, controls for educational level, employment status, profession, marital status, number of children and size of home town. All these characteristics are included as dummies using the response alternatives available in the WVS (see Appendix for details). The income data in the WVS refers to total household pre-tax income (including "pensions and other incomes") and is measured in ten country-specific income brackets based on self-reports. Income is consequently measured with error, but there is little reason to expect that measurement error is correlated with true income. The estimated income coefficients are therefore likely to be biased toward zero. Since several of the control variables are strongly correlated with income, the inclusion of these variables most likely exacerbates the attenuation bias. We should therefore expect smaller income coefficients in the long rather than in the short specification.

Proposition 1 predicts that the income coefficient β is negative for normal goods and positive for inferior goods. Moral values, however, may be correlated with other individual characteristics related to income. Although many such

characteristics are included as controls in the long regressions, the estimated income coefficients might be biased due to omitted variables. However, these omitted variables are likely to be correlated with income and moral values in the same way for both normal and inferior goods. This implies that the estimated income coefficients are not necessarily expected to have opposite signs, but that they should be higher for inferior than for normal goods.

Some questions in the WVS refer to goods or activities that are difficult to relate to income and have therefore been left out.[4] The remaining goods and activities referred to in the questions are classified into inferior and normal goods based on a priori concerns and available empirical evidence. Although some of these goods and activities are not typical consumption goods, the theoretical model can be seen as a reduced form of richer models that model each situation in more detail. For example, for several of the activities there is a risk of legal sanctions, but prices can be interpreted as a reduced form representation of the expected material cost of punishment.

One potential problem with the empirical analysis is that for two of the questions used, income might depend on moral values. People who are more tolerant toward benefit fraud and tax evasion will probably cheat more and might therefore report a higher income in the survey. Although this can rationalize a negative relationship between income and tax morale, it cannot explain a positive relationship between benefit fraud and income. In addition, the income from benefit fraud and tax evasion is likely to constitute a negligible fraction of reported incomes for most respondents.

Table 1 reports the income coefficients from the two different specifications with moral values as the dependent variables. The top section of Table 1 refers to activities that are likely to be inferior goods, and the bottom section refers to normal goods. A subset of the WVS is the European Values Survey (EVS), which contains some extra moral value questions for 32 European countries. Table 1 reports income coefficients estimated using the whole sample and the EVS countries separately.

Activities like benefit fraud, stealing cars and avoiding public transport

---

[4]These questions concern homosexuality, abortion, divorce, casual sex, euthanasia, suicide, lying, adultery, sex under legal age of consent, littering, political assassination, experiments on human embryos and genetic manipulation of food. Furthermore, two other questions have also been left out although they might be related to income: accepting bribes and buying stolen goods. Although poorer people have stronger economic incentives to engage in these activities (given diminishing marginal utility of money), they also have less money to spend on stolen goods and they might be less likely to be offered bribes. Finally, a question regarding alcohol consumption that was only asked in Muslim countries has been left out.

fares are likely to be inferior since the incentive to engage in these activities is higher the lower the income (given diminishing marginal utility of money). In line with the prediction of Proposition 1, all significant income coefficients for these questions are positive. Smoking is also an inferior good, at least in industrialized countries (Chaloupka and Warner, 2000), and we might therefore expect richer people to be less tolerant of smoking in public buildings. As can be seen from Table 1, this is not supported by the data, but the negative coefficients are not statistically significant.

The bottom six questions in Table 1 refer to goods and activities that are likely to be positively related to income. Although I have found no data on sex buyers, it seems most plausible that sex a normal good.[5] Alcohol and marijuana consumption is positively related to income, at least in the US (Saffer and Chaloupka, 1999).[6] Since driving a car is also a normal good, the propensity to drive under the influence of alcohol is probably increasing in income. Similarly, Shinar et al. (2001) show that there is a positive relationship between overspeeding and income in the US. The evidence on tax evasion is somewhat mixed, but several studies point at a positive relationship (see Andreoni et al., 1998, for a discussion). As can be seen from Table 1, all income coefficients for normal goods have the predicted negative sign, but they are not statistically significant for all questions in all specifications.

Adjusted R^2 for the regressions reported in Table 1 varies from 0.07 (the short specification for the drink and drive question) to 0.21 (the long specification for the prostitution question with all countries included). The income coefficients generally have the predicted signs and are statistically significant in most of the regressions. For example, based on the short specification for all countries, an increase in income from the lowest to the highest income category implies that individuals on average believe that prostitution is 0.6 more morally justifiable on a 1 to 10 scale. This corresponds to an increase of the moral value of one fourth of a standard deviation. Although this is a relatively small effect, there are several reasons why we should expect it to be small. First, income is poorly measured, which implies that the coefficients

---

[5] In the case of prostitution, it is typically poor women who work as prostitutes. Some people working as prostitutes are likely to be included in the sample, and the results might therefore be affected if prostitutes are more tolerant toward prostitution. However, buyers outnumber sellers by far, so this is likely to be a limited problem. Moreover, excluding women from the sample leads to somewhat stronger effects in three regressions and a somewhat weaker effect in one regression, but the coefficients remain negative and strongly significant in all four cases.

[6] As with prostitution it is possible to make a supply-side argument. If drug dealers on average are poorer we would expect these to be more morally tolerant of soft drugs. This effect is probably marginal since buyers are likely to outnumber sellers.

Table 1: Moral values regressions

	European countries		All countries	
	Short	Long	Short	Long
Benefit fraud	**4.65*****	**1.78*****	**4.25*****	**1.52*****
	10.34	3.04	12.59	3.46
	32/33161	32/24334	66/77673	56/54102
Avoiding fare	0.13	−0.62	**2.33*****	**1.27*****
	0.20	0.72	6.22	2.56
	17/19802	17/14970	53/66522	42/45388
Joyriding	**0.98*****	0.35		
	3.65	1.00		
	32/33729	32/24764		
Smoking in public	−0.55	−0.41		
	0.93	0.53		
	32/33201	32/24365		
Prostitution	**−7.26*****	**−3.62*****	**−6.05*****	**−3.57*****
	10.25	3.90	16.96	7.40
	19/19789	19/14420	53/65096	41/42675
Taking soft drugs	**−1.05*****	−0.23		
	2.71	0.48		
	32/33522	32/24601		
Drink and drive	**−0.86*****	**−1.06*****		
	3.11	2.92		
	32/33759	32/24783		
Overspeeding	**−5.06*****	**−3.23*****		
	12.03	5.80		
	32/33580	32/24655		
Pay cash for services	**−3.86*****	**−3.24*****		
	6.70	4.28		
	32/32357	32/23691		
Cheating on taxes	−0.52	−0.69	−0.04	**−1.03****
	1.00	1.20	0.14	2.33
	32/33263	32/24393	66/78561	55/53446

The table reports income coefficients multiplied by 100, absolute t values and the number of countries/observations used in estimating the coefficient. Controls in the short specification are country dummies, age and sex. The long specification in addition controls for marital status, educational level, employment status, occupation, size of home town and number of children. * = 10%, ** = 5% and *** = 1% significance level.

are likely to be biased toward zero. The coefficients are generally smaller in the long specification, suggesting that inclusion of extra controls exacerbates attenuation bias. Second, consumption of most of the goods listed in Table 1 is not particularly strongly related to income. Moreover, many people are likely to never have consumed some of the goods listed there, implying that only a subset of the population is used to identify the effect. Since the relation between income and consumption is weak, the relationship between income and moral values should also be weak.

As a robustness check, I separately run regressions for each country and question. Although the results differ slightly between countries, the pattern from Table 1 persists. With smoking in public buildings as the main exception, income coefficients are more often significant and positive than significant and negative for inferior goods, and the other way round for normal goods.[7]

The above analysis shows that the relationship between income and moral values observed in the data is consistent with Proposition 1. Could these findings be explained in some other way than by the theoretical model in this paper? The obvious candidate is that some variable correlated with both income and moral values has been omitted in the regressions. However, in order for such an omitted variable to rationalize the empirical findings above, the variable must be correlated differently with moral values depending on whether the good is normal or not. It is hard to see what kind of omitted variable this could be. A potential alternative interpretation of the empirical findings could be peer group effects – if the poor mainly socialize with the poor, they might adjust their moral values to each other. This alone cannot explain the empirical pattern, since it does not provide an account for why the poor should be more tolerant toward, for example, benefit fraud in the first place.

The main caveat with the empirical analysis is the classification of normal and inferior goods. Unfortunately, for several questions there is too little evidence available to precisely document the classification. In addition, based on cross-section data it is impossible to completely rule out the possibility that there are omitted variables that affect the results. The empirical results can therefore not be considered as conclusive evidence, but at least the results show that the data is not inconsistent with the model's predictions.

The ideal test of the theoretical model should involve measuring people's moral values before and after an exogenous change in income or prices. Naturally, it is hard to find such data in the field, and the cleanest test of the model

---

[7]The other exceptions are in the long specification for the question regarding soft drugs and the short specification regarding tax evasion, which show the opposite pattern.

probably requires experimental methods. Cognitive dissonance has long been studied experimentally by psychologists (several of these experiments are discussed by Aronson, 2003). One such study that is relatively close to the setting discussed here is the experiment on school children by Mills (1958). In the experiment, pupils participated in a classroom contest and were told that the best student would win a prize. The contest was such that experimenters could detect who had cheated. The children were asked both before and after the contest about their attitudes toward cheating. On average, those who had cheated also changed attitude toward thinking cheating was more acceptable, which is consistent with the model.

5   Concluding Remarks

We have seen that if a change in prices or income leads to higher consumption of an immoral good, the consumer becomes more morally tolerant toward that good. This suggests that there may be a "personal price" to pay for higher incomes – higher income not only leads to more consumption, but also to changes in moral values.

A priority for future research is an experimental test of the theory. On the theoretical side it would be interesting to further develop the extensions considered in Section 3, but also to extend the model to include the social environment. Social influences on moral values might for example be modeled as direct peer-group influence on moral values, or by incorporating social pressure in terms of social rewards and punishments. The model can also be incorporated in a general equilibrium framework to study indirect social effects on values and norms through prices.

Appendix: Derivation of the Slutsky Equation

Since there is at least one amoral good and material utility is increasing in consumption of all goods, it follows that the budget constraint will be binding (there is no "moral cost" of consuming more of the amoral good and the consumer will therefore be better off consuming more of that good). Let λ_B be the Lagrange multiplier for the budget constraint and let $\lambda_i \geq 0$ be the Kuhn-Tucker multipliers for the N non-negativity constraints. A solution $(\boldsymbol{x}^*, \boldsymbol{m}^*)$

to this maximization problem is given by the first order conditions

$$\frac{\partial u(\boldsymbol{x}^*)}{\partial x_i^*} - \frac{\partial d_i(m_i^*, x_i^*)}{\partial x_i^*} + \lambda_B p_i = 0,$$

$$-\frac{\partial d_i(m_i^*, x_i^*)}{\partial m_i^*} + 2\delta_i(\mu_i - m_i^*) \leq 0,$$

$$w - \boldsymbol{p} \cdot \boldsymbol{x}^* = 0,$$

where the second inequality is binding whenever $m_i^* > 0$. Rearranging the second first order condition gives the expression for the optimal moral value (1).

From here onward, the derivation of the Slutsky equation follows Section 3.E and 3.G in Mas-Colell et al. (1995) closely and I therefore only point out the required modifications of their proofs here. The consumer's expenditure minimization problem (EMP) is given by

$$\min_{\boldsymbol{x}, \boldsymbol{m}} \boldsymbol{p} \cdot \boldsymbol{x},$$

subject to

$$u(\boldsymbol{x}) - \sum_{i=1}^{N} \left[d_i(m_i, x_i) + \delta_i(\mu_i - m_i)^2 \right] \geq U,$$

$$\boldsymbol{x} \geq 0,$$

$$\boldsymbol{m} \geq 0.$$

As in the utility maximization problem (UMP), we focus on positive prices and we assume that the EMP is feasible, i.e., there is some $\boldsymbol{x} \geq 0$ and $\boldsymbol{m} \geq 0$ such that $U(\boldsymbol{x}, \boldsymbol{m}) \geq U$. Since the EMP is feasible and it can never be expenditure minimizing to choose moral values above $\boldsymbol{\mu}$, the constraint set is effectively compact guaranteeing that a solution to the EMP exists. Lemma A1 clarifies the exact relationship between the solutions to the EMP and UMP.

Lemma A1 *Suppose that $(\boldsymbol{x}^*, \boldsymbol{m}^*)$ solves the UMP at wealth w and let $U = U(\boldsymbol{x}^*, \boldsymbol{m}^*)$. Then $(\boldsymbol{x}^*, \boldsymbol{m}^*)$ is the solution to the EMP at utility level U and the minimized expenditure level is w. Conversely, suppose that $(\boldsymbol{x}^*, \boldsymbol{m}^*)$ solves the EMP at some feasible utility level U. Then $(\boldsymbol{x}^*, \boldsymbol{m}^*)$ is the solution to the UMP at wealth level $w = \boldsymbol{p} \cdot \boldsymbol{x}^*$ and the maximized utility level is U.*

Proof. The proof exactly mirrors the proof of Proposition 3.E.1 in Mas-Colell et al. (1995) once it is realized that local nonsatiation always holds since there is at least one amoral good. ∎

Lemma A1 establishes the duality between the UMP and EMP and implies that the EMP has a unique solution (because the UMP has a unique solution, the EMP does too by Lemma A1). Let the optimal consumption quantities determined by the EMP be denoted by $\boldsymbol{h}\left(\boldsymbol{p}, \boldsymbol{\mu}, U\right)$ and the minimized expenditure be denoted $\boldsymbol{e}\left(\boldsymbol{p}, \boldsymbol{\mu}, U\right)$. Note that the own-price effect is non-positive for compensated demand, i.e., if p_i increases but other prices remain unchanged, $h_i\left(\boldsymbol{p}, \boldsymbol{\mu}, U\right)$ cannot increase (the proof is exactly the same as the proof of Proposition 3.E.4 in Mas-Colell et al., 1995). The envelope theorem implies that compensated demand is the derivative vector of expenditure with respect to prices (see Proposition 3.G.1 in Mas-Colell et al., 1995). Finally, following the exact same steps as in the proof of Proposition 3.G.3 in Mas-Colell et al. (1995), it follows that the Slutsky equation is

$$\frac{\partial h_i\left(\boldsymbol{p}, \boldsymbol{\mu}, U\right)}{\partial p_j} = \frac{\partial x_i^*}{\partial p_j} + \frac{\partial x_i^*}{\partial w} x_j^*,$$

where U is the utility level at the optimum in the UMP at prices $\boldsymbol{p}$, income w and original moral values $\boldsymbol{\mu}$.

Appendix: Description of Data

The 69 countries in the WVS 1999-2004 wave for which at least one of the moral values questions are available are listed in Table A1. Northern Ireland is included as a separate country. The asterisks indicate the 32 European Values Survey countries where data is available for additional moral values questions.

The categorical variables used in the short regressions are gender (WVS code: x001) and age (WVS code: x003). In the long regression, dummies are included for employment status, educational attainment, occupation, size of home town and number of children of the respondent. Educational attainment (WVS code: x025) is measured in eight different categories, ranging from inadequately completed elementary education to university education with a degree. The employment categories (WVS code: x028) include full-time, part-time, self-employed, retired, housewife, student, unemployed and other. There are 16 occupational dummies (WVS code: x036), for example skilled manual worker, farmer and professional worker. The size of town dummies (WVS code:

Table A1: List of countries

Albania	France*	Macedonia	Singapore
Algeria	Germany*	Malta*	Slovakia*
Argentina	Great Britain*	Mexico	Slovenia*
Austria*	Greece*	Moldova	South Africa
Bangladesh	Hungary*	Morocco	South Korea
Belarus*	Iceland*	Netherlands*	Spain*
Belgium*	India	Nigeria	Sweden*
Bosnia-Herzegov.	Indonesia	Northern Ireland*	Tanzania
Bulgaria*	Iran	Pakistan	Turkey*
Canada	Ireland*	Peru	Uganda
Chile	Israel	Philippines	Ukraine*
China	Italy*	Poland*	USA
Croatia*	Japan	Portugal*	Venezuela
Czech Republic*	Jordan	Puerto Rico	Vietnam
Denmark*	Kyrgyzstan	Romania*	Zimbabwe
Egypt	Latvia*	Russia*	
Estonia*	Lithuania*	Serbia-Montenegro	
Finland*	Luxembourg*	Singapore	

Table A2: Moral values questions

Question	WVS Code	Wording of question
		Please tell me for each of the following statements whether you think it can always be justified, never be justified, or something in between:
Benefit fraud	f114	"Claiming government benefits to which you are not entitled"
Avoiding fare	f115	"Avoiding a fare on public transport"
Joyriding	f125	"Taking and driving away a car belonging to someone else (joyriding)"
Smoking in public	f133	"Smoking in public buildings"
Prostitution	f119	"Prostitution"
Taking soft drugs	f126	"Taking the drug marijuana or hashish"
Drink and drive	f130	"Driving under the influence of alcohol"
Overspeeding	f134	"Speeding over the limit in built-up areas"
Paying cash for services	f131	"Paying cash for services to avoid taxes"
Cheating on taxes	f116	"Cheating on taxes if you have a chance"

x049) contains eight different size brackets for the size of the town where the respondent lives. Finally, the number of children (WVS code: x011) of the respondent are included as dummies.

The wording of the moral values questions are reported in Table A2.

References

Akerlof, G. A. and Dickens, W. T. (1982), "The Economic Consequences of Cognitive Dissonance", *American Economic Review* 72(3), 307–319.

Andreoni, J., Erard, B. and Feinstein, J. (1998), "Tax Compliance", *Journal of Economic Literature* 36(2), 818–860.

Aronson, E. (2003), *The Social Animal*, 9 edn, Worth Publishers, New York.

Babcock, L. and Loewenstein, G. (1997), "Explaining Bargaining Impasse: The Role of Self-Serving Biases", *Journal of Economic Perspectives* 11(1), 109–126.

Beasley, R. K. and Joslyn, M. R. (2001), "Cognitive Dissonance and Post-Decision Attitude Change in Six Presidential Elections", *Political Psychology* 22(3), 521–540.

Bénabou, R. and Tirole, J. (2003), "Intrinsic and Extrinsic Motivation", *Review of Economic Studies* 70(3), 489–520.

Bénabou, R. and Tirole, J. (2006a), "Belief in a Just World and Redistributive Politics", *Quarterly Journal of Economics* 121(2), 699–746.

Bénabou, R. and Tirole, J. (2006b), "Incentives and Prosocial Behavior", *American Economic Review* 96(5), 1652–1678.

Bowles, S. (1998), "Endogenous Preferences: The Cultural Consequences of Markets and Other Economic Institutions", *Journal of Economic Literature* 36(1), 75–111.

Brekke, K. A., Kverndokk, S. and Nyborg, K. (2003), "An Economic Model of Moral Motivation", *Journal of Public Economics* 87(9–10), 1967–1983.

Chaloupka, F. and Warner, K. (2000), "The Economics of Smoking", in A. J. Culyer and J. P. Newhouse, eds, *Handbook of Health Economics 1B*, Elsevier, Amsterdam.

Dickens, W. T. (1986), "Crime and Punishment Again: The Economic Approach with a Psychological Twist", *Journal of Public Economics* 30(1), 97–107.

Ellingsen, T. and Johannesson, M. (2008), "Pride and Prejudice: The Human Side of Incentive Theory", *American Economic Review* 98(3), 990–1008.

Festinger, L. (1957), *A Theory of Cognitive Dissonance*, Stanford University Press, Stanford.

Frey, B. S. (1997), *Not Just For the Money: An Economic Theory of Personal Motivation*, Edward Elgar Publishing, Cheltenham & Brookfield.

Haagsma, R. and Koning, N. (2002), "Endogenous Mobility-reducing Norms", *Journal of Economic Behavior and Organization* 49(4), 523–547.

Hirschman, A. O. (1965), "Obstacles to Development: A Classification and a Quasi-vanishing Act", *Economic Development and Cultural Change* 13(4), 385–393.

Konow, J. (2000), "Fair Shares: Accountability and Cognitive Dissonance in Allocation Decisions", *American Economic Review* 90(4), 1072–1091.

Lieberman, M. D., Ochsner, K. N., Gilbert, D. T. and Schacter, D. L. (2001), "Do Amnesics Exhibit Cognitive Dissonance Reduction? The Role of Explicit Memory and Attention in Attitude Change", *Pscyhological Science* 12(2), 135–140.

Lindbeck, A. (1995), "Welfare State Disincentives with Endogenous Habits and Norms", *Scandinavian Journal of Economics* 97(4), 477–494.

Lindbeck, A. (1997), "Incentives and Social Norms in Household Behavior", *American Economic Review Papers and Proceedings* 87(2), 370–377.

Loewenstein, G. and Schkade, D. (1999), "Wouldn't it Be Nice? Predicting Future Feelings", in Kahneman, D., Diener, E. and Schwarz, N. (eds.), *Well-Being: The Foundations of Hedonic Psychology*, Russell Sage Foundation Press, New York, pp. 85–105.

MasColell, A., Whinston, M. D. and Green, J. R. (1995), *Microeconomic Theory*, Oxford University Press, Oxford.

Mills, J. (1958), "Changes in Moral Attitudes Following Temptation", *Journal of Personality* 26(4), 517–531.

Mullainathan, S. and Washington, E. (2009), "Sticking with Your Vote: Cognitive Dissonance and Political Attitudes", *American Economic Journal: Applied Economics* 1(1), 86–111.

Östling, R. (2008), *Bounded Rationality and Endogenous Preferences*, Ph.D. Thesis, Stockholm School of Economics.

Oxoby, R. J. (2003), "Attitudes and Allocations: Status, Cognitive Dissonance, and the Manipulation of Attitudes", *Journal of Economic Behavior and Organization* 52(3), 365–385.

Oxoby, R. J. (2004), "Cognitive Dissonance, Status and Growth of the Underclass", *Economic Journal* 114(498), 727–749.

Rabin, M. (1994), "Cognitive Dissonance and Social Change", *Journal of Economic Behavior and Organization* 23(2), 177–194.

Saffer, H. and Chaloupka, F. (1999), "The Demand for Illicit Drugs", *Economic Inquiry* 37(3), 401–411.

Shinar, D., Schechtman, E. and Compton, R. (2001), "Self-reports of Safe Driving Behaviors in Relationship to Sex, Age, Education and Income in the US Adult Driving Population", *Accident Analysis and Prevention* 33(1), 111–116.

Shleifer, A. (2004), "Does Competition Destroy Ethical Behavior?", *American Economic Review Papers and Proceedings* 94(2), 414–418.

Titmuss, R. M. (1970), *The Gift Relationship: From Human Blood to Social Policy*, Allen & Unwin, London.

Van de Ven, J. (2003), "Optimal Subsidies with Rationalizing Agents: Subsidize Enough, but Don't Subsidize Too Much", in *Psychological Sentiments and Economic Behaviour*, Ph.D. Thesis, Tilburg University, chapter 7, 157–177.