Demand for moderation of desire

Nathan Berg* and Jeong-Yoo Kim

Abstract: Many philosophical and religious traditions place great emphasis on teaching techniques to moderate ones desire for material consumption. Anyone who has attempted moderation knows that controlling ones desire for consumption often requires effort. This paper introduces a simple model in which the choice problem concerns how to allocate ones endowment into two distinct activities that raise utility through distinct channels. The first channel is material consumption. Applying effort toward acquisitiveness increases the expected level of material consumption. Holding preferences fixed, this increased material consumption increases utility. Alternatively, effort can be directed toward moderating ones desire for material consumption, which raises utility holding the level of material consumption fixed. Is it better to apply effort to increasing what we have, or controlling what we want? It depends on price, as well as risk, and ones aspiration levels that define contentedness. The model provides a characterization of the determinants of high and low demand for moderation of ones desire.

The preference-transformation technology considered here enables individuals to moderate desire by reducing the threshold level of material consumption required to achieve contentedness. Because the re-allocation of effort into self-control requires reductions in effort applied toward material consumption, and because contentedness is defined by material consumption relative to this controllable threshold, unusually strong income effects on the demand for moderation or self-control are present in the model. These strong income effects produce upward sloping demand curves in some cases, and highly unstable demand as a function of price. The paper considers whether these unstable price dynamics provide any insight into religious revolutions in which many peoples intensity of effort into endeavors aimed at moderating desire of material consumption change dramatically over a short time horizon.

Keywords: Self Control, Choosing Preferences, Ecological Rationality, Bounded Rationality

**Berg:* Associate Professor of Economics, School of Economic, Political, and Policy Sciences (EPPS), University of Texas at Dallas, GR 31 211300, Box 830688, Richardson, Texas 75083-0688. Phone 1-972-883-2088, Email nberg@utdallas.edu

Kim : Professor of Economics, Kyung Hee University, Phone , Fax , Email yookim@khu.ac.kr

1 Introduction

Is it better to control what you have conditional on what you want, or control what you want conditional on what you have? The exogenous preference model of neoclassical economics taught in most microtheory textbooks focuses on the first problem, seeking to explain how individuals maximize satisfaction taking desires and the resources available as given. Starkly different views about how humans achieve happiness can be found in Aristotelian (xxx), Buddhist and Confucian (xxx) traditions, which emphasize the second problem of moderating ones desires. Sometimes described in esoteric terms as a process of detaching from material acquisitiveness (xxx), the moderation of desire can be represented in straightforward terms within an otherwise standard utility function framework. Simply put, moderation of desire, or self-control, is the application of a preference-transformation technology that generates positive utility through an alternative mechanism—by transforming ones preferences to yield greater satisfaction from any fixed level of material consumption.

This paper approaches the question of how an individual chooses to amend his or her own preferences by an unusual methodological combination. We use a parsimonious and highly stylized neoclassical approach to pursue the behavioral problem of choosing how much self-control to exert on oneself. This implies that the individuals preferences are, at least partially, a choice variable (alongside quantities of material consumption) rather than the fixed, exogenously given data in the standard neoclassical framework. The model proceeds by applying a strictly neoclassical approach to this modified choice set in which an optimal allocation is chosen that assigns levels of effort to increasing material consumption and to reducing desire for (i.e. increasing satisfaction with any fixed level of) material consumption. We specify the general decision problem, solve the optimization problem for a special case in which most of the key relationships in the model are linear. Then we report comparative statics. Even under this simplest of model specifications, several surprising results emerge. First, moderation is an inferior good, with demand decreasing as a function of the wealth endowment. Second, because of the joint interaction of material consumption and the contentment threshold which is a choice variable under the decision makers control, demand for moderation exhibits unusually strong income effects, which produce upward sloping demand curves for many parameter specifications. Finally, because the globally-maximizing level of moderation switches from interior to corner solutions, the models price and demand dynamics are highly unstable in some regions, in which small changes in price produce dramatic shifts in allocations of effort toward moderation of desire.

2 The model

The model represents an individual's allocation problem regarding how to direct effort toward two distinct activities, both of which can potentially raise utility. The first channel through which effort may raise utility is material consumption, which is easily recognized as a standard assumption. The second effortful activity that can contribute positively to utility is moderation of ones desire for material consumption. Moderation, or self-control, refers to control over what one desires. The preference-transformation technology considered here enables one to reduce the quantity of material consumption required to feel contentment. This is equivalent to transforming ones preferences so that heightened utility is associated with any fixed level of material consumption.

The phrase "material consumption" is used to distinguish consumption of usual goods and services from the purchase of preference-transformation technology (e.g., lessons from a spiritual teacher focused on detachment from material possessions) enabling one to feel content with less, which falls under heading of moderation rather than material consumption. The same would apply to the purchase of a physical tool that aids in feeling content with less (e.g., a yoga mat, a piece of physical equipment used in spiritual ritual, or a self-help book helping addicts to moderate their desires), which would count as self-control and not as material consumption. Material consumption, then, is a residual consumption category, that includes all goods and services that are not aimed at moderating one's desires.

Let the random variable X represent the decision maker's uncertain level of material consumption, with cdf $F(x)$, density $f(x)$, and expected value $\mathrm{E}[X] = \mu$, which depends on the decision maker's allocation of effort into acquisitiveness (i.e., activities that enhance the expected level of material consumption), a. The function $\mu(a) = \mathrm{E}[X]$ describes the dependence of expected material consumption on the individuals choice of a, which we assume is an increasing and differentiable function ($\mu'(a) > 0$).

The other effortful activity that can increase utility—and the main focus of this paper—is denoted s (for self-control).[1] The mechanism by which s moderates desire is now assumed to take

[1] An alternative parameterization of self-control will be introduced in a subsequent section that develops a simplified linear version of the general model. In the simplified model, it is most natural to identify the term self-control with

the specific form of a controllable threshold-level of material consumption by contentment and discontentment are defined. The decision maker can, at the cost of effort which would otherwise go into increasing expected material consumption, lower the threshold that defines discontentment and therefore actively choose how to transform his own preferences. Increasing effort s reduces the probability of experiencing discontentment, all else equal. However, all else is not equal, because applying effort to self-control comes at a real cost in terms of expected material consumption. In this sense, the theory advanced here can be described as a contentment-aspiration theory of self-control.

We define the threshold $t = t(s, t_0)$ as the reference-point level of material consumption below which discontentment occurs. The notation makes explicit the dependence of t on self-control s and an exogenously given default threshold t_0 that specifies the minimum requirements in terms of material consumption for achieving happiness or contentedness. Because s is defined as effort that increases control over one's desires (i.e., reduces the threshold that separates the range of material consumption levels into discontented and contented outcomes), a natural formalization of this mechanism is:

$$\frac{\partial t(s, t_0)}{\partial s} < 0. \tag{1}$$

. It is natural to normalize units of self-control such that $t(0, t_0) = t_0$, implying that the default threshold t_0 defines discontentment whenever zero effort is allocated to self-control.

The magnitude by which material consumption falls short of the threshold t is $\max\{t - X, 0\}$m and utility is then specified as the function $u(X, \max\{t - X, 0\})$, which is increasing in the first argument and weakly decreasing in the second. The key trade-off in the model concerns the reduction in effort applied to acquisitiveness (a) and consequently lower expected material consumption X that results whenever an additional unit of effort is applied toward s to lower the contentment threshold t.

2.1 Translating from units of effort to units of material consumption

Because units of effort can be difficult to pin down, we translate effort variables a and s to more tangible units of material consumption, which also enables us to quantify the opportunity cost of self-control in units of expected material consumption. This translation requires specification of several important relationships.

the resulting distance by which preferences are transformed (in a sense that will be made concrete shortly) as a result of applying effort into self-control rather than s itself.

First, we can identify self-control with the distance by which the contentment threshold is reduced from t_0. This introduces an alternative definition of self-control in units of material consumption:

$$\delta = \delta(s) \equiv t - t_0. \tag{2}$$

Following the earlier definition of t and normalization of $t(0, t_0) = t_0$, it is straightforward that self-control δ satisfies:

$$\frac{\partial \delta(s)}{\partial s} > 0 \text{ and } \delta(0) = 0. \tag{3}$$

To measure the increase in s required to increase δ by one unit (i.e., reduce t by one unit), we need an inverse relationship, δ^{-1}, mapping δ to s, whose first derivative can be represented as $\frac{\Delta s}{\Delta \delta} = \frac{\partial \delta^{-1}}{\partial \delta}$.

To measure the reduction in acquisitiveness a that results from an increase in self-control effort s, implicit differentiation of a smooth aggregate effort constraint is required. This aggregate effort constraint is specified in terms of the differentiable effort aggregation function e and the endowment of aggregate effort parameter A (given exogenously by biological, sociological or institutional variables) as:

$$e(a, s) \leq A. \tag{4}$$

Assuming that this constraint binds, the cost of one unit of s in terms of a can be computed as:

$$\frac{\Delta a}{\Delta s} = -\frac{\partial e(a, s)}{\partial s} / \frac{\partial e(a, s)}{\partial a}. \tag{5}$$

Another measure required to compute the opportunity cost of δ in terms of μ is the sensitivity of μ to acquisitiveness effort a: $\frac{\partial \mu(a)}{\partial a}$, which can also be represented as the ratio of differentials $\frac{\Delta \mu}{\Delta a}$. Combining terms introduced above, the cost of reducing the contentment threshold by one unit of material consumption (i.e., increasing δ by one unit, which requires increased effort expenditure to raise s), expressed in units of forgone expected material consumption (resulting from reduced a), is:

$$\frac{\Delta \mu}{\Delta a} \frac{\Delta a}{\Delta s} \frac{\Delta s}{\Delta \delta} = \frac{\partial \mu(a)}{\partial a} \left(-\frac{\partial e(a, s)}{\partial s} / \frac{\partial e(a, s)}{\partial a}\right) \frac{\partial \delta^{-1}}{\partial \delta}. \tag{6}$$

2.2 General decision problem

Now we can specify the decision-maker's ex ante problem: choosing a and s to maximize expected utility subject to three constraints consisting of production technology, self-control technology, and the effort resource constraint:

$$\max_{(a,s) \in \Re_+^2} \mathrm{E}[U(X, \max\{t_0 - \delta(s) - X, 0\})], \text{ such that } E[X] = \mu(a), \ \delta = \delta(s), \text{ and } e(a, s) \leq A. \tag{7}$$

As long as the production and self-control technologies are defined by one-to-one (i.e., invertible) functions, the choice variables can be re-specified in terms of μ and δ:

$$\max_{(\mu,\delta)\in\Re^2_+} \mathrm{E}[U(X, \max\{t_0 - \delta - X, 0\})], \text{ such that } e(a^{-1}(\mu), s^{-1}(\delta)) \leq A. \tag{8}$$

As the next section shows, with appropriate re-scaling, the effort constraint in (8) can be re-expressed as a budget constraint in units of material consumption, where the opportunity cost in equation (6) is re-parameterized as p, the price of δ in terms of μ.

2.3 Simplified linear model

Suppose that production in material consumption is linear in acquisitiveness effort:

$$\mu = \mu_0 + v_a a, \Rightarrow \frac{\partial \mu}{\partial a} = v_a. \tag{9}$$

Suppose that self-control or preference-transformation technology provides a linear relationship between *change* in the contentment threshold (δ) and self-control effort (s):

$$\delta = v_s s, \Rightarrow \frac{\partial \delta^{-1}}{\partial \delta} = 1/v_s. \tag{10}$$

And suppose that the effort resource constraint is linear, too:

$$a + s = A, \Rightarrow -\frac{\partial e(a,s)}{\partial s} / \frac{\partial e(a,s)}{\partial a} = -1. \tag{11}$$

Then it follows that raising δ (i.e., reducing t) by one unit costs v_a/v_s in units of μ (i.e., $\frac{\partial \mu}{\partial \delta} = -v_a/v_s$), which we refer to as the price of self-control:

$$p = v_a/v_s. \tag{12}$$

2.4 The effort constraint is a budget constraint

With these simplified functional specifications, the effort constraint can be re-written in the easy-to-recognize form of an ordinary budget constraint. Multiplying (11) through by v_a and adding μ_0 to both sides, one obtains:

$$\mu + v_a s = v_a A + \mu_0. \tag{13}$$

The expression on the right-hand side represents the value of the aggregate effort endowment A in units of material consumption, which is equal to the maximum possible level of expected material consumption if all effort were allocated to acquisitiveness, denoted as $M \equiv v_a A + \mu_0$. Finally substituting $p\delta = (v_a/v_s)(v_s s) = v_a s$ produces the budget constraint:

$$\mu + p\delta = M. \tag{14}$$

2.5 Additively separable utility and binary contentment

Three further simplifications are introduced before solving the decision maker's constrained optimization problem and analyzing the statics of demand for self-control as a function of exogenous economic variables. First, the specification below assumes that U is additively separable in material consumption and discontentment. Second, the specification assumes that, apart from the discontentment term, preferences over distributions of X are risk neutral. The third simplification is to ignore the magnitude by which material consumption falls below the threshold t and represent discontentment as a binary indicator variable. This could be motivated by a model of psychic cost in which failing to surpass the threshold imposes a fixed cost in the form of a pervasive feeling of discontent, with high stakes consequences such as failing to be admitted to a select college, failing to attract a marriage partner in class that matches one's aspiration, or even as a change in mood that sours all other consumption experiences by a fixed amount. The resulting utility specification is:

$$U(X, \max\{t - X, 0\}) = X - \beta \mathbf{1}(X < t), \tag{15}$$

where $\beta > 0$ is the model's only preference parameter, which measures the extent to which the decision maker places weight on the possibility of falling below the contentment threshold when weighing benefits of different choices of δ and μ and their associated distributions of X. Note that this utility specification is a special case of the general specification above (8).

Substituting $\mu = M - p\delta$ from the budget constraint (14) simplifies to the constrained optimization problem with two choice variables to a univariate unconstrained objective function:

$$v(\delta) = M - p\delta - \beta F(t_0 - \delta). \tag{16}$$

Under the expectations operator, X is mapped into μ and then substituted in terms of δ, and the indicator function that depends on X is mapped into its expectation (i.e., $\mathrm{E}[\mathbf{1}(X < t)] = F(t) = F(t_0 - \delta)$).

The simple objective function $v(\delta)$ explicitly shows the trade-off at the heart of the model. Self-control as measured by the decision maker's choice of δ provides the benefit of reduced risk of discontentment (i.e., reducing t from t_0 to $t_0 - \delta$), but comes at the cost of a reduced level of expected material consumption μ. When zero effort is allocated to self-control ($\delta = 0$) and 100 percent is allocated to acquisitiveness, the objective function simplifies to $v(0) = M - F(t_0)$. The only way that self-control can provide improvements in utility is by reducing the β-weighted risk of discontentment more than the allocation into self-control costs (through $\mu = M - p\delta$).

It would be a mistake to compute the first-order condition for an interior maximizer based on differentiation of (16) with respect to δ using the notation given there, because there is a subtle dependence of $F(x)$ on δ. As δ changes, so too does $F(x)$, even holding the argument x constant. The indirect dependence of F on μ that is not explicit in the notation in (16) stems from the fact that δ affects μ through the budget constraint ($\mu = M - p\delta$).In the objective function in (16) in which the explicit argument of F also depends on δ, two terms are needed to differentiate $F(t_0 - \delta)$ with respect to δ:

$$\frac{dF(x)}{d\delta} = F'(x)\frac{\partial x}{\partial \mu} + \frac{\partial F(x)}{\partial \mu}\frac{\partial \mu}{\partial \delta} = f(t_0 - \delta)(-1) + \frac{\partial F(t_0 - \delta)}{\partial \mu}(-p). \tag{17}$$

By specifying $F(x)$ such that μ is strictly a position parameter that shifts the distribution function to the right or left (e.g., if X is normally distributed), the dependence of F on μ can be expressed directly in the argument of F.

We now assume that X is normally distributed with exogenously given $\text{Var}(X) = \sigma^2$, which allows us to re-write the cdf in terms of the standard normal cdf $\Phi(z)$ and pdf $\phi(z)$, both of which are independent of μ and σ:

$$F(t) = F(t_0 - \delta) = \Phi(\frac{t_0 - \delta - \mu}{\sigma}) = \Phi(\frac{t_0 - (1-p)\delta - M}{\sigma}), \tag{18}$$

where the last equality follows from substitution of $M - p\delta$ for μ from the budget constraint. After substituting away all appearances of μ (both implicit and explicit) in terms of δ, the objective function becomes:

$$v(\delta) = M - p\delta - \beta\Phi(\frac{t_0 - (1-p)\delta - M}{\sigma}). \tag{19}$$

It turns out that corner solutions occur frequently, and expressions will be presented below stating where zero self-control, $\delta = 0$, and maximum self-control, $\delta = (M - \mu)/p$, are global maximizers.

For an interior maximizer δ^*, $0 < \delta^* < (M-\mu)/p$, the first-order condition is:

$$-p + \beta\phi(\frac{t_0 - (1-p)\delta^* - M}{\sigma})\frac{1-p}{\sigma} = 0. \tag{20}$$

The second-order condition for an interior maximizer is:

$$-\beta\phi'(\frac{t_0 - (1-p)\delta^* - M}{\sigma})(\frac{1-p}{\sigma})^2 < 0, \tag{21}$$

which requires the pdf of X to be strictly increasing at $t_0 - \delta^*$ or, equivalently, that the argument of $\phi(\cdot)$ above is negative:

$$t_0 - (1-p)\delta^* - M < 0. \tag{22}$$

To solve (20) explicitly for δ^* requires taking the inverse of $y = \phi(z) = \frac{1}{(2\pi)^{1/2}}e^{-\frac{1}{2}z^2}$. The range of y values that this function hits is the half-open, half-closed real line segment $(0, \frac{1}{(2\pi)^{1/2}}]$. The inverse relation can be expressed as $z = \pm[-2\log((2\pi)^{1/2}y)]^{1/2}$. By virtue of the second-order condition, however, only left-of-mean values of z are relevant, implying that only the minus (and not the plus) value is needed to re-express the first-order condition (20) as follows:

$$\frac{p}{1-p}\frac{\sigma}{\beta} = \phi(\frac{t_0 - (1-p)\delta^* - M}{\sigma}). \tag{23}$$

This representation of the first-order condition is interpretable as requiring that self-control be chosen in a manner that equates the risk-benefit-weighted price of self-control with its marginal benefit in the form of a reduction in the probability of discontentment.

The inverse of ϕ, which is a relation but not a function, can be functionally identified with the left-most of the elements in its image (which is a set), to form a functional relationship enabling us to apply the inverse function to both sides of (23):

$$-[-2\log((2\pi)^{1/2}\frac{p}{1-p}\frac{\sigma}{\beta})]^{1/2} = \frac{t_0 - (1-p)\delta^* - M}{\sigma}. \tag{24}$$

When computing the inverse of $\phi(\cdot)$, it is crucial that the exogenous parameters lie in the admissible range: $0 < \frac{p}{1-p}\frac{\sigma}{\beta} < \frac{1}{(2\pi)^{1/2}}$. Otherwise, if $p = 0$, then self-control is free and everyone can avoid discontentment with probability one by choosing $\delta = \infty$. If the parameters are outside the admissible range in the other direction, $\frac{p}{1-p}\frac{\sigma}{\beta} > \frac{1}{(2\pi)^{1/2}}$, then self-control is too expensive relative to the β-weighted concern over the possibility of discontentment in the objective function to ever warrant purchase of self-control and, consequently, utility-maximizing δ is 0. This admissibility condition for an interior δ^* can be re-expressed as:

$$p < \frac{1}{\frac{\sigma}{\beta}(2\pi)^{1/2} + 1}, \tag{25}$$

which implies, too, that $p < 1$ must hold. Otherwise, if $p \geq 1$, it would cost more than one unit of expected material consumption to reduce the contentment threshold by one unit, which would be harmful on net to the decision maker: if μ falls by a magnitude greater than δ, then the probability of discontentment goes up, not down.

We can now use (24) to solve explicitly for δ^*:

$$\delta^* = \frac{1}{1-p}(t_0 - M + \sigma[-2\log((2\pi)^{1/2}\frac{p}{1-p}\frac{\sigma}{\beta})]^{1/2}), \tag{26}$$

provided all conditions on the exogenous parameters are met that guarantee a strictly interior maximizer. The exogenous parameters can be stacked in the vector $\theta \equiv [p, t_0, M, \beta, \sigma]$. It is tempting to try partitioning the parameter space in terms of the function $\delta^*(\theta)$ given in (26) and try claiming that zero self-control is optimal whenever $\delta^*(\theta) > 0$ and maximal self-control is optimal whenever $\delta^*(\theta) > M/p$. However, that claim turns out to be false, because of the highly unstable demand function for self-control that maximizes (20). The statics of self-control with respect to exogenous variables θ are analyzed next. The non-monotonic and in some cases highly unstable self-control responses to p are considered in detail with graphical analyses in a subsequent section.

2.6 Statics of $\delta^*(\theta)$

The determinants of high versus low applications of effort toward moderating desire are perhaps the most interesting static to investigate. It turns out that very small changes in p can cause optimal self-control to alternate between maximal and zero self-control. The responses of self-control to p turn out, however, to be the least transparent to analytic calculus-based characterizations. Before moving on to graphical analysis for this investigation, the analytic statics are recorded here. There is an obvious but important caveat to bear in mind when computing statics of $\delta^*(\theta)$, which is that all such statics are zero at interior points of the regions in the parameter space at which corner solutions prevail. The question of where in the parameter space corner solutions prevail is a primary point of interest addressed subsequently. The results reported in this section apply only to the region of the parameter space in which expected utility maximizing choices of δ are on the interior of the budget constraint and strictly positive amounts of both self-control and μ are chosen.

According to formula (26), an increase in a person or country's contentment threshold while holding material resources M and all other exogenous variables constant exerts a strictly positive effect on δ^*:

$$\frac{\partial \delta^*(\theta)}{\partial t_0} = \frac{1}{1-p} > 0. \tag{27}$$

In this model, it is the aspiration level defining contentment that is the primary driver of demand for self-control. When aspirations rise without any increase in resources, then the decision maker falls further behind in the sense of facing an unambiguously increased risk of discontentment, and the decision maker responds by increasing his or her allocation of effort to moderating desire. If, for example, one thinks about unequal levels of material consumption in a cross-section of countries, and increasing availability of information to residents of poorer countries concerning higher living standards in other countries, the model predicts a re-allocation of effort away from material consumption and toward self-moderation. Alternatively, if one thinks about the feudal economies under which organized Christianity developed, or the earlier economies in the East in which Buddhist thought flourished, the model would predict that these expansions of practices that teach allocating effort toward moderation of desire would coincide with increasingly available signals available to poor people about the high levels of material consumption being enjoyed by others (assuming that this information would exert upward pressure on the threshold t_0).

The effect of greater wealth (e.g., advances in labor-saving technology that increase the value of all individuals' endowment of total effort) exerts an unambiguously negative effect on the demand for moderation of desire:

$$\frac{\partial \delta^*(\theta)}{\partial M} = -\frac{1}{1-p} < 0. \tag{28}$$

Thus, in this model, self-control is unambiguously an inferior good. This negative income effect on demand for the moderation of desire is frequently strong enough to produce upward-sloping demand curves, which we return to in the next section.

Interestingly, if the contentment threshold and the real value of material resources grow by equal amounts, the two effects cancel each other. Changes in aspirations needed for contentment and changes in wealth only have an effect on self-control net of change in each other. In other words, shifting aspirations have effects on self-control when it changes by more than wealth, and wealth effects are important only if contentment aspirations do not track those changes in wealth.

For computing signs of derivatives, it is useful to denote the bracketed term in (26) as the function $h(y) \equiv [-2\log((2\pi)^{1/2}y)]^{1/2} = z$ (with y evaluated at $\frac{p}{1-p}\frac{\sigma}{\beta}$), which maps standard normal pdf values (y) into z on the positive real line. One notices that $h(y)$ is a positive valued and strictly decreasing function. Taking advantage of the inequality $h'(y) < 0$, it is straightforward to verify the intuitively obvious result that an increase in β, reflecting stronger subjective weight on the possibility of discontentedness, causes demand for self-control to increase, all else equal:

$$\frac{\partial \delta^*(\theta)}{\partial \beta} = -h'(\frac{p}{1-p}\frac{\sigma}{\beta})p(\frac{1}{1-p}\frac{\sigma}{\beta})^2 > 0. \tag{29}$$

The parameter σ represents the volatility or riskiness of material consumption. Another interpretation of σ is as a particular form of wealth inequality that is orthogonal to effort, since the distribution of material standards of living will be larger in environments with large σ, even if everyone were to choose the same level of effort. To characterize the sign of $\frac{\partial \delta^*(\theta)}{\partial \sigma}$, it is useful to notice that $h'(y) = -\frac{1}{yh(y)}$ and apply this result to simplify the following expression:

$$\frac{\partial \delta^*(\theta)}{\partial \sigma} = \frac{1}{1-p}[h(\frac{p}{1-p}\frac{\sigma}{\beta}) + \sigma h'(\frac{p}{1-p}\frac{\sigma}{\beta})\frac{p}{1-p}\frac{1}{\beta}] = \frac{1}{1-p}[h(\frac{p}{1-p}\frac{\sigma}{\beta}) - [h(\frac{p}{1-p}\frac{\sigma}{\beta})]^{-1}]. \tag{30}$$

The sign of $\frac{\partial \delta^*(\theta)}{\partial \sigma}$ depends on whether $h(\frac{p}{1-p}\frac{\sigma}{\beta})$ is greater or less than 1. By definition, $h(\phi(z)) = z$ for positive z. Therefore, the sign depends on whether $\frac{p}{1-p}\frac{\sigma}{\beta}$ is greater or less than $\phi(1) \approx 0.2420$. Keeping in mind that h is decreasing and therefore that $h(y) < 1$ if $y > \phi(1)$ and $h(y) > 1$ if $z < \phi(1)$, the direction of the desired effect can be characterized as:

$$\frac{\partial \delta^*(\theta)}{\partial \sigma} > 0 \text{ whenever } \frac{p}{1-p}\frac{\sigma}{\beta} < \phi(1), \text{ and } \frac{\partial \delta^*(\theta)}{\partial \sigma} < 0 \text{ whenever } \frac{p}{1-p}\frac{\sigma}{\beta} > \phi(1) \tag{31}$$

Thus, the effect of σ on $\delta^*(\theta)$ is nonmonotonic. For moderate levels of σ and p, the effect is positive, indicating increased demand for self-control when risk and inequality rise. The intuition for this is that, in low-risk low-inequality environments, a small increase in that risk or inequality increases the probability of discontentment on a relatively steep range of the pdf ϕ, implying that a response in the form of increasing δ has a proportionally large effect. In contrast, when beginning on a flat range of the pdf farther than one standard deviation from the mean, changing δ has a much weaker effect in terms of reducing the risk of discontentment.

Finally, the effect of p on $\delta^*(\theta)$ is, once again, nonmonotonic, as the figures in the next section will show. We record the analytic comparative static here:

$$\frac{\partial \delta^*(\theta)}{\partial p} = \frac{1}{(1-p)^2}[t_0 - M + \sigma h(\frac{p}{1-p}\frac{\sigma}{\beta}) + \frac{1}{1-p}\frac{\sigma^2}{\beta}h'(\frac{p}{1-p}\frac{\sigma}{\beta})]. \tag{32}$$

This can, again, be developed by using $h'(y) = \frac{1}{yh(y)}$. It is unclear, however, whether this leads to any new insights about regions of parameter space in which demand for self-control is upward- versus downward-sloping.

2.7 Demand for self-control

Figure 1 presents nine ordinary demand curves for self-control, or moderation. The x-axis is δ^*, the distance by which the individual chooses to shift his or her contentment threshold down from its default. The y-axis is the full range of prices from the upper bound given in (25) to zero. The resource endowment M is normalized to 1 for all graphs presented below. In later Figures, the contentment threshold default t_0 is varied up and down by two standard deviations. In Figure 1, t_0 is normalized also to 1.

Moving from top to bottom in Figure 1, the random component of material consumption X is increasing. The x-axes are re-scaled in the three rows to better see the detail of the curves. The nonmonotonicity of $\frac{\partial \delta^*(\theta)}{\partial \sigma}$ is evident moving down the first column of curves in Figure 1 in the downward vertical shift and simultaneous shifting out to the right. The downward vertical shift shows that at high prices, increasing σ causes quantities demanded to shift inward to the y-axis. At low prices, however, increasing σ causes demand to shift outward to the right. These effects are not artifacts of scaling the plot, because quantity demanded is exactly zero at prices above the points at which these demand curves intersect the y-axis.

Moving from left to right in Figure 1, the weight on discontentment in the objective function is increasing, and demand for self-control shifts out. The combination of low σ and large β produces demand curves with an upward-sloping portion, reflecting the unusually large income effects on the demand for moderating material consumption. The northeast demand curve in Figure 1 shows a severe instability in quantity demanded in the price range near $p = 0.9$. In this range, quantity demanded can shift from 0 to the maximum amount possible within a tiny price range, which is suggestive of religious revolutions whose philosophies depend heavily on a critique of excessive focus on material consumption. We return to the question of whether these large price effects in opposite directions might provide an explanation for the emergence of religious and spiritual movements observed in the historical record.

Figure 2 illustrates another set of demand curves, this time with a contentment threshold default that is very easy to reach: $t_0 = M - 2\sigma$. This corresponds to an environment in which more than 95 percent of people will achieve contentment with zero effort allocated to moderating desire. In this case, the possibility of discontentment is remote, and no unusual features are evident in the resulting demand curves.

Figure 3 illustrates an analogous set of environments but with a contentment threshold default that is hard to reach: $t_0 = M + 2\sigma$. In this case, if zero self-control is demanded ($\delta^* = 0$), then more than 95 percent of people will face discontentment, or the individual faces an almost certain risk of discontentment. Low-σ environments in which discontentment is virtually certain produce demand curves with large upward-sloping regions. The strong income effects are intuitive, because the high degree of certainty of discontentment on the convex portion of the pdf produces something tantamount to increasing returns to self-control. The more self-control is applied, the more productive self-control becomes at reducing the probability of discontentment.

Next we return to instability in quantity demanded along a single demand-for-self-control curve. We first plot expenditures-price curves to investigate what fraction of total income is allocated to self-control over the price range. Figure 4 presents expenditure-price curves. Recalling the budget constraint $\mu + p\delta = M$ with M normalized to 1, the quantity $100p\delta$ gives the percentage of the wealth endowment allocated to self-control. Figure 4 presents 9 curves with identical parameter values as Figure 1, showing that increasing and then decreasing expenditures on self-control as prices rise from zero is the rule rather than the exception. When price is very low, a large quantity of self-control can be purchased for a small expenditure, moving the contentment threshold to the concave portion of the pdf of X. When price is very large, the sacrifice in terms of μ is so great that only a small portion of the endowment is allocated to self-control. In the middle of the price range, however, a larger and larger share of wealth is allocated to self-control, reflecting movement along the convex (increasing-returns) portion of the pdf until the concave portion is reached, at which point total expenditures decline.

The root cause of instability, in the form of quantities demanded that move from zero to a substantial share of wealth in response to a small movement in price, is illustrated in Figure 5, which shows the univariate objective function $v(\delta)$ from equation (19) at three nearby values, $p = 0.70$, 0.72, and 0.74. The values of all other parameters are as in Figure 1, except for $t_0 = 0.5$ to reflect a scenario in which discontentment is defined by all levels of material consumption more than half a standard deviation below the default mean of X, μ_0. This ensures that both motives—to apply effort to increasing the mean of X and to decrease the threshold at which discontentment is defined—are reflected in the objective function. The three $v(\delta$ curves shift smoothly as price increases from 0.70. However, the v-maximizing value of δ shifts, first smoothly in the first two plots with interior maximizers, and then discontinuously to zero as the maximized objective function value approaches the zero-self-control objective function value $v(0)$.

3 Discussion

When the cost of self control is above a particular threshold (i.e., self control is too expensive in terms of the amount of expected consumption that the foregone effort in material consumption could have produced), then the individual chooses a wholly material path to happiness—that is, zero effort allocated to self control and maximum effort allocated to raising the level of material consumption. Below this threshold, however, which depends on both taste parameters and the distribution of stochastic material consumption given by the external environment, individuals allocate some effort away from material consumption and into self control.

If one views different cultures' philosophical traditions through the lens provided by this model, the prediction is a clean splitting between cultures, some of which are wholly committed to material consumption while others emphasize allocating effort to self control that moderates desire for material consumption. The reference-point level of material consumption plays a crucial role creating the condition under which the corner solution of zero self-control is chosen.

The model challenges a basic definition of bounded rationality often put forward in the behavioral economics literature based on three bounds: on willpower, computational capacity, and self-interest. Instead, the model describes an outcome in which agents and the cultures they populate explicitly choose to abandon self control, and this abandonment of self control is a rational and predictable function of opportunities and costs in the environment.

Although the model is behavioral in that it presumes willpower is costly and that preferences can be changed, the price theory analysis is neoclassical in flavor.

Figure 1: Ordinary Demand for Self-Control

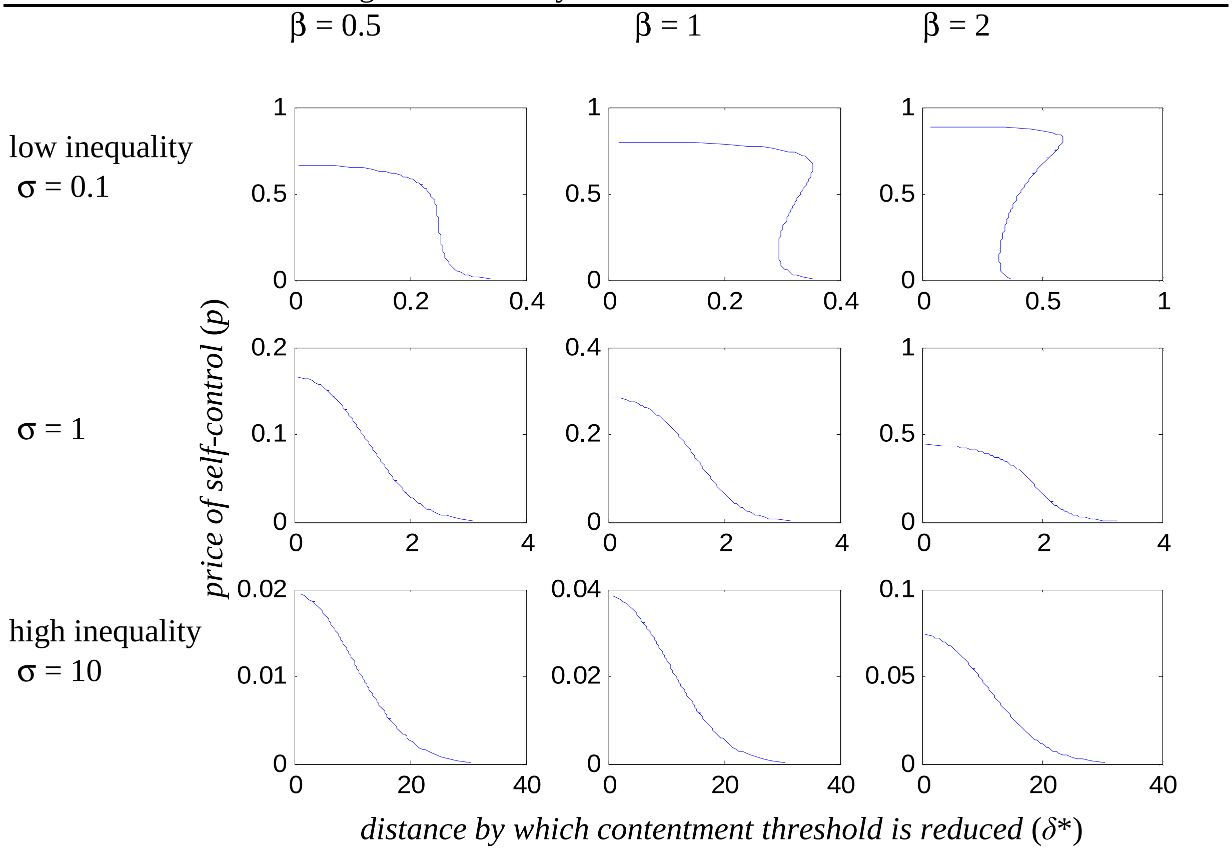

The demand curves above plot expected-utility maximizing values of delta over prices ranging from the upper bound in (25) to zero, with t_0 = 1 and M = 1.

Figure 2: Ordinary Demand for Self-Control with Easy-to-Reach Contentment Threshold Default ($t\_0 = M - 2\sigma$)

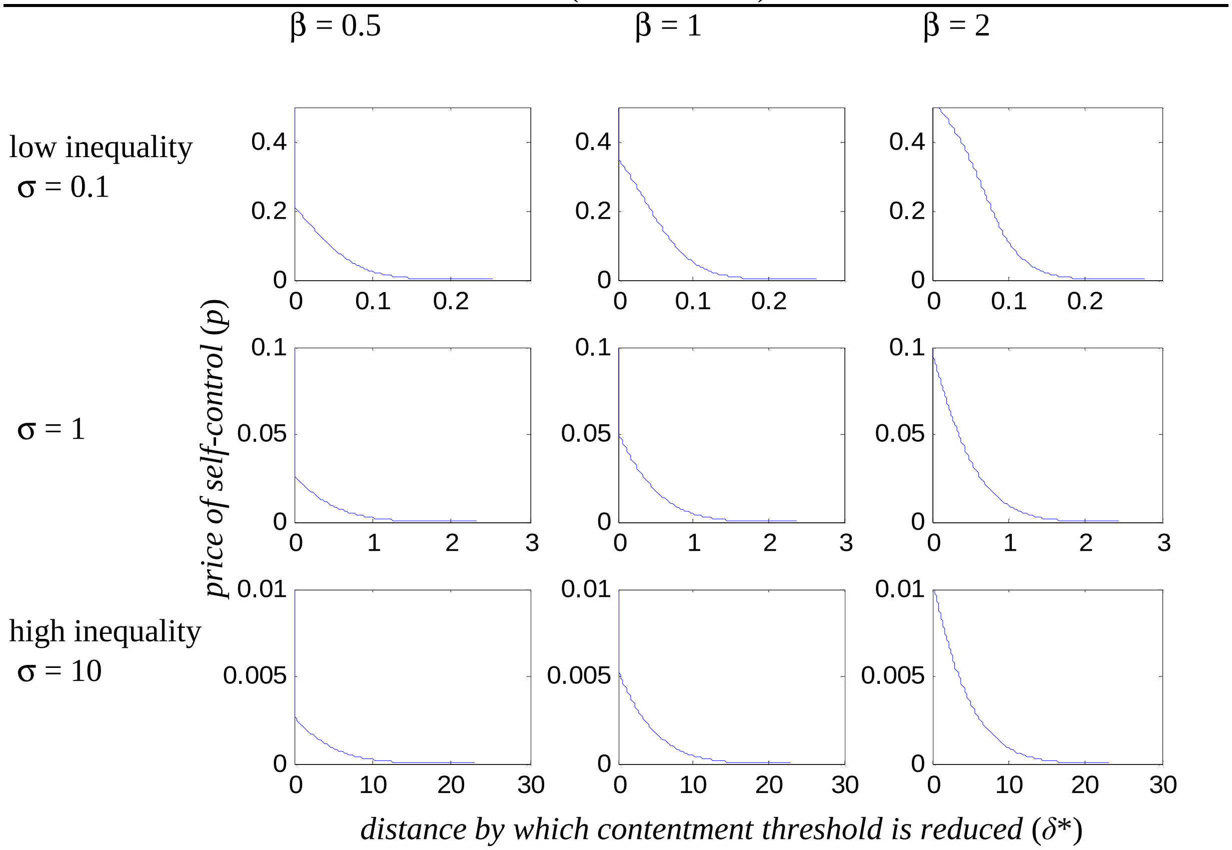

Figure 3: Ordinary Demand for Self-Control with Hard-to-Reach Contentment Threshold Default ($t\_0 = M + 2\sigma$)

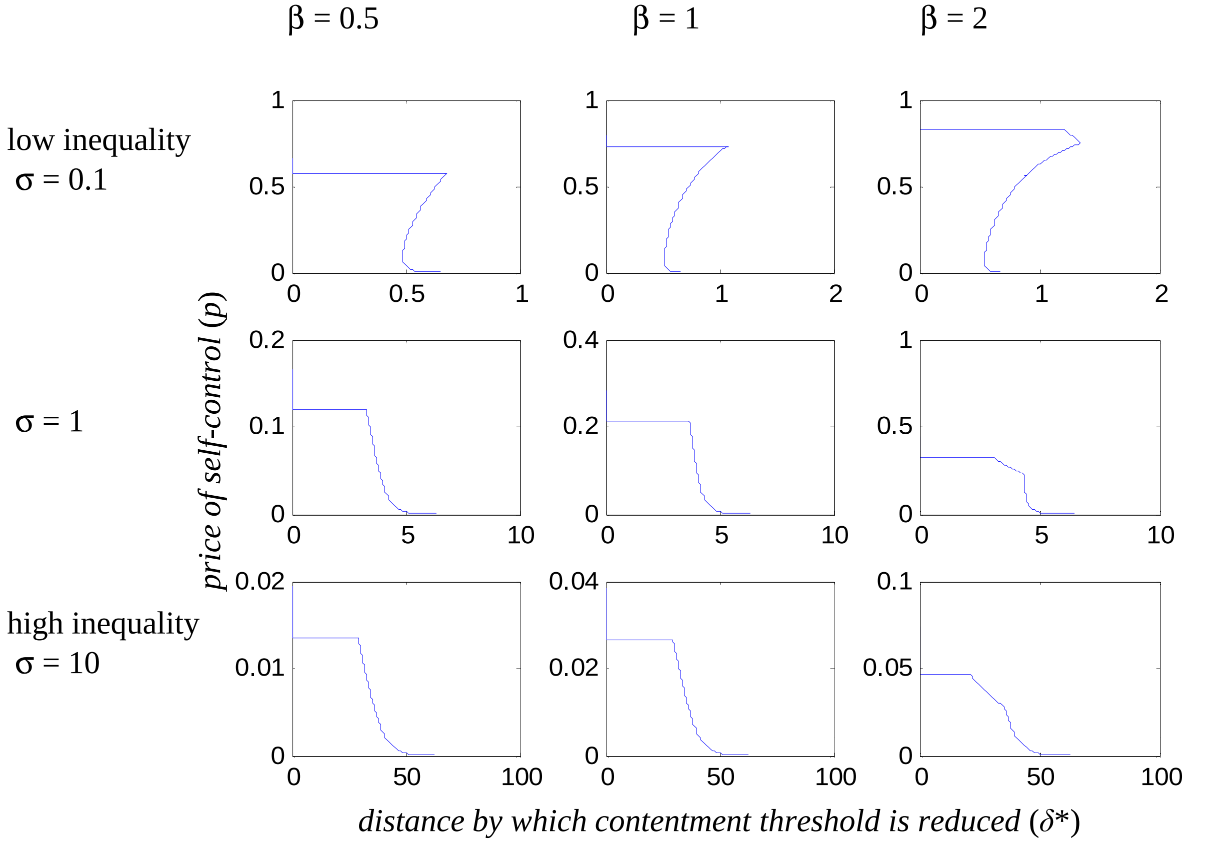

Figure 4: Percentage of *M* Expended on Self-Control and Price ($t\_0 = M$)

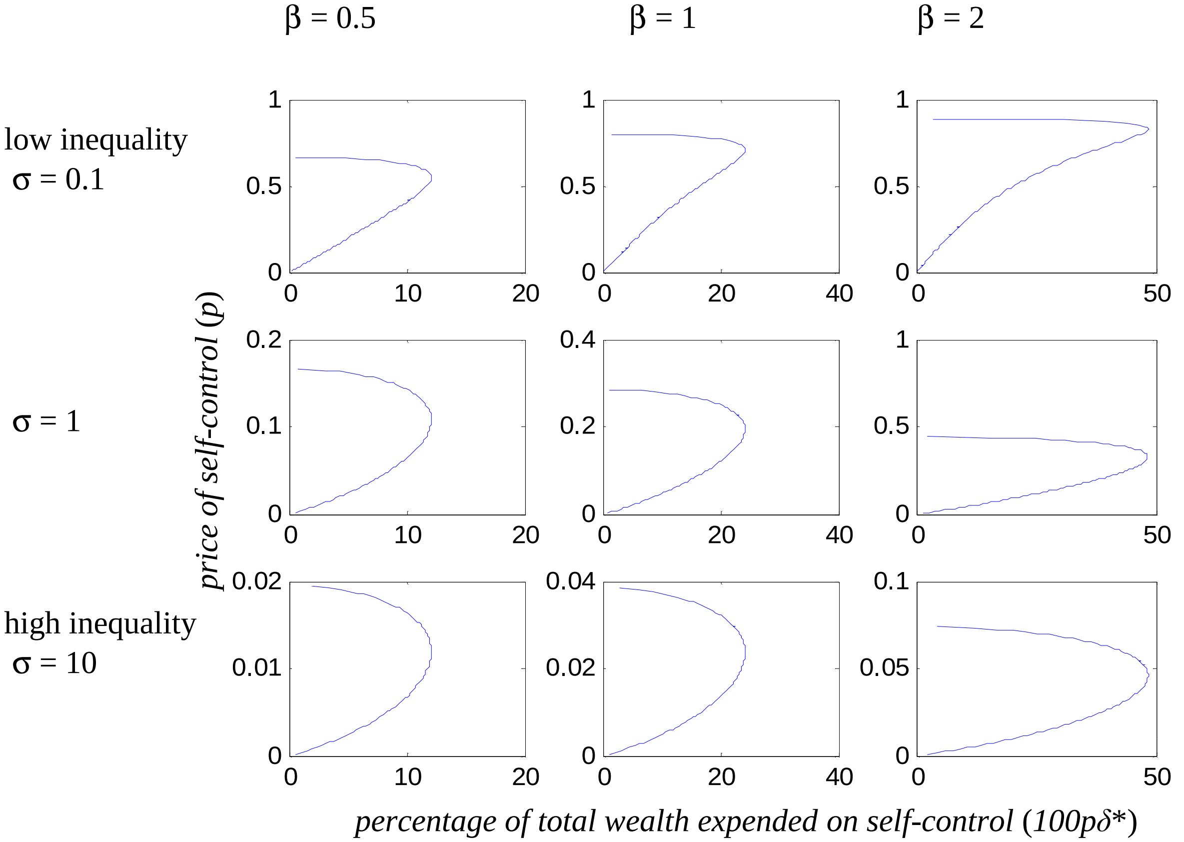

Figure 5: Unstable Demand for Self-Control: Large Fluctuation in Optimal Effort Allocated to Self-Control in Response to Small Changes in Price