A General Framework
of
Flawed Decision Making

Master Thesis
Department of Economics
Lund University
June 2014

Author: Thimmy Bågfelt
Supervisor: Håkan J. Holm

This text would not have been possible without the unwavering support and encouragement of my father, Jan-Olof Johansson, to whom I am forever grateful.

Abstract

Ubiquitous empirical evidence suggests that economic agents strive towards but often fail to maximize utility. Yet there is no consensus on how to best model deviations from optimal behaviour. This paper reviews and analyses the foundation of the utility maximization model and proposes a more general framework that is consistent with – but also allows for deviations from – utility maximization. The framework assumes a complete, reflexive, transitive and continuous binary preference relation, guaranteeing the existence of a utility function. This approach provides a clear theoretical perspective on behaviour that appears to be inconsistent with utility maximization. When applied to a standard supply-demand schedule, the framework implies that market clearing quantities are inefficiently low, even when agents make optimal decisions on average but deviate from optimal decisions in general.

Table of Contents

1. Introduction

Consumers spend more money when given five $1 bills than when given one $5 bill (Mishra et al 2006 and Raghubir & Srivastava 2009). People who are given a mug are willing to sell it only if they get twice as much as they would have paid for it, had it not been given to them (Kahneman et al 1991). Morning sunshine predicts increased stock returns (Hirshleifer & Shumway 2003). Evidence of bad choices on markets is mounting.

Still, the idea that people are perfect utility maximizers prevails in economics, for good reasons. A vast body of mainstream economic theory is based on the crucial assumption that agents consistently maximize a stable utility function representing their preferences. This paper reviews the utility maximization model (section 2.1) and addresses the merits and weaknesses of the assumptions on which it relies, i.e. reflexivity, completeness (section 2.2), transitivity (section 2.3) and continuity (section 2.4) of preferences.

A theoretical discussion about how the utility maximization model can be rejected is complemented by empirical evidence suggesting that utility indeed is not maximized in general (section 2.5). Empirical evidence also suggests that what appears to be deviations from optimal behaviour cannot be fully explained by measurement errors (section 3.1). This view seems to be shared by economists of all schools, but the extent to which deviations from utility maximization should be in focus is controversial.

The purpose of this paper is to investigate how the utility maximization model may fail, by reviewing both logical arguments and empirical evidence for and against the axioms that constitute the foundation of the utility maximization model. Some of the models that have been suggested as alternatives or variations of the utility maximizing model are discussed. There is plenty of interest in modelling the failure to maximize utility, but existing models do not serve as a common foundation for both perfect and flawed decision making (section 3.2).

Another purpose is to propose a general, realistic, useful and mathematically clear framework that is consistent with both the utility maximization model and the commonly observed failure to maximize utility. The suggested framework assigns a probability of choice to every bundle in a budget set, and allows researchers to choose what cognitive biases to include or exclude in their model (section 3.3). By itself, it fails to explain how the effect of any particular cognitive bias on choice is quantified or estimated. Finally, the paper aims to illustrate the implications of deviations from optimality in the demand-supply schedule of standard microeconomics.

Using fairly simple mathematics, the lessons of the framework is applied to the standard supply-demand schedule, leading to the conclusion that deviations from optimality leads to inhibition of trade (section 3.4). Sub-optimal consumers can only demand too low quantities because of their budget constraints. Even when both consumers and firms deviate from optimal decisions in general but make optimal decisions on average, inefficiently low trade ensues. This effect is especially large on small markets, but exists also on large markets.

Finally, a few simulations are conducted to demonstrate how stochastic supply, demand and resulting trade volume can be illustrated as probability density functions (section 3.5). It is made clear from the illustrations that the distribution of market clearing quantity differs significantly from the distributions of demand and supply. Thus, the suggested framework is useful beyond it outlining a clear mathematical representation and a general perspective on flawed decision making.

2. Utility Maximization

2.1. Review of the Utility Maximization Model

The following notation and reasoning is largely based on Mas-Collel (1995). Let X be a consumption set where all elements are mutually exclusive bundles that at least in principle are consumable. Let $\succsim$ be a preference relation defined on X, such that $x \succsim y$ means that an agent prefers x to y or is indifferent between the two. Then let $x \succ y$ represent the relation $x \succsim y \wedge \neg(y \succsim x)$ and $x \sim y$ represent the relation $x \succsim y \wedge y \succsim x$. Jaffray (1975) proves that the completeness, reflexivity, transitivity and continuity of $\succsim$ are sufficient conditions for the existence of a continuous utility function $u : X \rightarrow \mathbb{R}$ representing $\succsim$, so that $x \succsim y \Leftrightarrow u(x) \geq u(y)$ $\forall x, y \in X$. The proof is lengthy and complicated and will therefore not be reproduced here. Completeness, reflexivity, transitivity and continuity of preferences are defined as

$$x \succsim y \ \vee \ y \succsim x \quad \forall x, y \in X \qquad \text{(Completeness)}$$

$$x \succsim x \quad \forall x \in X \qquad \text{(Reflexivity)}$$

$$x \succsim y \wedge y \succsim z \ \Rightarrow \ x \succsim z \quad \forall x, y, z \in X \qquad \text{(Transitivity)}$$

$$\{x \in X : x \succsim y\} \text{ and } \{x \in X : y \succsim x\} \text{ are closed sets } \forall x, y \in X . \qquad \text{(Continuity)}$$

The self-interest of an agent implies that she chooses to consume a bundle that she prefers to all other feasible bundles. Formally, the individual chooses to consume an arbitrary bundle in the set $C(\succsim, B) = \{x \in B : x \succsim y \ \ \forall y \in B\}$, where $C(\cdot)$ is referred to as a choice rule and $B \subseteq X$ is a budget set consisting of all the bundles that are feasible to the agent. That is, $C(\cdot)$ is a procedure that the agent uses to generate a set of most preferred bundles $C(\succsim, B)$ from given preferences $\succsim$ and a budget set B.

Seeing that the agent is indifferent between all the bundles in $C(\succsim, B)$, it is reasonable to assume that the agent randomizes between the bundles in $C(\succsim, B)$. That is, whenever $C(\succsim, B)$ contains several bundles, the agent can reasonably choose an arbitrary bundle in $C(\succsim, B)$. Symbolically, if $p(x)$ is the probability of the agent choosing bundle x, then $p(x) > 0$ $\forall x \in C(\succsim, B)$ and

$p(x)=0 \ \ \forall x \notin C(\succsim, B)$. A slightly stronger version of this argument is that the agent is equally likely to choose any bundle in $C(\succsim, B)$, again because the agent by definition is indifferent between all bundles in this set. Symbolically, if $|C(\succsim, B)|$ is the cardinality of $C(\succsim, B)$, the prediction is that $p(x)=1/|C(\succsim, B)| \ \ \forall x \in C(\succsim, B)$ and $p(x)=0 \ \forall x \notin C(\succsim, B)$.

This abstraction does little more than to ensure a clear understanding of the theory for when the reader has reason to doubt the existence of a utility function. If a utility function indeed exists, the maximization of utility gives the same set while being mathematically more convenient. The utility maximization problem is $\max_x u$ subject to $x \in B$, and like $C(\succsim, B)$, the set $\{x \in B : u(x) \geq u(y) \ \ \forall y \in B\}$ contains all feasible bundles preferable to all other feasible bundles. If B is a compact set, the extreme value theorem implies that $\exists x \in B : u(x) \geq u(y) \ \ \forall \ y \in B$ and that $C(\succsim, B) \neq \emptyset$. The theory presented so far is the underlying core of most economics, and will hence forth be referred to as the utility maximization model.

As will be explained in further detail, ubiquitous empirical evidence suggests that people often fail to act in ways that are consistent with the utility maximization model. The inconsistency indicates that at least one assumption in the model is wrong. First, the assumptions of completeness, reflexivity, transitivity and continuity are discussed in an effort to investigate the extent to which relaxations of these assumptions can explain behaviour inconsistent with the utility maximization model.

2.2. Completeness and Reflexivity

Suppose that the completeness assumption is relaxed and that $\neg(x \succsim y) \wedge \neg(y \succsim x)$ for some $x, y \in X$. If an agent then is forced to choose between x and y, the agent can arbitrarily randomize between x and y and has no reason to deviate from such a strategy. The same reasoning is equally valid for when preferences are complete and $x \succsim y \wedge y \succsim x$ for some $x, y \in X$. Seeing that an indifference relation and a lack of preference relation results in identical behaviour, there seems to be no way of empirically testing whether an agent is indifferent between two bundles or if the two bundles are incomparable.

Aumann (1962), Bewley (2002), Eliaz & Ok (2006) and Mandler (2009) reject this view and offer models that differentiate between $\neg(x \succsim y) \wedge \neg(y \succsim x)$ and $x \succsim y \wedge y \succsim x$. They deduce their models from assumptions made about the nature of preferences and comparison. These assumptions vary between authors and each relies on a different psychological interpretation of incompleteness. The proposition $\neg(x \succsim y) \wedge \neg(y \succsim x)$ is described as “indecisiveness” by Eliaz and Ok (2006) and as “unranked” by Mandler (2008). A deeper psychological and philosophical account of the foundation of incompleteness is offered by Raz (1986) and Chang (2002) who use the terms “incommensurability” and “indeterminacy” respectively to describe incomplete preferences. Danan and Ziegelmeyer (2006) conduct an experiment concluding that preferences are incomplete because agents are willing to pay a small cost to postpone decisions.

If preferences indeed are incomplete, it remains unclear how $\neg(x \succsim y) \wedge \neg(y \succsim x)$ should be interpreted and how such lack of a preference relation can be measured. For example, in Danan and Ziegelmeyer's (2006) experiment, an agent may want to postpone a choice between x and y when $x \succsim y \vee y \succsim x$ because the agent expects his ability to make a decision consistent with his preferences to improve over time. Critics of the completeness axiom often point out the difficulty of agents to make comparisons between bundles that are either drastically different or very similar, but such difficulty is irrelevant if the comparison is made unconsciously. This ambiguity about the manifestation of incomplete preferences is likely the main reason that empirical evidence asserting violations of completeness is sparse and weak.

Nevertheless, there are situations in which completeness might be violated. For example, suppose that an agent has preferences such that car x is preferred to car y if and only if car x is both safer and faster and vice versa. Then, if car x is safer but slower than car y, neither $x \succsim y$ nor $y \succsim x$ is true. Alternatively, consider the more extreme case of a parent forced to choose between the death of either of her two children x and y. One can easily imagine that both alternative are so devastating to the parent that a comparison is not applicable. Now suppose that child x but not child y has life insurance. A parent that is indifferent between the death of child x and child y should then prefer the death of child x coupled with the insurance money to the death of child y without money. However, one can easily imagine that the difference in money fails to make the comparison any easier, which may indicate that x and y cannot be compared.

In any case, even if completeness is violated, Peleg (1970) proves that completeness is not necessary for the existence of a utility function representing preferences. In other words,

$x \succsim y \Rightarrow u(x) \geq u(y) \quad \forall x, y \in X$ still holds without completeness, but $u(x) \geq u(y) \Rightarrow x \succsim y$ $\forall x, y \in X$ is no longer true. Utility maximization can still be performed, but with a different interpretation. With completeness, a utility maximizing bundle is preferred to all other bundles, but without completeness, no bundles are strictly preferred to the utility maximizing bundle. Thus, whether preferences are complete or not is a minor concern for the purposes of this text because the utility maximization model remains largely intact even without completeness.

Reflexivity is sometimes claimed to follow directly from completeness (e.g. Mas-Colell 1995:6). While it is true that $x \succsim y \vee y \succsim x \quad \forall x, y \in X$ implies $x \succsim x$ if x is allowed to be equal to y, the normal interpretation of the completeness assumption is that cases with $x = y$ are not considered, meaning that completeness does not imply reflexivity. For example, the commonly known binary relation $>$ defined on any subset of $\mathbb{R}$ is complete because $x > y \vee y > x$ for all distinct x and y in that subset, but irreflexive because $x > x$ is false.

Regardless of whether reflexivity follows from completeness or not, $x \succsim x$ is probably as close to self-evident as any claim in economics can get. Imagining an agent with preferences such that x is not at least as good as itself is difficult, and if at all possible, such a scenario would deserve more attention in philosophy than in economics. For this reason, no further attention will be given to the possibility of preferences not being reflexive.

2.3. Transitivity

Fishburn (1991) summarizes the most common arguments for transitive preferences and then demonstrates that all of them are questionable reasons for assuming transitivity. Consider for example the money pump argument, where we assume that $x \succ y$, $y \succ z$ and $z \succ x$. An agent with such preferences would reasonably be willing to pay something to obtain y in place of z because $y \succ z$. Having made such a trade, the agent would also be willing to pay something to obtain x in place of y because $x \succ y$. Lastly, the agent would be willing to pay something to obtain z in place of x because $z \succ x$. However, trading in this manner, the agent pays to end up with the same bundle that she started with, which seems silly.

Fishburn agrees that it would be silly for an agent to lose money going through such a sequential cycle of pairwise decisions, but argues that the reasonable agent will only do so if she is deceived.

Fishburn goes on to suggest that the agent may use some choice rule other than choosing the most preferred bundle in pairwise sequential comparisons. For example, Quirk & Saposnik (1968:15) suggest that the agent randomizes between bundles when preferences are intransitive. Alternatively, one can easily imagine a choice rule where the initially endowed bundle is kept whenever it is an element in a set with intransitive preferences. Thus, the money pump argument is not as convincing as it first appears.

Intransitive preferences are logically possible if and only if there are at least three bundles in X, defined on at least two dimensions (Navarick & Fantino 1974, Ng 1977). This makes decision theory with violations of the transitivity assumption intrinsically complicated. In order to see how intransitive preferences may arise, consider as an example the budget set $\{x, y, z\}$, where each element represents a car defined on three dimensions with the characteristics seen in the following table.

	Speed	**Safety**	**Comfort**
Car x	Medium	Good	Poor
Car y	Poor	Medium	Good
Car z	Good	Poor	Medium

Table 1: If one car is strictly preferred to another if and only if it is better than the other car in two or more dimensions, preferences are intransitive.

The agent's preferences are such that one car is strictly preferred to another if and only if it is better than the other car in two or more dimensions. From the table above, pairwise comparisons of cars then make it clear that $x \succ y$, $y \succ z$ and $z \succ x$, implying that preferences are not transitive. In this example, dimensions are aspects of a single kind of good, but the dimensions may also be the quantities in a bundle of several commodities.

Numerous researchers have elicited what appears to be intransitive preferences in humans and animals by devising experiments where options explicitly or implicitly are defined in several dimensions (May 1954, Tversky 1969, Navarick & Fantino 1972, Loomes et al 1991, Sharif 1994, Roelofsma & Read 2000, Waite 2001, Lee et al 2009). Intransitive preferences may be pervasive also outside of the laboratory. Tversky (1969) gives an example of a car buyer who finds each add-on feature worth the slightly higher price, but then ends up preferring the original car with no add-on features to a more expensive car with many add-on features. Roelofsma & Read (2000) and Andreou (2007) propose that procrastination may be due to intertemporal intransitivity. Nooteboom (1984) argues that intransitive preferences are likely to be common in the retail industry because choices are made based on a wide range of dimensions such as proximity, cheapness and service attraction.

Despite abundant evidence against transitive preferences, the transitivity assumption prevails mainly for two reasons. First, the existent empirical evidence is sometimes criticized for being flawed or insufficient. For example, Regenwetter et al (2011) review 20 studies of intransitive preferences and conclude that most of the data is consistent with transitive preferences when considering the stochastic nature of the choices. Birnbaum & Schmidt (2008) find that intransitive patterns disappear in experiments when options are clarified and decisions are repeated, suggesting that the alleged evidence against transitivity is weak or has limited applicability in a real market outside of the laboratory.

The second and probably most important reason for relying on the transitivity assumption in spite of evidence against it, is that the disposing of the transitivity assumption would cause economics as we currently know it to completely collapse. Transitivity is a necessary condition for the existence of a utility function, and many economic models would be nonsensical without it. Thus, potential violations of the transitivity assumption is a major concern in the construction of the utility maximization model and for economics at large.

2.4. Continuity

It is intuitively best understood as the absence of jumps in preferences. To be clear however, a binary relation on a discrete set often is continuous. For example, if X is cars consumable only in non-negative integer amounts so that $X = \{0,1,2,...\}$, then $\succsim$ is continuous because all subsets of $\{0,1,2,...\}$ are closed.

Continuity is assumed because completeness, reflexivity and transitivity are not sufficient conditions for the existence of a utility function, but with the addition of a continuity assumption, a continuous utility function is guaranteed to exist. Continuity is not a necessary condition for the existence of a utility function and is not as crucial as the other assumptions. For example, if $X=[0,1]$ and $x \sim y \succ 0\ \forall x,y \in (0,1]$, then $\succsim$ is not continuous because $\{x \in X : y \succsim x\}$ is not a closed set . Still, any utility function $u : X \rightarrow \mathbb{R}$ with $u(x)=u(y)>u(0)\ \ \forall x,y \in (0,1]$ represents such preferences. In fact, if $\succsim$ is complete, reflexive and transitive, then any finite consumption set X can be represented by a utility function (Mas-Colell 1995:9).

Empirical evidence suggests that preferences often are discontinuous where price is zero. Shampanier et al (2007) show experimentally that people prefer a free good to a very low priced good so much and in such combinations that it suggests that preferences are discontinuous. In general, the zero price level seems to be qualitatively different from all other prices in the human mind. For example, Festinger & Carlsmith (1959) find that a given task tends to be more likeable when the agent performing it is not paid to do it. Gneezy & Rustichini (2000) find that performance is higher in a given task when the level of compensation is zero as opposed to a small amount. However, Seabright (2009) argues that what appears to be discontinuous preferences can be explained by agents signalling altruism, civic duty or some other trait to themselves or others. Following this line of thought, the fore-mentioned experimental results are consistent with continuous preferences. Thus, it is still unclear in what circumstances, if at all, preferences are truly discontinuous at zero price.

An important reason for assuming continuous preferences is that it rules out lexicographic preferences, which is a non-continuous preference relation making the existence of a utility function impossible (Mas-Colell 1995:46). Lexicographic preferences occur when bundles are compared in terms of two or more aspects and one aspect completely dominates another. For example, an agent that always prefers the bundle with more of biological diversity and only begins to consider money when comparing two bundles with the same biological diversity, has lexicographic preferences. This example illustrates that a refusal to trade one good for another regardless of the price indicates lexicographic preferences. Such preferences seem implausible when the commodity is some commonly traded consumer good, but can make more sense when the commodity is something more abstract and crucial.

Hanemann (1991) shows theoretically that the willingness to accept compensation for forgoing a good can approach infinity given that the good has no substitutes. As an extreme example, if an agent is offered money to end his own life, there may very well be no amount of compensation at which the agent wants to forgo his life, indicating lexicographic preferences. Spash (2000) suggests that some people might have lexicographic preferences for public environmental goods for ethical reasons, referring to multiple surveys such as Stevens et al (1991), Spash and Hanley (1995) and Lockwood (1998) where refusals to bid any amount to accept environmental damage might be interpreted as environmental value dominating monetary value.

Thus, theory and empirical evidence suggest that lexicographic preferences may be present in some people in very specific and unusual circumstances. Rosenberger et al (2003) discuss, without reaching a conclusive answer, the possibility that such results indicate little more than the surveys' failure to elicit extreme responses when preferences are extreme. In many cases a plausible explanation for what seems to be lexicographic preferences is that people value some aspects of a bundle much higher than other aspects, but without the former aspect completely dominating the latter. Consider for example, a person being offered $\$10^{99}$ as compensation for allowing the extinction of a particular species of bird. It is reasonable to doubt that the person's hypothetical refusal to accept the compensation would stand if such an extravagant offer would materialize. Indeed, if there is any finite amount that the agent would accept as compensation, preferences are not lexicographic and a utility function might exist. In conclusion, the possibility of lexicographic preferences or otherwise discontinuous preferences does not seem to hinder the existence of a utility function in most economic problems.

2.5. Testing the Utility Maximization Model

This section demonstrates how the utility maximization model, as presented in section 2.1. can be tested and potentially rejected. The model predicts that budget set B and preferences $\succsim$ cause $p(x)=1/|C(\succsim,B)| \quad \forall\, x\in C(\succsim,B)$ and $p(x)=0 \quad \forall\, x\notin C(\succsim,B)$. That is, $\succsim$ and B alone explains $p(x)$. The standard approach in economics is to say that $p(x)$ reveals $\succsim$, but this makes testing the utility maximization model impossible because it entails circular reasoning. Assuming that preferences $\succsim$ are revealed by choice $p(x)$ necessarily leads to the conclusion that preferences and choice are perfectly correlated, which had already been assumed. Going beyond revealed preference is therefore necessary when testing the utility maximization model.

Any experiment with the ambition to test the model must have a control group and at least one treatment group. Suppose for simplicity that only one treatment group is used and that $\succsim_1$, B_1 and

$p_1(x)$ represent the preferences, budget set and outcome of the control group. Then let $\succsim_2$, B_2 and $p_2(x)$ represent the preferences, budget set and outcome of the treatment group. A natural starting point is to consider varying preferences or the budget set so that inferences can be drawn from differences in outcomes.

First consider varying preferences to see what effect it has on choice. Imagine an experiment in which subjects in a treatment group are exposed to some stimulus that a control group is not exposed to, and that this affects preferences so that $\succsim_1 \neq \succsim_2$ and $B_1 = B_2$. Without a specification of preferences, the utility maximization model then offers no predictions on the expected difference between $p_1(x)$ and $p_2(x)$. It is unclear what stimulus could be used and how that stimulus should be expected to affect preferences. Whether the experiment indicates $p_1(x) = p_2(x)$ or $p_1(x) \neq p_2(x)$, the utility maximization model cannot be rejected.

Using this approach, the difference between $\succsim_1$ and $\succsim_2$ must somehow be estimated in order to make any prediction on the difference between $p_1(x)$ and $p_2(x)$, and to make the prediction meaningful, $p_1(x)$ and $p_2(x)$ cannot be used to estimate $\succsim_1$ and $\succsim_2$. Such estimation is difficult because preferences are hidden deeply in the psyche of the agent. Possible measures for estimating preferences include heart rate, blood pressure, skin conductance, pupil dilation, eye movements, muscle activity, dopamine levels, cortisol levels, brain waves (EEG), blood flow in the brain (fMRI) and self-reported perceptions.

Edgeworth (1881) suggested that technological advances would enable the mechanical measurement of preferences using psychophysical indicators, but the idea never gained much popularity. Fisher (1893) argued that preferences are better measured through their choices, and after Samuelson's (1938) ground breaking formalization of revealed preference, the consensus became that choice reveals preference and that any other estimate of preferences is inferior and irrelevant in economics. The psychophysical measurement of preferences is regaining some interest in the 21st century due to advances in brain scanning technology (Camerer et al 2005), but the scientific community is still far from being able to accurately estimate preferences without relying on choice.

Perhaps there is nothing closer to psychophysical measurements of preferences than Shizgal's (1997) model for estimating the utility resulting from electrical brain stimulation in rats. This said, economics is still lacking relevant models for estimating utility independently from choice safely in humans, perhaps because it can easily be argued that any estimate of preferences that is not based on the agent's relevant choice is less valid than the relevant choice itself. Testing the utility maximization model by varying preferences is then bound to result in little more than debates on the extent to which the used measure of preferences is valid, making this approach futile.

An alternative and more promising approach to testing the utility maximization model is to vary budget sets and draw inferences from choices made. In contrast to preferences, budget sets can be accurately estimated and readily controlled by an experimenter. Assuming that altering a budget set does not affect preferences, setting up an experiment where $B_1 \neq B_2$ and $\succsim_1 \equiv \succsim_2$ is easily accomplished.

To determine whether the resulting choices are consistent with the utility maximization model, we use the weak axiom of revealed preference (WARP) which – contrary to what its name seems to suggest – logically follows from the utility maximization model (Mas-Colell 1995:12). Then, by modus tollens, if WARP does not hold, the utility maximization model must be rejected.

WARP says that if x is chosen over y in one budget set, then there can be no other budget set containing x and y in which y is always chosen over x. Formally, if $x, y \in B_1$ and $x \in C(\succsim, B_1)$, then for any B_2 with $x, y \in B_2$ and $y \in C(\succsim, B_2)$, we must have $x \in C(\succsim, B_2)$. An experiment where x is chosen over y when the agent is facing $x, y \in B_1$, but y is chosen over x when $x, y \in B_2$ then is evidence against the utility maximization model. That said, such an inference is somewhat questionable because y might have been chosen over x when $x, y \in C(\succsim, B_2)$.

Battalio et al (1973), Simonson & Tversky (1992), Sippel (1997), Harbaugh & Krause (2000), Harbaugh et al (2001) and Andreoni & Miller (2002) test WARP experimentally with results suggesting that choices often are inconsistent with WARP. In other words, empirical evidence indicates that the utility maximization model is flawed.

A third way of testing the utility maximization model is to hold both preferences and the budget set fixed while varying something that should be independent of both preferences and the budget set. The model can then be rejected if this something affects the dependent variable $p(x)$ in a way that is not predicted by the model. In experimental notation, the null hypothesis in this approach is

$$H_0:\ \succsim_1 \equiv \succsim_2 \text{ and } B_1 = B_2 \ \Rightarrow\ p_1(x) = p_2(x)\,,$$

and the alternative hypothesis is

$$H_1:\ \succsim_1 \equiv \succsim_2 \text{ and } B_1 = B_2 \ \Rightarrow\ p_1(x) \neq p_2(x)\,.$$

Devising an experiment in which preferences and budget sets are fixed, only $p_1(x)$ and $p_2(x)$ need to be estimated. Such simplicity surely plays a big part in making this approach popular in bringing forth flaws of the utility maximization model. Many studies justify the rejection of the utility maximization model on the grounds of this approach, and a few simple and striking phenomena established in those studies will be briefly discussed here.

The framing effect is evident when subjects facing identical problems make different choices, depending on how alternatives are presented in the problem. Examples of the framing effect are ubiquitous and well founded in empirical research. Tversky & Kahneman (1981), McNeil (1982), Cormier O'Connor et al (1985) and Gächter et al (2009) find that people make systematically different choices in logically equivalent decision problems, depending on how the decision problems are worded.

Tversky & Kahneman (1981) asked 152 participants of an experiment to imagine the US preparing for the outbreak of an unusual Asian disease expected to kill 600 people. The participants of the experiment then faced the choice of favouring either program A or program B. In the treatment group, program A was described as *“400 people will die”* and program B was described as *“There is 1/3 probability that nobody will die, and 2/3 probability that 600 people will die”*. In the control group, program A was described as *“200 people will be saved”* and program B was described as *“There is 1/3 probability that 600 people will be saved, and 2/3 probability that no people will be saved”*. 78% of the participants in the treatment group chose program B over program A but only 28% of the participants in the control group chose program B over program A, as seen in illustration 1.

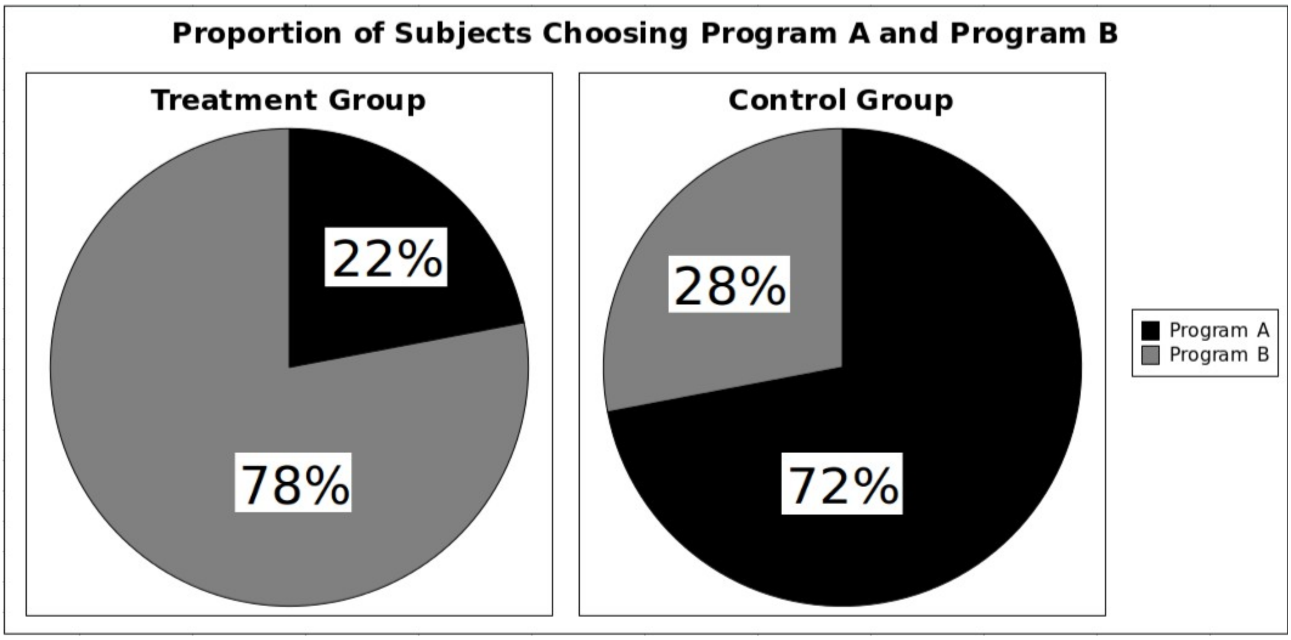

Illustration 1: In an experiment by Tversky & Kahneman (1981) subjects facing a problem expressed in terms of the probability of death favoured a risky program 78% of the time. Subjects facing the same problem but framed in terms of the probability of survival only favoured the same risky program 28% of the time.

The participants favoured the risky alternative significantly more when the problem was framed in terms of the probability of deaths than when it was framed in terms of the probability of saved lives, although program A and program B are logically equivalent between the control group and the treatment group.

Kahneman et al (1991) propose that in general, people choose differently when gambles are framed as potential losses than when they are framed as potential gains because losses are emotionally more severe than equivalent gains, a tendency known as loss aversion. Tom et al (2007) and Canessa et al (2013) find that there is a neural basis for loss aversion in risky decision problems and that losses are treated qualitatively different from gains in the human brain.

Staw (1976), Arkes & Blumer (1985), Garland & Newport (1991), Heath (1995) and Soman (2001) find that sunk costs make a significant difference to what decisions are made. For example, Garland & Newport (1991) find that the probability of undergraduate and graduate business students investing money in a project is positively correlated with how much money has already been invested in that project relative to the budget of that project. Agents taking sunk costs into account suggest that their choice depends not only on preferences and budget, but also on the costs that previously have been incurred. Zeng et al (2013) find evidence for a neural basis for sunk costs having a significant effect on investment decisions.

Tversky & Kahneman (1974), Jacowitz & Kahneman (1995), Strack & Mussweiler (1997), Mussweiler & Englich (2005) and Epley & Gilovich (2006) find that decisions on a numerical scale are significantly affected by agents' exposure to numbers that are obviously irrelevant to the problem at hand. For example, if one is told that the ancient city-state Babylon had 50 000 inhabitants and then asked to estimate how much a car is worth, one's estimate is likely to be lower than it would have been, had one been told that Babylon had 200 000 inhabitants, even though it is obvious that the population of Babylon has no relation to the value of a car. This effect is commonly known as the anchoring effect. Neither cognitive ability, expertise nor monetary incentives removes the anchoring effect, although it may reduce it (Wright and Anderson 1989, Englich et al 2005, Simmons et al 2010, Bergman et al 2010).

The above outline of empirical research is by no means a complete account of all evidence, but it clearly indicates that people often fail to maximize utility in predictable ways. Seeing that agents by definition *want* to maximize utility, the evidence suggests that agents at least sometimes fail to maximize utility because they are *unable* to maximize it.

3. Alternatives to Utility Maximization

3.1. Judgement Error or Measurement Error

In the last section, agents were deemed unable to maximize utility, which can be explained by agents making mistakes. This explanation of apparent sub-optimal behaviour will from here on be referred to as the *judgement error explanation* because it says that agents are flawed as decision makers. An alternative explanation is that what appears to be behaviour inconsistent with utility maximization actually is due to measurement errors. In the same way that a physicist is bound to make small errors when measuring the weight, length or volume of an object, economists plausibly make small errors when measuring or estimating preferences and choices.

Assuming this alternative explanation, researchers are unable to see that utility is successfully maximized because they fail to sufficiently take all aspects of preferences into account. It is possible that an agent genuinely strictly prefers x to y despite these bundles' so called logical equivalence, because x is described in more pleasant terms than y, or because resources previously have been spent trying to obtain x. Two seemingly identical bundles presented differently may not be identical after all. This latter explanation will from here on be referred to as the *measurement error explanation* because it says that it is the researcher's measurement of preferences rather than the agent's utility calculation that is flawed.

On its own, the measurement error explanation is effectively a reaffirmation of the assumption that x being chosen over y implies that $x \succsim y$. If the measurement error explanation fully explains all observed anomalies, the utility maximization model may hold even when it appears that agents fail to maximize utility. Most economic models make this assumption explicitly or implicitly, and characterizes empirical observations of what appears to be deviations from utility maximization as measurement errors. Indeed, “preferred to” is often taken to be synonymous to “chosen over” in economics.

The judgement error explanation is compelling because it intuitively appears more realistic than the measurement error explanation on grounds of introspection and anecdotal evidence. While most economists would consider it likely that people are at least slightly flawed decision makers, most non-economists would consider the same proposition obvious. Anyone who suspects that he or she ever made a mistake has reason to hold the judgement error explanation true, however minor that mistake might have been.

Formal empirical evidence also supports the judgement error explanation. Sokol-Hessner et al (2012) find that loss aversion can be reduced by using emotion regulation strategies such as reframing decisions and outcomes in one's mind. That is, it seems that agents facing decision problems constructed to misguide them (e.g. the decision problems posed by Tversky & Kahneman, 1981) can make choices more consistent with utility maximization if they are provided with the appropriate mental tools.

Soman & Cheema (2001) find that the sunk cost effect diminishes with repetition. Wilson et al (1996) and Bergman et al (2010) find that the anchoring effect is moderated by expertise and intelligence. To be sure, this is not to say that the anchoring effect completely disappears for experienced and intelligent people. Indeed, Englisch et al (2006) find that highly educated and experienced legal professionals in Germany are susceptible to anchoring in giving sentences to criminals. Nevertheless, the lessening of susceptibility to these anomalies suggests that agents become better at maximizing utility whenever they can, supporting the judgement error explanation.

These findings are not as easily reconciled with the measurement error explanation because there is little reason to believe that providing agents with intelligence or training in any way affects the reliability of measurements. However, it is possible that intelligence and training affect agents' preferences through some more or less obscure mechanism. It is also possible that some omitted variable simultaneously affects both preferences and other variables such as intelligence. For example, the level of dopamine in the human brain is intricately related not only to intrinsic rewards and motivation (Schultz 1998, Pessiglione 2006), but also to intelligence (Previc 1999, Nieoullon 2002). Thus, dopamine might affect both intelligence and preferences simultaneously in such a way that intelligence predicts preferences and choices, without there ever being any error in judgement. For a detailed account of how dopamine might relate to utility and decisions in economics, see Caplin & Dean (2008).

Similarly, if training somehow is related to preferences, it may be that agents are consistently and perfectly maximizing utility even when it appears as if though they are improving their ability to maximize utility through training. For example, the deliberate reframing of a decision in an agent's mind, as in Sokol-Hessner et al (2012), might have changed the agent's preferences rather than his ability to make better choices in accordance to his fixed preferences. Repeated exposure to several decisions similar to each other may also change the way in which one perceives the desirability of bundles in a budget set. For example, an experienced poker player may make decisions differently from a beginner poker player, not only because the experienced poker player makes better decisions but also because the experienced poker player's attitude towards risks and potential outcomes might have changed as a result of the training.

Even evidence of violations of WARP can potentially be explained within the utility maximization model. In some experiments choices are made at long intervals, meaning that preferences might change in subjects from one time to another, which would nullify the test. For example, Battalio et al (1973) alter the budget sets facing subjects of their experiment every week. It is then possible that a subject's preferences change from one week to another in such a way that choices appear inconsistent under the assumption that preferences remain constant, even if utility is consistently maximized.

Sippel (1997) addresses this methodological weakness by devising a lottery where subjects have all the time they need to make ten consumption choices, each with a different budget set and each having a 1/10 probability of being realized. Observations where $x \in C(\succsim, B_1)$ with $x, y \in B_1$ and $y \in C(\succsim, B_2)$ with $x, y \in B_2$ are then interpreted as violations of WARP. However, if $x, y \in C(\succsim, B_2)$, this outcome is consistent with WARP. Also, Hjertstrand & Swofford (2013) propose a statistical test in this line of thought that when applied to previously published data in effect fails to reject the utility maximization model.

The possibility that agents are perfect utility maximizers cannot be ruled out, but introspection, anecdotal evidence, and the formal empirical evidence presented in this text indicate that there is at least some truth to the judgement error explanation. While scholars differ on the significance and importance of the judgement error explanation and the measurement error explanation respectively, the validity of both approaches in economics should be uncontroversial. Agents seem to be at least somewhat prone to error, and even in the most exact science any measurement is prone to error.

Any set of choices that appears to contradict completeness, transitivity or continuity might be due to an agent making sub-optimal choices or to a researcher that infers erroneous conclusions from observed behaviour. Thus, either of the two explanations is sufficient to explain all observations that can be interpreted as evidence for preferences not being complete, transitive or continuous. The most plausible story of how the anomalies described in this text arise, is reasonably a combination of the judgement error explanation and the measurement error explanation.

Contrary to the impression given by some behavioural economists' criticism of classical economics, most economists seem to agree that it is reasonable to believe that economic agents are flawed utility maximizers, and that the perfect utility maximizer is an ideal that can be more or less useful for constructing parsimonious models.

For example, Simon (1955:104) is very critical of the classical utility maximization model when he says that *“there is a complete lack of evidence that, in actual human choice situations of any complexity, [utility maximization] computations can be, or are in fact, performed”*. To be sure, he then admits the possibility that the unconscious might be better at maximizing utility than the conscious. On the opposite side of the spectrum, Friedman (1953:15) assures that the traditional assumptions are not necessarily realistic but rather useful for their ability to predict outcomes. In essence, his argument is that *'the relevant question to ask about the “assumptions” of a theory is not whether they are descriptively “realistic,” for they never are, but whether they are sufficiently good approximations for the purpose in hand'*.

A modern example of a controversy on the focus of economic models is based on Kahneman et al (2003) who outline neurological findings suggesting that agents systematically make sub-optimal choices, emphasizing the usefulness of what is here called the judgement error explanation of economic anomalies. Gul & Pesendorfer (2005) respond to such interest in neurology by asserting that the underlying reasons for an agent's choice are irrelevant because choices revealing preferences is a methodological stance rather than an assumption in mainstream economics.

Economists assume perfect utility maximizing agents because it is a reasonable starting point for explaining and predicting economic behaviour, and because it allows for the construction of parsimonious models. A refinement that implies forgoing this assumption is typically a sacrifice of parsimony in an effort to increase realism or predictive power. Much of the controversy is in deciding what weight should be given to realism as opposed to parsimony in an economic model, and this depends on the problem at hand. Section 3.3. proposes a probabilistic framework that allows researchers to balance parsimony against realism by incorporating any amount of cognitive biases into their model, and that includes the utility maximization model as a special case.

3.2. Existing Models of Sub-Optimal Behaviour

Several attempts have been made to model sub-optimal behaviour, some more successful than others. One early example is Simon (1955) who argues that the utility maximization model requires too much from the agent for it to be accurate. He points out that any organism is unlikely to have perfect information or the cognitive ability required to perfectly evaluate the utility of all feasible bundles. He especially focuses on the difficulty of an agent to consider and evaluate all elements in a budget set containing many bundles. Instead, he proposes that the agent uses heuristics or simplifications to choose the best bundle that she can only from the subset of bundles that she considers. The chosen bundle is then not necessarily the most preferred bundle but rather a sufficiently good bundle from the agent's perspective.

Tversky (1972) propose and find some empirical support for a model where the agent sequentially compares bundles in terms of one aspect at a time and eliminates bundles according to some criteria for every iteration until only one bundle remains. Kahneman & Tversky (1979) develop prospect theory, in which preferences and choices depend on potential losses and gains compared to a reference point, rather than the final outcome alone as in the utility maximization model. Thanks to its predictive powers and generality, prospect theory is probably the most influential model in microeconomics that is inconsistent with standard theory.

Smith & Walker (1993) argue that an agent must weigh the benefits of making good choices against the cost of making good choices. The time and effort of deliberating is an opportunity cost to the agent, because that time could be spent doing something more lucrative or enjoyable. Their review of experimental studies reveals that agents deviate more from optimal choices when pay-offs increase or opportunity costs decrease, suggesting that the trade-off described by Smith & Walker contributes to deviations from utility maximization.

Walker & Ben-Akiva (2002) propose a general model where utility depends not only on preferences and a budget constraint but also a random parameter with unknown distribution, making utility random. The random nature of utility can then explain how odd behaviour is consistent with utility maximization. Tsang (2008) on the other hand sees deviations from utility maximization as being caused by limited computational power, so rather than having a random utility function, the agent fails to maximize deterministically determined utility due to shortcomings in cognitive ability.

Other examples of models intended to account for behaviour that seems to violate the utility maximization model are provided by Marley (1968,1982), McFadden (1980), Machina (1985) and Gigerenzer & Goldstein (1996). Related discussions of such models are offered by Loewenstein & Lerner (2003), Rieskamp et al (2006) and Weirich (2008). A more in depth summary and comparison of the merits and weaknesses of these models, albeit interesting, is a project too ambitious to be covered in this text.

The studies mentioned here is by no means an exhaustive list and considering how many models have been suggested during the last six decades, it is likely that some important models unintentionally but still inappropriately have been omitted in this text. In any case, there is clearly a justified interest in modelling violations of the utility maximization model. A consensus on the best model for utility maximization violations is lacking and none of the models listed here stands out as accurate and general enough to form a foundation for a theory that takes flawed decision making into consideration.

Thaler (1980) points out that prospect theory is particularly popular, much thanks to it being applicable to many different kinds of decision problems. List (2004) finds that when prospect theory is applicable, experience of agents correlates positively to the predictive power of standard theory and negatively to prospect theory, indicating that agents can learn to reduce the violations of utility maximization that prospect theory represents. Prospect theory fails however, to explain and predict outcomes in other contexts. For example, it offers no explanation for an agent's choice being affected by an irrelevant numerical value due to an anchoring effect (e.g. Tversky & Kahneman, 1974).

Smith & Walker (1993) clarifies that empirical evidence suggests that their model of effortful deliberation incurring an opportunity cost only explains a part of deviations from the utility maximizing choice. Ariely et al (2009) even finds that performance in cognitive tasks declines when incentives are excessive. It is then clear that the Smith & Walker model by itself is insufficient to explain deviations from utility maximizing behaviour.

The generalized random utility model suggested by Walker & Ben-Akiva (2002) fails to specify the causes for deviations from utility maximization and therefore lacks predictive power, but it enables the combining of models into a common framework. Machina (1985:576) criticizes random utility models for being unclear on how a random element is realised and to what extent one realised error might depend on an earlier realised error. Even if a clarification is provided, the generalized random utility model is unintuitive because it fails to recognize flaws in decision making. It maintains the view that all choices are optimal, even in the face of phenomena such as the framing effect, the anchoring effect and the reduction in the susceptibility to error that comes with experience. Nevertheless, a generalized random utility model is a promising approach because it has the potential to explain all choices of an economic agent and to combine several models into a single fairly simple framework.

Fudenberg (2006) argues that more general models is exactly what is needed for insights in behavioural economics to be widely accepted and integrated into standard economics. He suggests that the specific assumptions of the many models in behavioural economics should be reduced to a few core assumptions that form a more general foundation, and names prospect theory among others as an example to be emulated. McFadden (1999) acknowledge not only the flaws but also the merits of the utility maximization model and argues that a hybrid between standard theory and behavioural economics is an appropriate approach. The framework presented in the next section addresses the need for a theoretical foundation compatible with both behavioural economics and the standard theory of utility maximization.

3.3. Introducing A General Framework

As in the utility maximization model, assume that the binary preference relation $\succsim$ defined on the consumption set X is complete, reflexive, transitive and continuous. The arguments for and against these assumptions are the same as those outlined in sections 2.2 - 2.4. Preferences are then guaranteed to be represented by a utility function $u: X \rightarrow \mathbb{R}$ such that $x \succsim y \Leftrightarrow u(x) \geq u(y)$ $\forall x, y \in X$, and the optimal choice is given by $C(\succsim, B)$ = $\{x \in B: x \succsim y \quad \forall y \in B\}$ = $\{x \in B: u(x) \geq u(y) \quad \forall y \in B\}$. However, rather than assuming that the agent always makes an optimal choice as in the utility maximization model, allow actual choice to deviate from optimal choice due to some flaw in the agent's decision making process.

Let an additional argument $\boldsymbol{\gamma}$ represent all the reasons that actual choice deviates from optimal choice, which can be called quirks, anomalies, cognitive biases or errors. If the agent's choice is affected by exactly J quirks, $\boldsymbol{\gamma}$ can be defined as a vector $\boldsymbol{\gamma} = (\gamma_1, ..., \gamma_J)$ where γ_j is the effect of quirk j on the utility that results from the decision in question, expressed on the real number scale. An actual choice set $C(\succsim, B, \boldsymbol{\gamma})$ then depends on preferences, budget constraints and as of yet unspecified quirks that may cause an agent's behaviour to deviate from $C(\succsim, B)$.

Seeing that the utility function is ordinal, $\boldsymbol{\gamma}$ must be scaled to correspond to any one of the utility functions that represent preferences. This choice of utility function is arbitrary and implies no loss of generality. Regardless of which utility function is used, the agent's actual choice can be seen as the set of chosen bundles after having accounted for the effect of J quirks so that for any $x \in C(\succsim, B, \boldsymbol{\gamma})$ and $y \in C(\succsim, B)$, we have $u(x) \;=\; u(y) + \Sigma_{j=1}^{J} \gamma_j$. Equivalently, the set of chosen bundles after considering mistakes can be defined as $C(\succsim, B, \boldsymbol{\gamma}) \;=\; \{x \in B : u(x) = u(y) + \Sigma_{j=1}^{J} \gamma_j \quad \forall\, y \in C(\succsim, B)\}$. Note that $\gamma_j \leq 0 \quad \forall\, \gamma_j \in \boldsymbol{\gamma}$ holds by definition and that this implies $\Sigma_{j=1}^{J} \gamma_j \leq 0$ and $u(x) \leq u(y)$. Perfect utility maximization then is the special case where $\Sigma_{j=1}^{J} \gamma_j = 0$.

This definition implicitly states that situations in which an agent's quirks causes her to choose bundles that are infeasible never occur. Imposing a restriction saying that $\boldsymbol{\gamma}$ is defined such that $C(\succsim, B, \boldsymbol{\gamma}) \subseteq B$ is then necessary to avoid situations in which $C(\succsim, B, \boldsymbol{\gamma}) = \emptyset$. If this restriction was not imposed, a disturbing consequence would have to be addressed. By definition, only bundles in B can be chosen, so allowing agents to mistakenly choose bundles that are not in B would be absurd. Therefore, when quirks cause an agent to choose or try to choose a bundle that is not in B , an empty choice set would mathematically be the result, but applied in the real world this would be little more than admitting a failure to understand the causes of choice.

The distinction between the sets $C(\succsim, B)$ and $C(\succsim, B, \boldsymbol{\gamma})$ is so far little more than a formal acknowledgement of the cognitive biases for which evidence is mounting in behavioural economics. It allows researchers to begin with the utility maximization model as a reasonable first approximation and then adjust for relevant quirks depending on the problem at hand. Within this framework, realism typically increases while parsimony decreases, as the number of quirks taken into account increases. It is a synthesis of classical microeconomics and behavioural economics in the sense that models from both schools can be incorporated into the framework.

An increasingly large body of empirical evidence and models accounting for many different cognitive biases suggests that it is reasonable, and perhaps even necessary, to see $\boldsymbol{\gamma}$ as an expression for many intricately interrelated quirks. The specification and organisation of quirks and their interactions is an important and huge research project that is beyond the scope of this text. The emphasis here is rather on the possibility and usefulness of specifying quirks and connections between quirks within a clear general framework.

If all the components of $\boldsymbol{\gamma}$ fully explains choice, there is nothing random about it. However, when formulating models in practice it is more feasible for researchers to pick and define the components of $\boldsymbol{\gamma}$ as they deem appropriate. It is likely that no model accounts for all of an agent's quirks and if relevant quirks are omitted, choice should be seen as a random outcome. The deterministic nature of the utility maximization choice rule $C(\cdot)$ then does not easily translate to a the realistic framework that this text aims to describe. $C(\cdot)$ produces a subset of B, but something more is required to represent the stochastic nature of actual decision making.

One way of formally expressing a probabilistic choice mechanism is to define $\bar{C}(\succsim, B, \boldsymbol{\gamma})$ as a vector with every component $p(x_i)$ representing the probability of the agent choosing bundle x_i, where the index i runs from 1 to I in a budget set consisting of $|B| = I$ mutually exclusive bundles. The framework can then be summarized as

$$\bar{C}(\succsim, B, \boldsymbol{\gamma}) = [p(x_1), p(x_2), \ldots, p(x_I)] .$$

This expression says that the probability of an agent choosing a bundle from the set of all feasible bundles is somewhere between zero and one, depending on preferences, cognitive biases and chance.

Actual choice deviates systematically from optimal choice if $\Sigma_{j=1}^{J} \gamma_j < 0$ and it deviates randomly from optimal choice if relevant quirks are omitted from $\boldsymbol{\gamma}$. Thus, unless there is a way to make sure that all relevant quirks are included in $\boldsymbol{\gamma}$ and that $\Sigma_{j=1}^{J} \gamma_j = 0$, actual choice is guaranteed to deviate from optimal choice. The utility maximization model is the special case in which the agent is not affected by systematic quirks or random whims. The probability of a bundle being chosen would then be strictly positive if and only if that bundle maximizes utility.

This framework is so general that it is consistent with almost any decision model. Exceptions include models relaxing completeness, transitivity or continuity of preferences. For example, Ok (2002) propose a vector valued utility function to account for incomplete preferences, which is inconsistent with the real valued utility function that the suggested framework here implies. Random utility models such as the one by Walker & Ben-Akiva (2002) are inconsistent with a framework in which a stable utility function representing preferences is guaranteed by completeness, reflexivity, transitivity and continuity.

Furthermore, seeing that $\boldsymbol{\gamma}$ can explain any decision or combination of decisions, the framework is consistent with any empirical evidence. What appears to be a violation of completeness, transitivity or continuity can instead be explained by flaws in the agent's decision making process. For example, when an agent chooses x over y, y over z and z over x a researcher can conclude that the agent makes at least one decision that fails to maximize utility, due to flawed heuristics represented by $\boldsymbol{\gamma}$, or random whims. The framework being consistent with any empirical evidence is a weakness more than a strength because it implies the impossibility of testing the truthfulness of it. In any case, the framework has little value in isolation but is rather intended as a stepping stone towards improved models, or as an extension of existing models.

The suggested approach is applicable to individual agents whose ultimate goal it is to maximize utility. It may or may not be possible to generalize the framework further. The preferences of a group of people might not be complete or transitive even if all individuals in the group have complete and transitive preferences, so a utility function does not necessarily exist for groups of people. The difficulty of aggregating preferences is worth considering not only for consumers but also for firms.

In many situations of uncertainty it is reasonable to view the agent as maximizing *expected* utility rather than utility, but the framework is not readily applicable to the maximization of expected utility without first addressing some important issues. With expected utility, every bundle is a lottery with a probability assigned to every outcome in that lottery. The expected utility of a lottery is a weighted average of the utility of all possible outcomes in that lottery.

Neumann & Morgenstern (1944) propose that expected utility is weighted according to the objective probabilities of outcomes, which is now the predominating view in economics. Savage (1954) instead proposes that expected utility is weighted according to the agent's perception of the probabilities of outcomes, which is an approach that is now known as subjective expected utility. Before applying the suggested framework to expected utility, we must clarify whether the agent maximizes objective expected utility, subjective expected utility or some other weighted average.

Furthermore, expected utility theory typically assumes independence in addition to completeness, reflexivity, transitivity and continuity, i.e. that for any $x, y, z \in X$ with $x \succsim y$, we have $\alpha x + (1-\alpha) z \succsim \alpha y + (1-\alpha) z \quad \forall \alpha \in (0,1]$. The independence assumption is particularly controversial because of ubiquitous empirical evidence contradicting it. The interested reader can find a more in depth review of expected utility theory and related empirical research in Schoemaker (1982). Machina (1982) suggests that expected utility theory is possible with assumptions less restrictive than independence. In any case, guaranteeing an expected utility function requires more than what is required to guarantee a utility function.

Firms are typically seen as maximizing profit in standard microeconomics, but such a view is unlikely to fully reflect the complex reality. In many cases, firms intentionally and explicitly consider goals other than monetary gains, such as promoting a clean environment or economic equality. Introducing uncertainty, it is not clear whether firms should maximize expected profit or the expected utility that the beneficiaries of the firm gain from it. Lin et al (1974) finds that the maximization of expected utility predicts agricultural firms' decisions better than profit maximization.

Even if a firm stipulates profit maximization as its sole ambition, the agents that make up the firm may have conflicting goals, implying some ambiguity on what the firm is actually optimizing. For example, an employee typically has an incentive to increase her salary at the expense of her employer's profit, so decisions made by that employee on behalf of the employer may well act to reduce the employer's profit. Berle & Gardiner (1932) argue that profit maximization becomes less plausible when directors on salaries exercise more control over a firm than its many small and less influential share owners, which is common for many large modern corporations.

The suggested framework is relevant only for problems where an agent optimizes a function and mistakes representable by $\boldsymbol{\gamma}$ can be quantified on the same dimension. If a function exists and an agent's or institution's only objective is to maximize it, a framework analogous to the suggested one can easily be imagined. The problem is that it is often not clear what decision makers optimize, if anything. The idea that actual choice deviates from optimal choice due to decision makers' imperfections is tenable for many decision problems other than utility maximization, but if there is no function defined on the real number scale to optimize, the suggested approach fails to describe economic decisions.

3.4. Stochastic Supply-Demand Schedule

The application of the framework presented in section 3.3. on the supply-demand schedule shows that it leads to tangible and useful conclusions. This section demonstrates how the supply-demand schedule can be enriched by seeing it in the light of the suggested framework. The demand function is one of the most important applications of the utility maximization model. It answers how much of a good an agent chooses to consume, given prices and the agent's wealth. While the consumption set X can be any set of consumable bundles in general, the demand function is applicable only for a good that is consumed in quantities defined on the real number scale. The demand function also requires that each quantity of the good in question has a real-valued price assigned to it and that the agent has a real-valued wealth to spend on goods in X .

The conventional way of constructing a demand function begins with assuming $X=\mathbb{R}_+^L$ so that every $x\in X$ is interpreted as a bundle of L goods, with every good $l\in L$ being consumed in an infinitely divisible quantity $x_l \geq 0$. The next step is to assume that every good $l\in L$ has a price $p_l > 0$ per unit so that nothing is entirely free in infinite amounts and the cost of consuming x_l is $p_l x_l$. Now let $\boldsymbol{x}=(x_1, ..., x_L)$ be a quantity vector for all goods in X and let $\boldsymbol{p}=(p_1, ..., p_L)$ be the price vector for those goods. The total cost of any bundle $\boldsymbol{x}\in X$ then is the dot product $\boldsymbol{p}\cdot\boldsymbol{x} = \Sigma_{l=1}^{L} p_l x_l$. The agent cannot spend any more than a finite wealth $w\geq 0$, so his budget set is $B=\{\boldsymbol{x}\in X: 0\leq \boldsymbol{p}\cdot\boldsymbol{x}\leq w\}$, which is often referred to as a Walrasian budget set.

Assuming strictly convex preferences in addition to the outlined market structure guarantees the existence of a demand function giving a unique utility maximizing bundle for every price vector $\boldsymbol{p}$ and wealth w (find proof in appendix A). Strictly convex preferences means that the set $\{\boldsymbol{x}\in X: \boldsymbol{x}\succsim \boldsymbol{x}'\}$ is strictly convex for any $\boldsymbol{x}'\in X$, which implies diminishing marginal utility in every good $l\in L$.

A Walrasian budget set and strictly convex preferences are reasonable assumptions in many typical decision problems in economics but are by no means always true. If either the budget set or preferences are not convex, a demand function cannot be guaranteed. A volume discount for a good l with a marginal price that is lower for high levels of x_l than for low levels of x_l for example, would give a non-convex budget set. Strictly convex preferences are violated if for example either a glass of milk or a glass of orange juice is preferred to a mixture of the two, and this seems plausible. This caveat is as valid in the standard utility maximization model as it is in a framework acknowledging human errors in decision making.

If an agent attempts but fails to maximize utility, the quantity actually demanded will differ from the quantity that is optimally demanded. Using the notation from section 3.3, if $\Sigma_{j=1}^{J} \gamma_j < 0$ or if factors affecting choice are excluded from $\boldsymbol{\gamma}$, then actual demand generally differs from optimal demand. However, this notation is unnecessarily complex here because, by the definition of demand, deviations from an optimal choice can be expressed on a real number scale representing the consumed quantity of a good.

A convenient way of expressing demand and errors in demand is to let an asterisk denote optimality so that if an agent optimally demands $x_l{*}(\boldsymbol{p},w)$ units of a given good, the agent ultimately chooses to demand $x_l{*}(\boldsymbol{p},w)+\varepsilon_l$ units of the good, where ε_l is a random variable with distribution $\varepsilon_l \sim D_l(\mu_l,\sigma_l)$, still subject to $\boldsymbol{p}\cdot\boldsymbol{x} \in [0,w]$. The random variable ε_l expresses the same deviations from optimality as $\bar{C}(\succsim, B, \boldsymbol{\gamma})$, but in a more specific context.

If the mean $\mu_l=0$ the agent's average choice is utility maximizing, and if $\mu_l \neq 0$ there is a systematic bias in the agent's decision making. If $\mu_l = 0$ and the standard deviation $\sigma_l=0$ the agent is a perfect utility maximizer, but if either $\mu \neq 0$ or $\sigma_l > 0$ the agent's decisions are generally flawed. This description differs from the standard demand function only in that a random element representing the agent's mistakes is formally included. This may initially seem like an innocent stance, but turns out to have important implications.

Having introduced a random element to the individual's demand function, the next natural step is to consider a market demand function representing the aggregated demand of several individuals. Suppose that each consumer demands $x_{lm}{*}(\boldsymbol{p},w_m)+\varepsilon_{lm}$ where the index m represents the m'th consumer in a set of M consumers. Preferences, wealth and error terms may vary but all consumers face the same prices $\boldsymbol{p}$. There is now one error term $\varepsilon_{lm} \sim D_{lm}(\mu_{lm}, \sigma_{lm})$ for each good and consumer so that the error distribution of one individual may differ from the error distribution of another individual. It is also possible that the error distribution of one consumer depends on the error distribution of other consumers.

The market demand function for a group of utility maximizers is $d_l{*}(\boldsymbol{p},(w_1,\dots,w_M))$ = $\Sigma_{m=1}^{M} x_{lm}{*}(\boldsymbol{p}, w_m)$ and the stochastic version of market demand subsequently is d_l = $\Sigma_{m=1}^{M} x_{lm}{*}(\boldsymbol{p}, w_m)$ + ε_{lm} . If ε_{lm} is an i.i.d. variable, d_l is approximately normally distributed for a group of many people by the central limit theorem, and the standard deviation of the mean $\sigma_l / \sqrt{M}$ decreases as M increases, meaning that d_l approaches $\Sigma_{m=1}^{M} x_{lm}{*}(\boldsymbol{p}, w_m) + \mu_m$ as M approaches infinity.

Inconveniently, empirical evidence suggests that errors are not identical or independent across agents (Moulton 1986). Rather, individuals tend to emulate others' behaviour even when it makes no sense, leading to conformity in markets and otherwise (Asch 1951, Bikhchandani & Sharma 2000, Hong et al 2004). Thus, ε_{lm} being an i.i.d. variable is a dubious assumption, and the error distribution of the market demand function can take many forms other than a normal distribution. Strictly speaking, a normally distributed error term is impossible because it spreads out from negative infinity to positive infinity while a decision only can be made in the finite interval $[0, w]$. A normal error distribution can be a reasonable approximation when the mean of the error is at least a few standard deviations from the limits of the budget set, but an optimal decision typically entails exhausting the budget.

If for any bundle $\boldsymbol{x} \in X$, there is always a bundle $\boldsymbol{y} \in X$ close to $\boldsymbol{x}$ such that $\boldsymbol{y} \succ \boldsymbol{x}$, preferences are said to be locally non-satiated. Assuming locally non-satiated preferences is enough to imply that all wealth is optimally spent (Mas-Colell 1995:52). Remember from section 3.3. that $\boldsymbol{\gamma}$ and $\bar{C}(\succsim, B, \boldsymbol{\gamma})$ are defined so that only bundles in an agent's budget set can be chosen. If the utility maximizing bundle is on the upper limit of a consumer's budget set, any mistake that entails spending more than the given wealth cannot be realised.

Equivalently, if the consumer optimally spends all wealth, she can only make mistakes such that $\Sigma_{l=1}^{L} p_l \, \varepsilon_l \leq 0$. If wealth is allocated between several different goods, an overconsumption of one good may be financed by an underconsumption of another good, so $\sigma_{lm} > 0$ does not necessarily imply $\mu_{lm} < 0$. Nevertheless, deviations from a utility maximizing bundle on the upper bound of a budget set can only lead to an inefficiently low total consumption.

Next, a supply function is introduced. In spite of the reservations expressed in section 3.3, assume that firms are sub-optimal profit maximizers. Like consumers, firms can be thought of as having quirks representable by a vector like $\boldsymbol{\gamma}$. Unless all of the quirks are perfectly specified, firms also appear to make random choices on the quantity they should supply to maximize profit, given prices. Firms may well face restrictions on how much they produce, similar to how consumers are constrained by a budget set, which then limit the quantity they can supply. However, it is conventional to assume that firms can produce any amount they see fit to maximize profit.

A single firm's supply function is $q_l{*}(p_l) = \operatorname{argmax}_{x_l} p_l \cdot x_l - c(x_l)$ where $c(x_l)$ is the firm's cost of selling x_l units of good l. Using the same reasoning as with the demand function, consider a market with several sub-optimal firms. Market supply then is $s_l = \Sigma_{n=1}^{N} \operatorname{argmax}_{x_{ln}} p_l \cdot x_{ln} - c(x_{ln}) + \varepsilon_{ln}$ where n refers to the $n'\text{th}$ firm in a set of N firms and $\varepsilon_{ln} \sim D_{ln}(\mu_{ln}, \sigma_{ln})$.

Consumers reasonably prefer more to less of a good when demand exceeds supply, given constant prices and wealth. Similarly, firms want to increase trade of the same good when supply exceeds demand. That is, consumers buy as much of good l as they can when $s_l < d_l$, and firms sell as much of good l as they can when $d_l < s_l$. Thus, for any market size and any degree of proneness to error, the quantity traded given prices $\boldsymbol{p}$ is $\min\{s_l, d_l\}$, or $\min\{s_l{*}(p_l) + \varepsilon_{ls}$, $d_l{*}(\boldsymbol{p}, (w_1, \ldots, w_M)) + \varepsilon_{ld}\}$, where the notation $\varepsilon_{ls} = \Sigma_{n=1}^{N} \varepsilon_{ln} \sim D_{ls}(\mu_{ls}, \sigma_{ls})$ and $\varepsilon_{ld} = \Sigma_{m=1}^{M} \varepsilon_{lm} \sim D_{ld}(\mu_{ld}, \sigma_{ld})$ is used to distinguish between market supply errors and market demand errors.

The exact distribution of the market clearing quantity $\min\{s_l, d_l\}$ can be derived directly from the distributions of s_l and d_l. The expected value of $\min\{s_l, d_l\}$ is invariably lower than any weighted average of s_l and d_l, so the minimum of supply and demand being traded acts as an inhibitor of trade at all levels of prices and wealth. This implies that the average quantity produced and consumed on a free market is lower than what is socially optimal even when all market agents individually make optimal decisions on average.

Consider especially a market where there are no transaction costs and prices are given such that an equilibrium quantity $s_l*(p_l) = d_l*(\boldsymbol{p}, w_1, \ldots, w_I)$ results. Suppose that both firms and consumers make sub-optimal decisions but maximize profits and utility on average, so that $\mu_{ls}, \mu_{ld} = 0$ and $\sigma_{ls}, \sigma_{ld} > 0$. We can then expect the actual quantity produced and consumed to be lower than the equilibrium quantity predicted by standard microeconomic theory. If errors are i.i.d, this effect becomes smaller as the size of a market increases.

3.5. Simulations of Stochastic Supply and Demand

A demand curve as it is conventionally drawn on a two-dimensional plane is not satisfactory as an illustration of demand that accounts for the random imperfections of consumers. First, considering any given price, the probability of an agent demanding a range of quantities can be illustrated using a probability density function. Putting a continuous interval of prices together will then form a surface in a three dimensional space illustrating demand with errors. If errors are normally distributed, the probability distribution function will be bell shaped (illustration 2) and the surface consisting of a continuum of probability density functions for all prices will be a ridge (illustration 3).

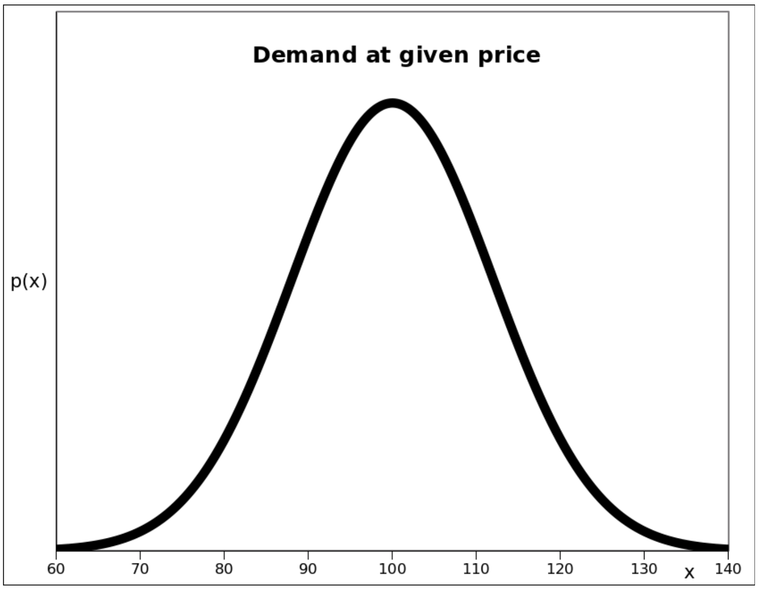

Illustration 2: Probability density function of demand with $d_l* = 100$ *and* $\Sigma_{m=1}^{M} \varepsilon_{lm} \sim N(0,12)$

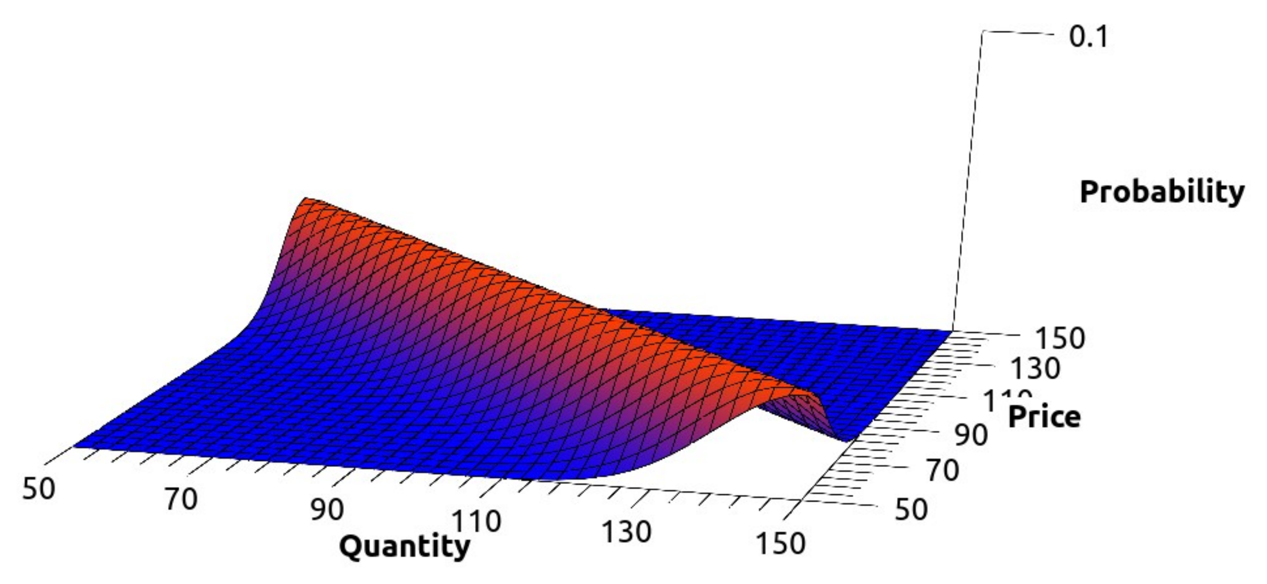

Illustration 3: The demand ridge gives the probability that a quantity will be demanded for a given price. Any price can be seen as giving a probability density function, or visually a slice of the ridge, which by definition has an area equal to one.

Similarly, if firms are allowed to deviate from their profit maximizing choice, a supply curve on a two-dimensional plane is not a satisfactory illustration of the random nature of firms' supply. A probability density function can illustrate the probability that a range of quantities is supplied, given a price (illustration 4). Displaying all prices together in a three dimensional graph gives a surface illustrating stochastic supply for a continuum of prices (illustration 5).

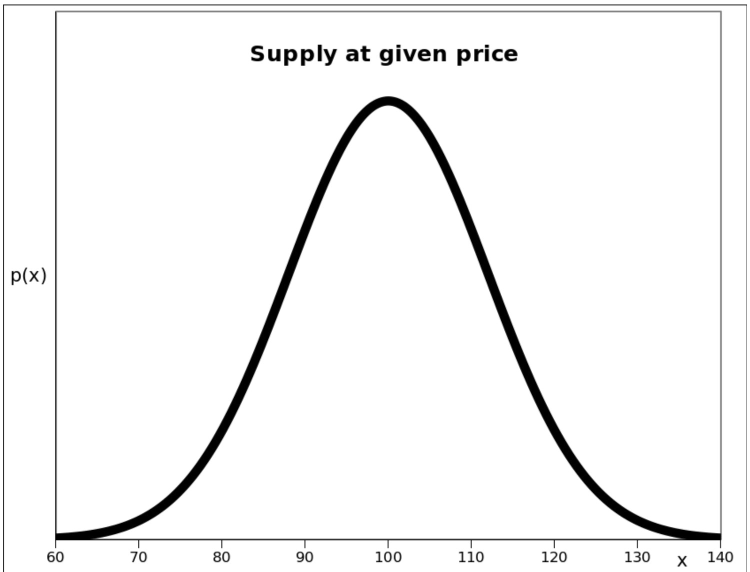

Illustration 4: Probability density function of supply with $s_l{*} = 100$, $\Sigma_{n=1}^{N} \varepsilon_{\mathrm{ln}} \sim N(0{,}12)$

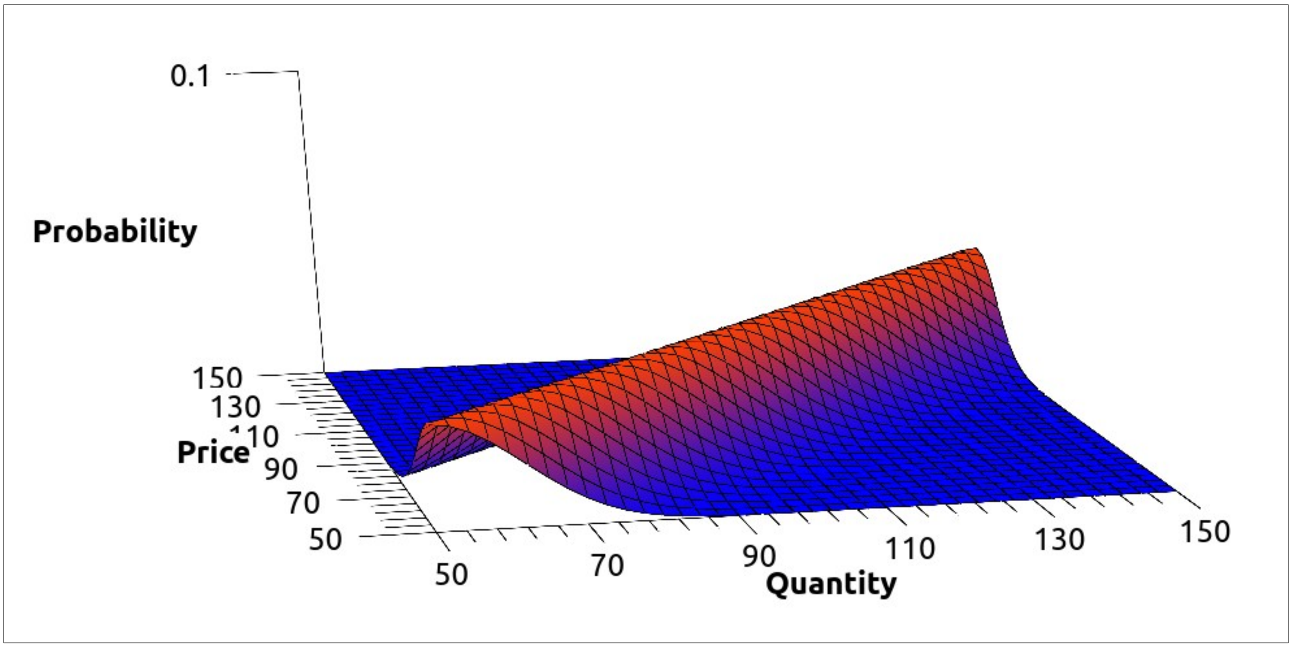

Illustration 5: The supply ridge gives the probability that a quantity will be supplied for a given price. Any price can be seen as giving a probability density function, or visually a slice of the ridge, which by definition has an area equal to one.

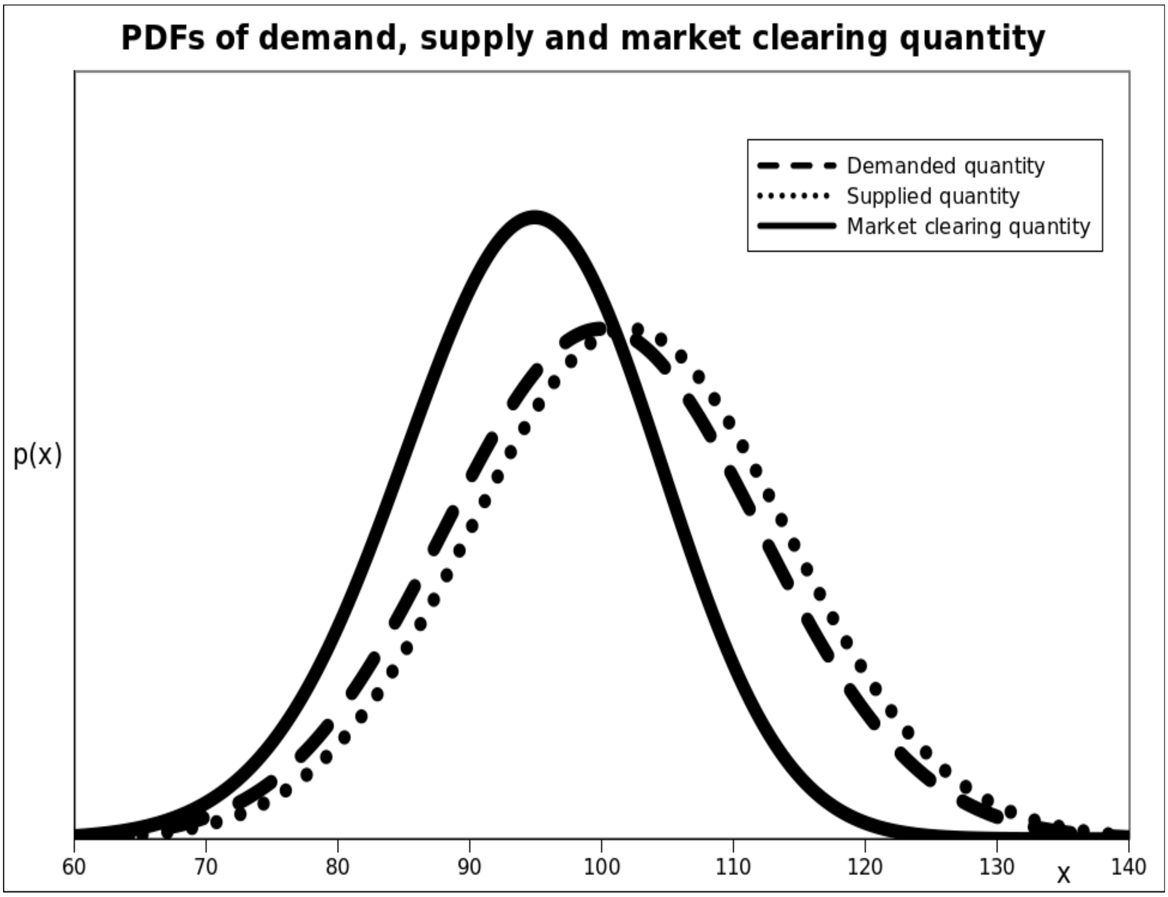

Illustration 6: Probability density functions of demand with $d_l^* = 100$, $\Sigma_{m=1}^{M} \varepsilon_{lm} \sim N(0,12)$, *supply with* $s_l^* = 102$, $\Sigma_{n=1}^{N} \varepsilon_{\ln} \sim N(0,12)$ *and the resulting minimum of demand and supply, or market clearing quantity.*

For any given prices, there is a probability density function for demand and another for supply. The probability density function of the minimum of the two random variables s_l and d_l serves to illustrate the probability that a good is traded over a range of quantities given constant prices.

Illustration 6 shows a simulation of probability density functions of demand with $d_l{*}=100$, $\Sigma_{m=1}^{M}\varepsilon_{lm} \sim N(0,12)$, supply with $s_l{*}=102$, $\Sigma_{n=1}^{N}\varepsilon_{\ln} \sim N(0,12)$ and the resulting minimum of demand and supply, for a given price. Demand and supply have almost overlapping distributions, but the distribution of the market clearing quantity $min\{s_l, d_l\}$ is significantly different from the distributions of both demand and supply. Specifically, the distribution of the market clearing quantity has a lower mean and lower variance than both demand and supply.

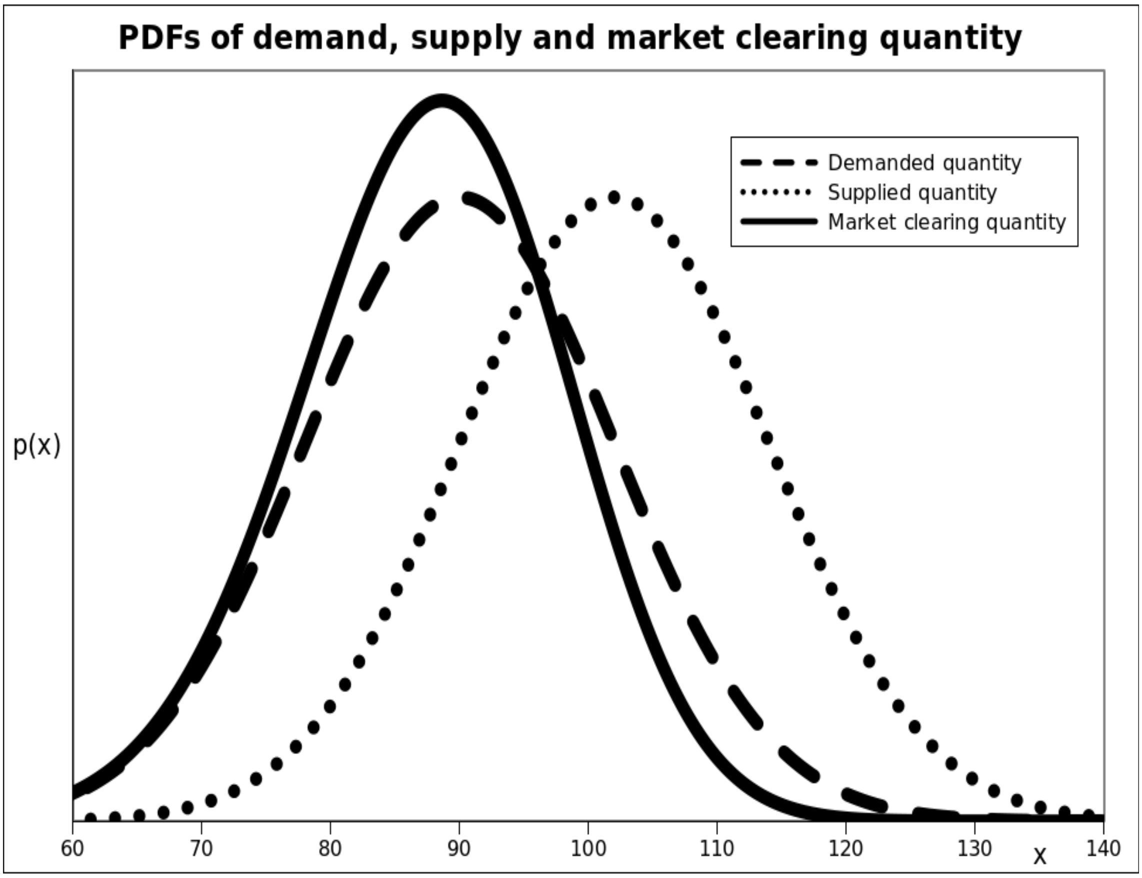

Illustration 7: Probability density functions of demand with $d_l{*}=90$, $\Sigma_{m=1}^{M}\varepsilon_{lm} \sim N(0,12)$, *supply with* $s_l{*}=102$, $\Sigma_{n=1}^{N}\varepsilon_{\ln} \sim N(0,12)$ *and the resulting minimum of demand and supply, or market clearing quantity.*

Illustration 7 shows a simulation of probability density functions of demand with $d_l^* = 90$, $\sum_{m=1}^{M} \varepsilon_{lm} \sim N(0,12)$, supply with $s_l^* = 102$, $\sum_{n=1}^{N} \varepsilon_{\ln} \sim N(0,12)$ and the resulting minimum of demand and supply. The distribution of the market clearing quantity approaches the distribution that has the lower expected value as the absolute difference between expected demand and supply increases.

Finally, illustration 8 shows a simulation of probability density functions of demand with $d_l^* = 90$, $\sum_{m=1}^{M} \varepsilon_{lm} \sim N(0,18)$, supply with $s_l^* = 102$, $\sum_{n=1}^{N} \varepsilon_{\ln} \sim N(0,9)$ and the resulting minimum of demand and supply. In this simulation, both the distribution of demand and the distribution of supply are symmetrical, but the graph clearly shows that the resulting distribution of the market clearing quantity is heavily skewed, with its left tail closely following the demand distribution.

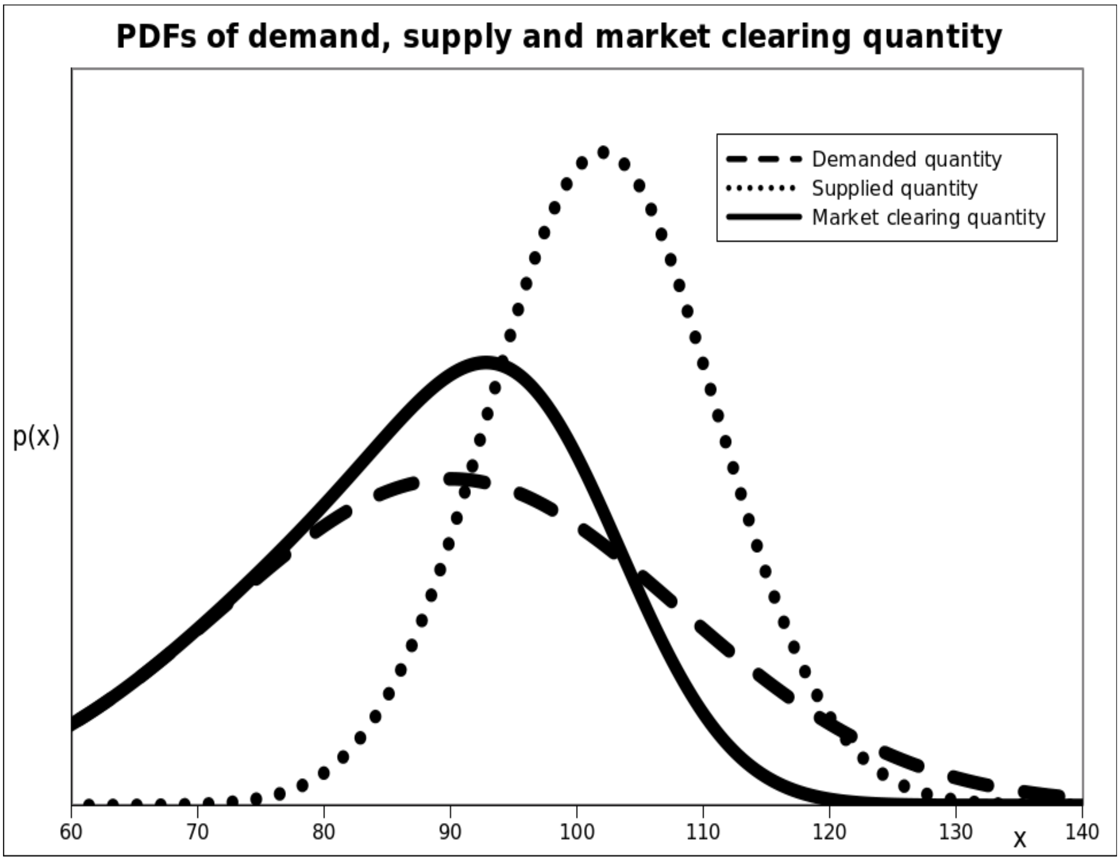

Illustration 8: Probability density functions of demand with $d_l^* = 90$, $\sum_{m=1}^{M} \varepsilon_{lm} \sim N(0,18)$, *supply with* $s_l^* = 102$, $\sum_{n=1}^{N} \varepsilon_{\ln} \sim N(0,9)$ *and the resulting minimum of demand and supply, or market clearing quantity.*

4. Conclusion

This text reviews the utility maximization model and the assumptions on which it relies. Some evidence suggests that these assumptions may not hold, meaning that the existence of a utility function reasonably can be questioned. Yet more evidence suggests that if a utility function does exist, agents often fail to maximize it. A general framework accounting for agents' failure to optimize preferences is suggested. As in the utility maximization model, preferences are assumed to be complete, reflexive, transitive and continuous, which guarantees the existence of a utility function representing preferences. The possibility of violations of these assumptions cannot be denied, but is disregarded in the suggested framework, partly because what appears to be violations can be explained as mistakes within the suggested framework, and partly because of the fundamental importance of the utility function that follows from the assumptions.

Empirical evidence is consistent with agents striving towards but failing to maximize utility, and deviations from utility maximization cannot be characterized as measurement errors alone. Rather, economists are justified in acknowledging that economic decisions often are flawed, i.e. that agents' actual choices often deviates from their utility maximizing choices. Some existing models accounting for sub-optimal choices are discussed, but none of them is accurate and general enough to serve as a theoretical foundation for both classical microeconomics and behavioural economics.

The suggested framework asserts that there is a utility maximizing bundle in a compact budget set and that the agent strives towards choosing such a bundle. Deviations from the optimal choice are described as being caused by a vector $\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_J)$ with each component quantifying how much a quirk j reduces utility. When quirks affecting deviations from the optimal choice are omitted from $\boldsymbol{\gamma}$, choice is probabilistic rather than deterministic, so the framework assigns a probability of choice to every bundle in the budget set. This generalisation of the utility maximization model allows researchers to start with the standard utility maximization model and adjust predictions to account for flaws that cause agents' choices to deviate from the optimal bundle.

The suggested framework clarifies how the utility maximization model can be tweaked to account for flaws in decision making. This sets the stage for a holistic view of sub-optimal behaviour, but extensive research is needed to quantify how the many quirks of economic agents are interrelated. A better understanding of how cognitive biases for which evidence is mounting are connected, might allow for a synthesis of many models in behavioural economics which are currently applicable only to specific situations.

Next, the implications of the general framework on the supply-demand schedule are examined. First, the framework implies that deviations in demand is asymmetrical because the agent typically spends her entire budget on the optimal bundle, so deviations from optimal demand cannot entail spending too much money. Exacerbating this tendency of consumers to underspend rather than overspending, the interaction between sub-optimal consumers and sub-optimal firms is shown to inhibit trade. Expected trade is inefficiently low when consumers and firms make mistakes in general but make optimal choices on average, even without transaction costs and with prices that produces an efficient equilibrium in the standard model. Merely allowing for deviations from optimal choice implies that the quantity produced and consumed on a free market is lower than what is socially optimal.

Many big and important questions remain unanswered. Is an analogous framework tenable for expected utility and for firms? Do some agents' error distribution differ from others' in systematic ways? How does sub-optimal choices affect outcomes when prices are chosen by firms or negotiated between buyers and sellers? Does the shape of consumers' error distributions matter to a firm that sets price? How can the large variety of cognitive biases be simplified into a parsimonious but still general model? Is there a better framework than the suggested one?

Adding an error distribution dimension to choice theory requires increasingly complex analysis and more computational power, and it is not clear how further research on this path will benefit economics or science at large. Researchers finding it worthwhile examining error distributions and their implications on economics in closer detail might find increasingly powerful computers a useful or even necessary tool. For now, using fairly simple mathematics, this paper demonstrates that acknowledging the mistakes of economic agents has important implications in economic theory. Acknowledging deviations from optimal choices is not as trivial or innocent a stance as it superficially may appear.

5. References

Andreoni, James, and John Miller. "Giving according to GARP: An experimental test of the consistency of preferences for altruism." Econometrica 70.2 (2002): 737-753.

Andreou, Chrisoula. "Understanding procrastination." Journal for the Theory of Social Behaviour 37.2 (2007): 183-193.

Ariely, Dan, et al. "Large stakes and big mistakes." The Review of Economic Studies 76.2 (2009): 451-469.

Arkes, Hal R., and Catherine Blumer. "The psychology of sunk cost." Organizational behavior and human decision processes 35.1 (1985): 124-140.

Asch, Solomon E. "Effects of group pressure upon the modification and distortion of judgments." Groups, Leadership, and Men. S (1951): 222-236.

Aumann, Robert J. "Utility theory without the completeness axiom." Econometrica: Journal of the Econometric Society (1962): 445-462.

Battalio, Raymond C., et al. "A test of consumer demand theory using observations of individual consumer purchases*." Economic Inquiry 11.4 (1973): 411-428.

Bergman, Oscar, et al. "Anchoring and cognitive ability." Economics Letters 107.1 (2010): 66-68.

Berle, Adolf A., and C. Gardiner. "Means, The modern corporation and private property." New York: Macmillan 3 (1932): 45.

Bewley, Truman F. "Knightian decision theory. Part I." Decisions in economics and finance 25.2 (2002): 79-110.

Bikhchandani, Sushil, and Sunil Sharma. "Herd behavior in financial markets: A review." (2000).

Birnbaum, Michael H., and Ulrich Schmidt. "An experimental investigation of violations of transitivity in choice under uncertainty." Journal of Risk and Uncertainty 37.1 (2008): 77-91.

Camerer, Colin, George Loewenstein, and Drazen Prelec. "Neuroeconomics: How neuroscience can inform economics." Journal of economic Literature (2005): 9-64.

Canessa, Nicola, et al. "The Functional and Structural Neural Basis of Individual Differences in Loss Aversion." The Journal of Neuroscience 33.36 (2013): 14307-14317.

Caplin, Andrew, and Mark Dean. "Dopamine, reward prediction error, and economics." The Quarterly Journal of Economics 123.2 (2008): 663-701.

Chang, Ruth. "The Possibility of Parity*." Ethics 112.4 (2002): 659-688.

Danan, Eric, and Anthony Ziegelmeyer. Are Preferences Complete?: An Experimental Measurement of Indecisiveness Under Risk. Max-Planck-Inst. for Research into Economic Systems, Strategic Interaction Group, 2006.

Edgeworth, Francis Ysidro. Mathematical psychics: An essay on the application of mathematics to the moral sciences. C. Keagann Paul, 1881.

Eliaz, Kfir, and Efe A. Ok. "Indifference or indecisiveness? Choice-theoretic foundations of incomplete preferences." Games and Economic Behavior 56.1 (2006): 61-86.

Englich, Birte, Thomas Mussweiler, and Fritz Strack. "Playing dice with criminal sentences: The influence of irrelevant anchors on experts' judicial decision making." Personality and Social Psychology Bulletin 32.2 (2006): 188-200.

Epley, Nicholas, and Thomas Gilovich. "The anchoring-and-adjustment heuristic Why the adjustments are insufficient." Psychological science 17.4 (2006): 311-318.

Festinger, Leon, and James M. Carlsmith. "Cognitive consequences of forced compliance." The Journal of Abnormal and Social Psychology 58.2 (1959): 203.

Fishburn, Peter C. "Nontransitive preferences in decision theory." Journal of Risk and Uncertainty 4.2 (1991): 113-134.

Fisher, Irving. "Mathematical Investigations in the Theory of Value and Prices." Bull. Amer. Math. Soc. 2 (1893), 204-211 DOI: http://dx. doi. org/10.1090/S0002-9904-1893-00145-6 PII (1893): 0002-9904.

Friedman, Milton. "The methodology of positive economics." Essays in positive economics 3.8 (1953).

Fudenberg, Drew. "Advancing Beyond" Advances in Behavioral Economics"." Journal of Economic Literature (2006): 694-711.

Gächter, Simon, et al. "Are experimental economists prone to framing effects? A natural field experiment." Journal of Economic Behavior & Organization 70.3 (2009): 443-446.

Garland, Howard, and Stephanie Newport. "Effects of absolute and relative sunk costs on the decision to persist with a course of action." Organizational Behavior and Human Decision Processes 48.1 (1991): 55-69.

Gigerenzer, Gerd, and Daniel G. Goldstein. "Reasoning the fast and frugal way: models of bounded rationality." Psychological review 103.4 (1996): 650.

Gneezy, Uri, and Aldo Rustichini. "Pay enough or don't pay at all." The Quarterly Journal of Economics 115.3 (2000): 791-810.

Gul, Faruk, and Wolfgang Pesendorfer. "The case for mindless economics." The foundations of positive and normative economics (2008): 3-39.

Hanemann, W. Michael. "Willingness to pay and willingness to accept: how much can they differ?." The American Economic Review (1991): 635-647.

Harbaugh, William T., and Kate Krause. "Children's altruism in public good and dictator experiments." Economic Inquiry 38.1 (2000): 95-109.

Harbaugh, William T., Kate Krause, and Timothy R. Berry. "GARP for kids: On the development of rational choice behavior." American Economic Review (2001): 1539-1545.

Heath, Chip. "Escalation and de-escalation of commitment in response to sunk costs: The role of budgeting in mental accounting." Organizational Behavior and Human Decision Processes 62.1 (1995): 38-54.

Hirshleifer, David, and Tyler Shumway. "Good day sunshine: Stock returns and the weather." The Journal of Finance 58.3 (2003): 1009-1032.

Hjertstrand, Per, and James L. Swofford. "A Stochastic Test of the Number of Revealed Preference Violations." (2009).

Hong, Harrison, Jeffrey D. Kubik, and Jeremy C. Stein. "Social interaction and stock-market participation." The journal of finance 59.1 (2004): 137-163.

Jacowitz, Karen E., and Daniel Kahneman. "Measures of anchoring in estimation tasks." Personality and Social Psychology Bulletin 21 (1995): 1161-1166.

Jaffray, Jean-Yves. "Existence of a continuous utility function: an elementary proof." Econometrica: Journal of the Econometric Society (1975): 981-983.

Kahneman, Daniel, and Amos Tversky. "Prospect theory: An analysis of decision under risk." Econometrica: Journal of the Econometric Society (1979): 263-291.

Kahneman, Daniel, Edward Diener, and Norbert Schwarz, eds. Well-being: Foundations of hedonic psychology. Russell Sage Foundation, 2003.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. "Anomalies: The endowment effect, loss aversion, and status quo bias." The journal of economic perspectives (1991): 193-206.

Lee, Leonard, On Amir, and Dan Ariely. "In search of homo economicus: cognitive noise and the role of emotion in preference consistency." Journal of Consumer Research 36.2 (2009): 173-187.

Lin, William, G. W. Dean, and C. V. Moore. "An empirical test of utility vs. profit maximization in agricultural production." American Journal of Agricultural Economics 56.3 (1974): 497-508.

List, John A. "Neoclassical theory versus prospect theory: Evidence from the marketplace." Econometrica 72.2 (2004): 615-625.

Lockwood, Michael. "Integrated value assessment using paired comparisons." Ecological Economics 25.1 (1998): 73-87.

Loewenstein, George, and Jennifer S. Lerner. "The role of affect in decision making." Handbook of affective science 619.642 (2003): 3.

Loomes, Graham, Chris Starmer, and Robert Sugden. "Observing violations of transitivity by experimental methods." Econometrica: Journal of the Econometric Society (1991): 425-439.

Machina, Mark J. "" Expected Utility" Analysis without the Independence Axiom." Econometrica: Journal of the Econometric Society (1982): 277-323.

Machina, Mark J. "Stochastic choice functions generated from deterministic preferences over lotteries." The economic journal (1985): 575-594.

Mandler, Michael. "Indifference and incompleteness distinguished by rational trade." Games and Economic Behavior 67.1 (2009): 300-314.

Marley, A. A. J. "Random utility models with all choice probabilities expressible as 'functions' of the binary choice probabilities." Mathematical Social Sciences 3.1 (1982): 39-56.

Marley, Anthony Alfred John. "Some probabilistic models of simple choice and ranking." Journal of Mathematical Psychology 5.2 (1968): 311-332.

Mas-Colell, Andreu, Michael Dennis Whinston, and Jerry R. Green. Microeconomic theory. Vol. 1. New York: Oxford university press, 1995.

May, Kenneth O. "Intransitivity, utility, and the aggregation of preference patterns." Econometrica: Journal of the Econometric Society (1954): 1-13.

McFadden, Daniel. "Econometric models for probabilistic choice among products." Journal of Business 53.3 (1980): S13-S29.

McFadden, Daniel. "Rationality for economists?." Journal of Risk and Uncertainty 19:1-3 (1999): 73-105.

McNeil, Barbara J., et al. "On the elicitation of preferences for alternative therapies." The New England journal of medicine 306.21 (1982): 1259-1262.

Mishra, Himanshu, Arul Mishra, and Dhananjay Nayakankuppam. "Money: A bias for the whole." Journal of Consumer Research 32.4 (2006): 541-549.

Moulton, Brent R. "Random group effects and the precision of regression estimates." Journal of econometrics 32.3 (1986): 385-397.

Mussweiler, Thomas, and Birte Englich. "Subliminal anchoring: Judgmental consequences and underlying mechanisms." Organizational Behavior and Human Decision Processes 98.2 (2005): 133-143.

Navarick, Douglas J., and Edmund Fantino. "Stochastic transitivity and unidimensional behavior theories." Psychological Review 81.5 (1974): 426.

Neumann von J, Morgenstern O. "Theory of games and economic behavior." Princeton University Prress, Princeton (1944).

Ng, Yew-Kwang. "Sub-semiorder: A model of multidimensional choice with preference intransitivity." Journal of Mathematical Psychology 16.1 (1977): 51-59.

Nieoullon, André. "Dopamine and the regulation of cognition and attention." Progress in neurobiology 67.1 (2002): 53-83.

Nooteboom, Bart. "Intransitive preferences in retailing." Service Industries Journal 4.1 (1984): 82-92.

O'Connor, Annette M. Cormier, et al. "Eliciting preferences for alternative cancer drug treatments the influence of framing, medium, and rater variables." Medical Decision Making 5.4 (1985): 453-463.

Ok, Efe A. "Utility representation of an incomplete preference relation." Journal of Economic Theory 104.2 (2002): 429-449.

Peleg, Bezalel. "Utility functions for partially ordered topological spaces." Econometrica: Journal of the Econometric Society (1970): 93-96.

Pessiglione, Mathias, et al. "Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans." Nature 442.7106 (2006): 1042-1045.

Previc, Fred H. "Dopamine and the origins of human intelligence." Brain and cognition 41.3 (1999): 299-350.

Quirk, James P., and Rubin Saposnik. Introduction to general equilibrium theory and welfare economics. New York: McGraw-Hill, 1968.

Raghubir, Priya, and Joydeep Srivastava. "The denomination effect." Journal of Consumer Research 36.4 (2009): 701-713.

Raz, Joseph. The morality of freedom. Oxford University Press, 1986.

Regenwetter, Michel, Jason Dana, and Clintin P. Davis-Stober. "Transitivity of preferences." Psychological Review 118.1 (2011): 42.

Rieskamp, Jörg, Jerome R. Busemeyer, and Barbara A. Mellers. "Extending the bounds of rationality: evidence and theories of preferential choice." Journal of Economic Literature (2006): 631-661.

Roelofsma, Peter HMP, and Daniel Read. "Intransitive intertemporal choice." Journal of Behavioral Decision Making 13.2 (2000): 161-177.

Rosenberger, Randall S., et al. "Measuring dispositions for lexicographic preferences of environmental goods: integrating economics, psychology and ethics." Ecological Economics 44.1 (2003): 63-76.

Samuelson, Paul A. "A note on the pure theory of consumer's behaviour." Economica (1938): 61-71.

Savage, Leonard J. "The foundations of statistics." (1954).

Schoemaker, Paul JH. "The expected utility model: Its variants, purposes, evidence and limitations." Journal of economic literature (1982): 529-563.

Schultz, Wolfram. "Predictive reward signal of dopamine neurons." Journal of neurophysiology 80.1 (1998): 1-27.

Seabright, Paul B. "Continuous preferences and discontinuous choices: How altruists respond to incentives." The BE Journal of Theoretical Economics 9.1 (2009).

Shafir, Sharoni. "Intransitivity of preferences in honey bees: support for'comparative'evaluation of foraging options." Animal Behaviour 48.1 (1994): 55-67.

Shampanier, Kristina, Nina Mazar, and Dan Ariely. "Zero as a special price: The true value of free products." Marketing Science 26.6 (2007): 742-757.

Shizgal, Peter. "Neural basis of utility estimation." Current opinion in neurobiology 7.2 (1997): 198-208.

Simmons, Joseph P., Robyn A. LeBoeuf, and Leif D. Nelson. "The effect of accuracy motivation on anchoring and adjustment: Do people adjust from provided anchors?." Journal of personality and social psychology 99.6 (2010): 917.

Simon, Herbert A. "A behavioral model of rational choice." The quarterly journal of economics 69.1 (1955): 99-118.

Simonson, Itamar, and Amos Tversky. "Choice in context: Tradeoff contrast and extremeness aversion." Journal of Marketing Research (JMR) 29.3 (1992).

Sippel, Reinhard. "An experiment on the pure theory of consumer's behaviour*." The Economic Journal 107.444 (1997): 1431-1444.

Smith, Vernon L., and James M. Walker. "Monetary rewards and decision cost in experimental economics." Economic Inquiry 31.2 (1993): 245-261.

Sokol-Hessner, Peter, Colin F. Camerer, and Elizabeth A. Phelps. "Emotion regulation reduces loss aversion and decreases amygdala responses to losses." Social cognitive and affective neuroscience 8.3 (2013): 341-350.

Soman, Dilip, and Amar Cheema. "The effect of windfall gains on the sunk-cost effect." Marketing Letters 12.1 (2001): 51-62.

Soman, Dilip. "The mental accounting of sunk time costs: Why time is not like money." Journal of Behavioral Decision Making 14.3 (2001): 169-185.

Spash, Clive L., and Nick Hanley. "Preferences, information and biodiversity preservation." Ecological economics 12.3 (1995): 191-208.

Spash, Clive L. "Ecosystems, contingent valuation and ethics: the case of wetland re-creation." Ecological economics 34.2 (2000): 195-215.

Staw, Barry M. "Knee-deep in the big muddy: A study of escalating commitment to a chosen course of action." Organizational behavior and human performance 16.1 (1976): 27-44.

Stevens, Thomas H., et al. "Measuring the existence value of wildlife: what do CVM estimates really show?." Land Economics 67.4 (1991).

Strack, Fritz, and Thomas Mussweiler. "Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility." Journal of Personality and Social Psychology 73.3 (1997): 437.

Thaler, Richard. "Toward a positive theory of consumer choice." Journal of Economic Behavior & Organization 1.1 (1980): 39-60.

Tom, Sabrina M., et al. "The neural basis of loss aversion in decision-making under risk." Science 315.5811 (2007): 515-518.

Tsang, Edward PK. "Computational intelligence determines effective rationality." International Journal of Automation and Computing 5.1 (2008): 63-66.

Tversky, Amos, and Daniel Kahneman. "Judgment under uncertainty: Heuristics and biases." science 185.4157 (1974): 1124-1131.

Tversky, Amos, and Daniel Kahneman. "The framing of decisions and the psychology of choice." Science 211.4481 (1981): 453-458.

Tversky, Amos. "Elimination by aspects: A theory of choice." Psychological review 79.4 (1972): 281.

Tversky, Amos. "Intransitivity of preferences." Psychological review 76.1 (1969): 31.

Waite, Thomas A. "Intransitive preferences in hoarding gray jays (Perisoreus canadensis)." Behavioral Ecology and Sociobiology 50.2 (2001): 116-121.

Walker, Joan, and Moshe Ben-Akiva. "Generalized random utility model." Mathematical Social Sciences 43.3 (2002): 303-343.

Weirich, Paul. "Utility Maximization Generalized." Journal of Moral Philosophy 5.2 (2008): 282-299.

Wilson, Timothy D., et al. "A new look at anchoring effects: basic anchoring and its antecedents." Journal of Experimental Psychology: General 125.4 (1996): 387.

Wright, William F., and Urton Anderson. "Effects of situation familiarity and financial incentives on use of the anchoring and adjustment heuristic for probability assessment." Organizational Behavior and Human Decision Processes 44.1 (1989): 68-82.

Zeng, Jianmin, et al. "An fMRI study on sunk cost effect." Brain research 1519 (2013): 63-70.

Appendix A – Proof of the Existence of a Demand Function

Proposition: If B is compact and convex, and $\succsim$ is complete, reflexive, transitive, continuous and strictly convex, then a demand function exists.

Proof: The existence of a utility function representing preferences is guaranteed by completeness, reflexivity, transitivity and continuity. A demand correspondence gives the set of bundles that maximizes utility given a budget set defined by prices and wealth, or formally $\boldsymbol{x}^*(\boldsymbol{p},w) = \operatorname{argmax}_{\boldsymbol{x}\in B(\boldsymbol{p},w)} u(\boldsymbol{x})$. The demand correspondence is a demand function iff $\exists!\,\boldsymbol{x}\in\boldsymbol{x}^*(\boldsymbol{p},w)$ for every $\boldsymbol{p}$ and w. Thus, all that is required to prove the existence of a demand function is to show that there is a unique utility maximizing bundle $\boldsymbol{x}$ for all prices and wealth levels.

Note first that the Walrasian budget set $B=\{\boldsymbol{x}\in X: 0\leq \boldsymbol{p}\cdot\boldsymbol{x}\leq w\}$ is compact. B is closed because it contains all of its own limit points, or equivalently because its complement $\{\boldsymbol{x}\in X: \boldsymbol{p}\cdot\boldsymbol{x}<0 \vee w<\boldsymbol{p}\cdot\boldsymbol{x}\}$ is an open set. B is bounded from below by the assumption that all quantities are non-negative, and bounded from above by the assumption that no more than a finite wealth can be spent on goods with strictly positive prices. Thus, B is closed and bounded, or in other words compact.

Knowing that B is compact, the extreme value theorem establishes that the agent's utility maximization problem $\max_x u(\boldsymbol{x})$ subject to $\boldsymbol{x}\in B$ implies $\exists \boldsymbol{x}\in B: \ u(\boldsymbol{x}) \geq u(\boldsymbol{x}') \ \ \forall\, \boldsymbol{x}' \in B$, i.e. that the agent always can choose a utility maximizing bundle in B. Thus $\boldsymbol{x}^*(\boldsymbol{p},w)\neq\emptyset$, establishing the existence of a demand correspondence.

To guarantee that the demand correspondence is a demand function, we begin by showing that B is convex. Consider any $\boldsymbol{x},\boldsymbol{x}'\in B$ and define $\boldsymbol{x}''=\alpha\boldsymbol{x}+(1-\alpha)\boldsymbol{x}'$ with $\alpha\in[0,1]$. By the definition of B, the cost of $\boldsymbol{x}$ is $0\leq \boldsymbol{p}\cdot\boldsymbol{x}\leq w$ and the cost of $\boldsymbol{x}'$ is $0\leq\boldsymbol{p}\cdot\boldsymbol{x}'\leq w$. The cost of $\boldsymbol{x}''$ is $\boldsymbol{p}\cdot\boldsymbol{x}'' = \boldsymbol{p}\cdot[\alpha\boldsymbol{x}+(1-\alpha)\boldsymbol{x}'] = \alpha\boldsymbol{p}\cdot\boldsymbol{x}+(1-\alpha)\boldsymbol{p}\cdot\boldsymbol{x}'$ and we conclude that $0\leq \alpha\boldsymbol{p}\cdot\boldsymbol{x} + (1-\alpha)\boldsymbol{p}\cdot\boldsymbol{x}' \leq w \ \ \forall\alpha\in[0,1]$ or simply $0\leq\boldsymbol{p}\cdot\boldsymbol{x}''\leq w$, implying that $\boldsymbol{x}''\in B$. The established result $\boldsymbol{x},\boldsymbol{x}'\in B \ \Rightarrow \ \alpha\boldsymbol{x}+(1-\alpha)\boldsymbol{x}'\in B \ \ \forall\alpha\in[0,1]$ is the definition of convexity.

Next, assume strictly convex preferences, meaning that for any $\boldsymbol{x}, \boldsymbol{x}' \in X$ with $\boldsymbol{x} \neq \boldsymbol{x}'$ we have $\alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{x}' \succ \boldsymbol{x}, \boldsymbol{x}'$ $\forall \alpha \in [0,1]$. In terms of a utility function, for any $\boldsymbol{x}, \boldsymbol{x}' \in X$ with $\boldsymbol{x} \neq \boldsymbol{x}'$ we have $u(\alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{x}') > u(\boldsymbol{x}) = u(\boldsymbol{x}')$ $\forall \alpha \in [0,1]$.

Now consider any $\boldsymbol{x}, \boldsymbol{x}' \in \boldsymbol{x}^*(\boldsymbol{p}, w)$ with $\boldsymbol{x} \neq \boldsymbol{x}'$ and define $\boldsymbol{x}'' = \alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{x}'$. Seeing that $\boldsymbol{x}^*(\boldsymbol{p}, w) \subseteq B$, the convexity of B implies $\boldsymbol{x}'' \in B$. By the definition of the demand correspondence we have $u(\boldsymbol{x}) = u(\boldsymbol{x}')$ and by the definition of strictly convex preferences we have $u(\boldsymbol{x}'') > u(\boldsymbol{x}) = u(\boldsymbol{x}')$. But the existence of an $\boldsymbol{x}'' \in B$ giving strictly higher utility than $\boldsymbol{x}$ and $\boldsymbol{x}'$ contradicts $\boldsymbol{x}, \boldsymbol{x}' \in \boldsymbol{x}^*(\boldsymbol{p}, w)$. Thus there can only be one bundle in $\boldsymbol{x}^*(\boldsymbol{p}, w)$ so that $\exists! \boldsymbol{x} \in \boldsymbol{x}^*(\boldsymbol{p}, w)$ for every $\boldsymbol{p}$ and w, establishing the existence of a demand function.

$\blacksquare$