AALBOR UNIVERSITET

Optimizing Step by Step

A Theoretical Framework for Behavioural Economics Based on the Variational Free Energy Principle

Morten Henriksen ______________________________

03-06-2019

Specialevejleder: Jørgen Stamhus

Abstract

The purpose of this thesis is to offer a new framework for understanding action, optimization and choice when applied to economic theory more generally. By drawing upon the concept known as the variational free energy principle, the thesis will explore how this principle can be used to temper rational choice theory by re-formulating how agents optimize. The approach will result in agent behaviour that encompasses a wide range of so called cognitive biases, as seen in the scientific literature of behavioural economics, but instead of using these biases as further indications of market inefficiencies or market failures, the thesis will likewise attempt to show the limits to which these biases can inform or critique standard economic theory. The thesis therefore offers up a “middle of the road” approach, in which the neoclassical agent is not quite as rational as rational choice theory assumes, but at the same time, not quite as irrational as behavioural economics would often have us believe.

Table of Contents

Optimizing Step by Step

A Theoretical Framework for Behavioural Economics Based on the Variational Free Energy Principle

"The curious task of economics is to demonstrate to men how little they know about what they imagine they can design" Hayek (The Fatal Conceit, 1988 p. 76)

1. Introduction

The purpose of this thesis is to offer a new framework for understanding action, optimization and choice when applied to economic theory more generally. By drawing upon the concept known as the variational free energy principle, the thesis will explore how this principle can be used to temper rational choice theory by re-formulating how agents optimize. The approach will result in agent behaviour that encompasses a wide range of so called cognitive biases, as seen in the scientific literature of behavioural economics, but instead of using these biases as further indications of market inefficiencies or market failures, the thesis will likewise attempt to show the limits to which these biases can inform or critique standard economic theory. The thesis therefore offers up a "middle of the road" approach, in which the neoclassical agent is not quite as rational as rational choice theory assumes, but at the same time, not quite as irrational as behavioural economics would often have us believe.

The variational free energy principle is a relatively new theory in the neuroscientific arena, and is still, by some neuro-scientists, considered to be a fringe theory, although it is steadily gaining traction. This thesis is therefore the first attempt to apply the theory to economics, and despite my best efforts, it may as a consequence be difficult to follow at times. For this I apologize.

1.1. Conflicting Views

Standard neo-classical microeconomics has since the publishing of Paul Samuelson's *Foundations of Economic Analysis* (1946), been working with a highly mathematized and thermodynamic interpretation of economics, wherein individuals are formally treated as highly rational and optimizing agents all completely interchangeable with each other. In his book, Samuelson proposed a new methodology for the economic science called operationalism.

The main focus of this methodology was to attempt to transform qualitative phenomenon onto quantitative units through mathematical operations, ultimately in an attempt to measure the seemingly un-measurable such as preference or utility. This led to the development of revealed preference, in contrast to demonstrated preference, where the possibility of predicting future preference schemes or anticipating a specific utility gain from consumption, is made explicit due to the ability to accurately formulate a linear function of past behaviour (Blaug, 1992).

Yet in our respective functions as acting human beings, to act upon the world in any meaningful sense according to operationalism, thus presumes inaction on behalf on the actor since *to act* is *to change*, instantly defeating any notion of stability and equilibrium with respect to time. What is meant by this is that action is a form of intervention, geared towards rectifying or alleviating felt distress, discomfort or dissatisfaction with states of the world affecting the individual. It is therefore more accurate to conceptualize the nature of reality as a quasi-deterministic process ruled by entropy, whereby action serves to counterbalance this ever present condition of nature.

When framing economics in terms of mutually determined action/reaction relations, we automatically imply a stable and predictable relationship between consequent and antecedent. For this reason, we need a model of economics in which the antecedent, here represented as an agent, has similar stable characteristics from which consequents flow in a predictable manner. The logic is clear and concise. If markets are efficient, then agents operating in the market are rational. Markets appear to be efficient when free from coercion; therefore agents must be rational when similarly free to choose. But what if these stable characteristic prove to be not so stable on closer examination? What if the hyper rational agent of the neo-classical doctrine fails to display hyper rational behaviour in his or her every day actions and interactions? Only two possible answers come to mind. Either markets cannot be efficient, and therefore has to be coerced. Or alternatively there must be something wrong with the specification of the terms “efficient” and “rational agent” as they apply to economics. There is off course a third answer to this question, but that would mean ignoring or dismissing a rather substation body of work in behavioural economics. A body that has since the early 80’s consistently demonstrated numerous flaws in the rational quality of the neoclassical agent.

It does not escape me that the above posed question would be answered in a manner of seconds, and without further deliberation, by practically every economist I have ever had the privilege of knowing. That we humans are fallible and our outlooks uncertain, strongly implies that markets cannot be stable or efficient without third party intervention. But to my mind, that simply reframes the problem. Who then is not fallible, and who then has certain outlooks to so assuredly know when or even how to intervene? Usually this will be answered by referencing the careful study of empirical data, noting what resulted from what measures under what circumstances. This however assumes the aforementioned action/reaction relations underpinning the neo-classical doctrine, and must likewise presuppose stable and predictable agents who are, if not rational, then perhaps predictably irrational as the psychologist and behavioural economist Dan Ariely titled his book from 2008.

Behavioural economics is however not a school of thought in the traditional sense, but more a collection of ad-hoc observations of human behaviour in various decision making scenarios (Wilkinson & Klaes, 2012). For this reason some economists, especially of the behavioural persuasion, argue that the insights from behavioural studies and experiments are meant to compliment and update standard neoclassical analysis (Rabin, 2002; Thaler, 2015; Loewenstein & Camerer, 2002). As such, there is no epistemological conflict between the two approaches, even though the methodological differences are substantial off course. It is therefore perceived to be very unfortunate that standard neoclassical analysis have been highly dismissive of the various insights and contributions from behavioural economic research (Loewenstein & Angner, 2006). Yet if we consider the sheer number of biases that has to be accounted for according to psychologists and behavioural economists alike, a number rapidly approaching 200, then it seems unlikely that standard neoclassical analysis could ever hope to include the entirety of these interjections. Richard A. Posner, a Harvard law professor who has written extensively on rational choice theory as it applies to both law and economics, has concluded that the most important contributions from behavioural studies have already been included into rational choice theory for quite some time now.

Posner argues that contributions such as *Bounded rationality, bounded willpower, Bounded self-interest, framing effects, hindsight bias, the endowment effect* and *hyperbolic discounting,* can all be analysed within the framework of rational choice, and offers up numerous examples on how this is typically done. If we take the concept of hyperbolic discounting for example, this phenomenon can be perfectly explained and analysed within rational choice theory by referencing the cost of information. Similarly, the concept of bounded willpower can be completely rational if agents were permitted to be comprised of multiple sub-agents with opposing preferences (Posner, 1998; Posner, 2002). There would therefore be a clear divide between the young self and the old self, as well as the future oriented self who values health and longevity, and the present oriented self how wants to indulge in immediate pleasures, none of which would require the need to invoke irrationality. Consequently Posner argues that many behavioural economists caricature rational choice theory and underestimate its applicability to real world situations (Posner, 1998; Posner, 2002).

The schism between rational choice and behavioural economics is not one easily mended. It is true that rational choice theory on its own terms permits an absurd conceptualization of human beings from the perspective of settled psychology, neurology and even common sense, of which the idea of perfect knowledge is perhaps the worst offender. Other than that, how the visual, auditory and olfactory system synthesises information creates interpretive ambiguity, and while this ambiguity can be represented as probability distributions, or rather probability densities, the noise and lack of precision in how sensory inputs are interpreted must render highly rational behaviour ultimately irrational, and visa versa, by the sheer passing of time. Consider the irrationality of shorting mortgage backed securities back in 2005, or even 2006, for instance, or the post 2008 irrationality of the opposite trade. It is off course quite possible to make a compelling case for rational choice theory in both instances, but it would be equally correct to take the behavioural stance depending on how the various motives and actions involved are technically framed. An old probability riddle by the mathematician Martin Gardner comes to mind. The riddle, known as *the two children problem*, might help to illustrate this point more clearly.

If a man has two children and at least one of them is a girl, what then is the probability that both children are girls? Since we know that children can only come in two categories, (boy or girl), the obvious answer would be 50 %. But it would be equally correct to construe the problem as having three categories, (boy, girl), (girl, boy) and (girl, girl), in which case the answer is 33.3 %. Because of the ambiguity in the information presented, both answers are correct at the same time.

Now the man states that one of his girls is named Lisa, a piece of information that intuitively should not influence the probabilities at all, but in fact it does. The categories in the problem is now: (boy, Lisa), (Lisa, boy), (Lisa, girl) and (girl, Lisa), but since the last two categories satisfies the condition of both children being girls, there must be only two relevant categories: ((boy, Lisa), (Lisa, boy)) and ((Lisa, girl), (girl, Lisa)) = 50%.

The interesting point is that we used seemingly irrelevant information in order to remove the ambiguity present in the original formulation of the riddle, here the irrelevant information being the name of the girl that we already knew were a girl.

Naturally it would not be apt to apply the predicate irrational to the riddles presented above; one can be perfectly in accordance with rational choice theory regardless of having answered 50 or 33.3 %. But when we consider the prevalence of systemic errors, or clusters of errors, in regards to human decisions and judgement, it becomes clear that rational or irrational choices cannot be conditions for optimal or suboptimal outcomes when factoring in time.

A good example of this is one often used by behavioural economists. When investing, rational choice theory suggests that individuals use relevant information in order to maximize utility, here utility being the value of a portfolio of various assets. When new information presents itself, the individual is expected to update his or her beliefs and accordingly rearrange the composition of his or her portfolio to match the changed states of the world; we are expected to take advantage of useful information, so to speak. This however is often touted to be the wrong strategy from a long run perspective. Over time, the investors who systematically disregard new information, and do not take strides to update their beliefs, tend to fare better overall in the stock market (Kahnemann, 2011; Thaler, 2015).

1.2. Complex Systems

That rational or irrational cannot be conditions for optimal or suboptimal outcomes over time, is one of the hallmarks of complex systems, where the concept of emergence destroys the linear relationship between antecedent and consequent or the micro level and the macro level. Complex systems are defined as non-linear systems in which its elements are interrelated, diverse, co-dependent and adaptive, giving the system the ability to generate unpredictable networks, emergent properties and spontaneous order (Turner & Baker, 2019).

A way to visualize a complex system is by imagining a simple problem, like how much flour to add when baking bread. This problem has an optimal solution, in that the function that describes the problem has a clear and definable peak. Now we consider complex problems. These problems have many peaks and valleys of various heights, and the difficult part is to figure out when a local peak might be less optimal than an alternative solution. A complex system on the other hand, is a system in which the topographical landscape itself is constantly shifting and moving.

One of the best examples of a complex system is the brain itself, but the theory of complex systems applies equally well to ecological systems, biological systems as well as economics. In order to get a complex system, its elements have to be tuned in a certain manner. Too low or too high degrees of interrelatedness, divisiveness, co-dependency and adaptability create systems that are either too chaotic or too stable. A too chaotic system has no purpose in the colloquial sense, will never reach equilibrium and will disperse and die out relatively quickly due to entropy. A too stable system will reach equilibrium too quickly, and will either be highly susceptible to chocks and disturbances and shatter due to its rigidness, or never leave equilibrium due to its high level of adaptability. In order to get a complex system, you will therefore need to be somewhere in the middle. Not too chaotic, and not too stable and/or rigid. The variable adaptability is very relevant to the overall discussion of rationality and irrationality, as it can also be interpreted as the level of intelligence of the systems' individual elements (Turner & Baker, 2019).

It is safe to say that rational choice theory operates under the assumption of too high levels of adaptability, in that the equilibrium is an ever present condition of the system, where possible deviations can only come from exogenous chocks. That the chocks have to be exogenous is very important, because if they were generated endogenously, they would have been predicted in advance and diffused before they could become an issue. This means that every action of one of the elements, or individuals, will cause the entire system to adapt. But that is not all. Every action would also be expected to cause this adaptation, which means that every action would take this into account ad infinitum, ultimately creating a perfectly static and dead equilibrium stuck forever on a single peak solution. We can see this more clearly when realizing that rational choice theory actually relies on mutual and instantaneous action, something that will explored more carefully later in the theses. This point on the equilibrium is basically the same one that professor Mises outlined seventy years ago in his book "*Human Action*" from 1949, where he warned the economic profession against becoming too emphatic about the relatively new concept of equilibrium analysis. Professor Mises envisioned the equilibrium as a hypothetical construct, analytically useful only in so far as it establishes a direction towards which everything is tending, but can never reach (Mises, 2007 (1949)). As such, the system will never be in equilibrium because of constant disturbances, analogous to a constantly changing topographic landscape, the very definition of a complex system.

A simple yet relevant example of our limited adaptability is the game invented by Alain Ledoux, where participants are asked to guess what 2/3 of the average combined guesses in the game will be. The participants pick a number between 0 and 100, and the winner will be the individual who comes closest to 2/3 of the average guesses. In rational choice theory, participants are expected to pick the number 0, since this is the Nash equilibrium. This is however not what individuals do. In fact, picking the number 0 is a sure way to lose the game, as the wining picks tends to centre around 22 according to a study by Jean-Robert Tyran and Frederik Roose Øvlisen from the economics department of Copenhagen University (Øvlisen, 2009). In the study we see that the rational choice is not the same as the optimal choice because of the diversity of the constituent elements within the game. Knowing that individuals on average are somewhat rational is the key, which implies a diverse range of adaptability in individuals not commensurate with rational choice theory. As it turns out, the study showed that a non-trivial proportion of the participants actually picked a number above the threshold of 66.67, which should be impossible given the rules of the game (Øvlisen, 2009).

1.3. Too Far in the Other Direction

In contrast to the rational choice theory, behavioural economists have a tendency to push the argument a little too far in the opposite direction. The concepts of bounded rationality, bounded willpower and bounded self-interest are all well and good, but there does not seem to be any limit to where these boundaries actually are. In his book "*predictably irrational"* Dan Ariely uses the cognitive bias know as *anchoring* in an attempt demonstrate how the central law of supply and demand is a flawed concept. He argues that new goods and services entering a market will have no reference prices to which they can be evaluated. Therefore, the producer, importer or purveyor of these goods and services are free to set a price at his or her own discretion, counting on clever marketing to create demand, like associating the good he or she wants to sell with other highly luxurious but unrelated goods. After the price has been set, this will function as an anchor in the minds of the consumer, believing that any price below this anchor must be a good deal, regardless of whether or not it objectively is a good deal based on comparable goods and services (Ariely, 2008).

To demonstrate the concept of anchoring, Ariely asked his students to write down the last two digits of their social security number. He thereafter instructed the students to bid on various goods by asking the students what they thought the items were worth. According to Ariely, the students whose last two digits were high, routinely overvalued the items compared to the students whose last two digits were low (Ariely, 2008).

In his book, Ariely uses the Tahitian black pearl as an example of arbitrary pricing followed by the anchoring effect, although he never considers the numerous of goods and services that have been rejected by the consumer over the years, and therefore never got to make a market impact. It is reasonable to assume that the producers, importers and purveyors of these goods and services all wanted to "trick" the consumer into buying what they sold, but if price setting is arbitrary, and consumers simply follows anchors when evaluating these prices, then why did they fail? The concept of anchoring might be a real psychological phenomenon, but it fails to appreciate that prices have never resided solely within the mind of the individual, and to my knowledge, no economist have ever held this view. Prices are formed from the subjective evaluations of countless individuals, based on many different factors. No one knows what any specific object is worth, that is the very point of a market process that allows for price discovery. And even if Ariely is right about the arbitrary nature of prices, that is a far cry from concluding that a shift in either the demand schedule or the supply schedule will not affect prices, which is all the law of supply and demand really states.

There are many more examples of where the behavioural economics literature takes its arguments a bit too far, and seems to operate under the assumption of a too low level of adaptability in humans, as will be shown and addressed later in the theses. Individuals do not appear to act in the world as much as they simply react to the world, all highly irrational but "luckily" highly manipulable, susceptible to suggestion and nudges, depending off course on whom is doing the nudging.

"On the one hand, the picture of the human being that Jolls, Sunstein and Thaler draw is one of unstable preferences and, what turns out to be related, infinite manipulability (…) It seems then that the politically insulated corps of experts that JST favor would be charged with determining the populace's authentic preferences, which sounds totalitarian to me. On the other hand, nothing in JST's analysis exempts "experts" from the cognitive quirks, from weakness of will, or from concerns with fairness. The expert, too, is behavioral man. Behavioral man behaves in unpredictable ways. Dare we vest responsibility for curing irrationality in the irrational?" (Posner, 1998 p. 1575).

1.4. Structure of the Thesis

So far I have attempted to outline some of the ways in which neoclassical economics and behavioural economics operate under what can be called either to strict or to loose assumptions on human agency. In the following sections I will present a "middle of the road" approach based on the variational free energy principle, but it will unfortunately require a considerable amount of setup in order to do so.

The variational free energy principle is a relatively new theory proposed by neuroscientist Karl Friston, and argues that biological systems counterbalance entropic forces by minimizing the property know as variational free energy. Since entropy can be seen as a time average of surprise, minimizing variational free energy will therefore be the mechanism by which surprise or uncertainty is reduced in the system. This can basically be done in two ways. Either by adjusting position in accordance with beliefs through more standard Bayesian inference and learning models, that is, moving uphill towards a peak in the topographic landscape, or by actively changing the world to match beliefs, that is, by changing the topographic landscape itself.

The ability to actively affect the world is the strength of the variational free energy principle, and why will be explained in depth later. For now it will be enough to state that the variational free energy principle gives us the ability to frame optimization problems in terms of action sequences or policies. This provides an agent with the ability to adjust course as well as influence states in the world during this action sequences. The theory is however based on Bayesian statistics and something called the Bayesian brain hypothesis, which is why we will have to introduce these two concepts to the reader in order to give a clearer picture of the thesis.

Once the variational free energy principle has been introduced, we then have to explain how this principle applies to agency in general, by creating a practical framework that can be compared with both neoclassical and behavioural economics. We therefore also need to go through rational choice theory and the expected utility hypothesis, as well as various cognitive biases from the behavioural economics literature.

First rational choice and expected utility will be introduced, followed by Bayesian statistics and the Bayesian brain hypothesis. Next the variational free energy principle will be presented, along with an explanation of how this principle deviates from more traditional Bayesian methods, methods also used in the expected utility hypothesis.

After this, a mode of agency will be proposed explaining the process by which an agent has to optimize, presupposing that the agent is trying to optimize within a complex system by minimizing variational free energy. This will have serious implications for the rational choice theory, which is why we will then have to extract out some principles of action from our new mode of agency, in contrast to the axioms of rational choice. This will be in order to conserve the basic foundation of economics, often described as the laws of economics, in so far as they are not in violation of the variational free energy principle.

Finally this “middle of the road” approach will be compared to the behavioural economics literature in an attempt to see how well the approach handles cognitive biases. If this proves out to be successful, then we may have found a theoretical framework for behavioural economics that can unify the ad-hoc collection of theories and insights that the branch presently is comprised of.

The question that this thesis will seek to answer will therefore be:

Can the principle of variational free energy be used to explain how agents optimize, while simultaneously allowing for the presence of cognitive biases?

2. Rational Choice and the Expected Utility Hypothesis

In this section, rational choice and the expected utility hypothesis will be presented, followed by an initial critique of its mathematical foundation. The critique will be centred on the notion of mutually determined action, a direct implication of the expected utility hypothesis not permissible under the variational free energy principle. Later this critique will be expanded to include the axioms of rational choice, as well as the idea of positive affine transformations, the very process by which subjective value is transformed into ordinal and/or cardinal utility.

2.1. Axioms of Rational Choice

Rational choice and the expected utility hypothesis comprise the backbone of neoclassical economics, and likewise provide a mathematical foundation for operationalistic practises in economics. In essence the hypothesis states that individuals act as if trying to maximize a function of expected utility with due consideration to risk expressed as probability distributions over various gambles. As demonstrated by Morgenstern and Neumann (1947), in order for the hypothesis to work, Individuals need to act in accordance with a set of assumptions, which enables the economist to use positive affine transformations on subjective value, giving utility itself an ordinal, if not cardinal value. More formally these assumptions are referred to as axioms of rational choice, and even though the axioms may not be descriptive in nature, they are never the less widely held to be the normative standard by which individuals ought to act in order for the neoclassical system to be efficient. Given that economics can only be expressed formally through the use of representative agents with uniform characteristics, the fact that markets appear to be efficient when free from coercive forces, must therefore demonstrate the comparable efficiency or rationality of the actual individual, and hereby confirm the axioms of expected utility. Before we analyse this claim, it will however first be necessary to state the axioms as well as delving a little deeper into expected utility under risk and uncertainty.

Completeness: A preference ordering is complete if and only if, for any two outcomes X and Y, an individual prefers X to Y, prefers Y to X, or is indifferent between the two. Formally this can be represented as $(x_1, x_2) > (y_1, y_2)$ and $(x_1, x_2) < (y_1, y_2)$ and $(x_1, x_2) \sim (y_1, y_2)$, with respect to a class of outcomes or a basket of goods.

Transitivity: For any three outcomes, X, Y and Z, if X is preferred to Y, and Y is preferred to Z, then X must be preferred to Z as well. $if\ (x_1, x_2) > (y_1, y_2)\ and\ (y_1, y_2) > (z_1, z_2)\ then\ (x_1, x_2) > (z_1, z_2)$.

Reflexivity: $(x_1, x_2) \geq (x_1, x_2)$, meaning that any outcome is at least as good as an identical outcome, or any good is at least as good itself. This could be considered a somewhat trivial axiom, but worth adding as it establishes the direct proportionality and indirect directionality of preference ordering, that is necessary when optimizing towards equilibrium.

Revealed preferences: If it is revealed that outcome X is preferred to outcome Y, then it cannot be revealed that outcome Y is preferred to outcome X. In this case an individual will always choose to consume X over Y if both choices are present on the budget line m so that $p_1x_1 + p_2x_2 = m\ and\ p_1y_1 + p_2y_2 = m$, unless relative prices change. Formally this can be expressed as $px(p', m') \leq m \wedge x(p', m') \neq x(p, m) \Rightarrow p'x(p, m) > m'$.

Given these axioms, consumer behaviour can be expressed through indifference curves that illustrate the concept of consumer equilibrium.

In addition to the four axioms listed, two more must be discussed as they, together with *completeness* and *transitivity*, are essential for the expected utility hypothesis. These two axioms are *independence* and *continuity*. Both *completeness* and *transitivity* deal with preference under certainty, whereas the next two axioms deal with preferences under risk/uncertainty.

Independence: Let $p \in [0,1]$ and X, Y and Z be outcomes or probability distributions over outcomes, so that a weak preference over X to Y occurs if and only if $pX + (1-p)Z$ is weakly preferred to $pY + (1-p)Z$. This means that preferences over outcomes cannot be influenced by factors that are not relevant to the initial preference order. Suppose one is faced with the choice of two lotteries, $(p)[q(V) + (1-q)(U)] + (1-p)(Z)$ and $(p)(Y) + (1-p)(Z)$. Independence says that, if the first lottery is preferred over the second, then what determines this preference cannot have anything to do with the commonalities (p) and $(1-p)(Z)$. The choice is thus solely a choice between $q(V)+(1-q)(U)$ and Y. Embedding a direct lottery in a compound lottery, without changing the initial conditions, should therefore have no relevance in the valuation of the respective lotteries.

Continuity: Using the example from the independence axiom, continuity can be expressed as an indifference qualifier. If one prefers X to Y to Z, then there must exist a particular $p \in [0,1]$ such that one is indifferent between the lottery $pX+(1-p)Z$ and the certainty of Y. Continuity thus states that a unique probability distribution can always be found, such that one is indifferent between the probability of receiving the most preferable outcome, plus the probability of receiving the least preferable outcome, and the certainty of receiving the middle outcome (Levin, 2006).

2.2. Differences in Expected Utility

It might be pertinent here to interject the difference between risk and uncertainty, since the expected utility hypothesis can be formulated with respect to both objective risk and subjective uncertainty. Risk is here defined as an arbitrary set of outcomes of which the probability distribution is known with certainty, often described as lotteries. As such, these outcomes can be represented as $X=\{x_1,x_2,\dots,x_n\}$, and their relatively risky prospects as $(x_1,p_1;x_2,p_2;\dots;x_n,p_n)$, where each of the probabilities of the outcomes x_i, is determined by their relative probabilities p_i. Denoting a utility function as u, representing tastes or preferences, expected utility under risk can be expressed quite straightforwardly as

$$\sum_{i=1}^{n} u(x_i)p_i$$

Equation 1

Under uncertainty, expected utility has to take into account that the probability distributions are unknown to the agent, and therefor depends on subjective evaluations, or more precisely, the agent's beliefs about the variable states of the world that the agent occupies. Based on these beliefs, the agent must choose what acts to undertake, in order to maximize utility. Formally this can be expressed as

$$\sum_{i=1}^{n} u(x_i)\pi\left(f^{-1}(x_i)\right)$$

Equation 2

Where π represents the subjective beliefs about the likelihood of ending up in any particular state, and the invers function of x_i can be thought of as the corresponding course of action (Karni, 2014).

For both the expected utility hypothesis under risk and uncertainty, *completeness, transitivity, independence* and *continuity* are fundamental axioms that outline the behavioural assumptions necessary for the hypothesis to work. Though the axioms can be thought of with either a positive or normative predicate attached, depending on one's beliefs, the axioms are essential none the less. For now, let us stick to expected utility under risk as laid out by von Neumann and Morgenstern in their 1947 book, *Theory of Games and Economic Behaviour*, and later expand the concept to include uncertainty.

2.3. Expected Utility under Risk

The first two axioms, *completeness* and *transitivity*, can be thought of as behavioural assumptions, while *transitivity* and *continuity* are mathematical assumptions. For the same reason, the axiom *continuity* is often labelled the *Archimedean axiom*, because of its relation to Euclidean geometry. If an individual's preferences follow the four axioms, then utility can be represented numerically, as long as the order of the utility set is preserved.

The reason for this is that expected utilities are identical to positive affine transformations if, and only if, the necessary axioms are not violated. We can see this by writing $au + b$, where $a > 0$, u is utility, and b is any real number. If $u_1 > u_2$, then it would also be the case that $au_1 + b > au_2 + b$ as long as a is positive. Being able to represent utility numerically then places utility within a probability space, or σ-algebra, containing the collection of events that can be assigned probabilities as well as associated conditional probabilities.

The σ-algebra on a set of outcomes or states S, is a collection of subsets X of S such that $\emptyset, S \in X$. If $x \in X$, then the compliment of the subsets of x in X, denoted by x^c, will also be elements of X, such that $x^c \in X$.

In addition to this, if $x_i \in X$ for $i \in \mathbb{N}$, ($\mathbb{N}$ being all natural numbers), then $\bigcup_{i=1}^{\infty} x_i \in X$, $\bigcap_{i=1}^{\infty} x_i \in X$ (Neely, 2012).

This means that a collection of subsets is σ-algebra, if the subsets contain the empty set $\emptyset$, and is closed when taking the compliments, countable unions and countable intersections. It is important to note that in order for something to be measurable, the subsets has to be σ-algebra, since the relative probabilities described by unions and intersects, must be closed sets in order for them to be properly defined.

If utility is subject to positive affine transformation, then utility can be expressed in accordance with measure-theory, given that any real number exist within the space $\mathbb{R}^n$. Concordantly, a real value function for u must be present, depending on the outcomes $\{x_1, x_2, \ldots, x_n\}$, and their relative probabilities $(x_1, p_1; x_2, p_2; \ldots; x_n, p_n)$. A number within $\mathbb{R}^n$ can then be thought of as a point within a space defined by set-theory, and expressed as a particular state or outcome belonging to a set of states or outcomes that is a superset of various sets of acts. We can then say that a specific act is an element of a set of outcomes, or simple $f \in C$, and sets of acts are subsets of outcomes, or $F \subset C$ outcome.

The function that map a domain to a codomain, $f: X \rightarrow Y$ if $A \subseteq X$, is the image of the area in Y in which the results of A resides. Similarly, the union and/or intersect of different sets in the same domain, is an image that maps the function of the union and/or intersect in the codomain. Although intersections are not always preserved when mapping a domain to a codomain, the invers image of the function is however preserved (Ling, 2006). Later it will be shown why inverse images are very important for the field of subjective expected utility, but for now, let us continue with expected utility under risk.

In expected utility under risk, we can define a choice set $\mathbb{C}$ that is a convex subset of linear space $\mathbb{R}^n$. If X is an arbitrary set of outcomes, and $\Delta(X)$ is the Laplace operator on X, then elements of $\Delta(X)$ can be referred to as lotteries. Defining two lotteries a and b with $p \in [0,1]$, then the convex combination $pa + (1-p)b \in \Delta(X)$ will be given by $(pa + (1-p)b)(x) = pa(x) + (1-p)b(x)$ for all $x \in X$. Here $\Delta(X)$ is a convex subset of linear space $\mathbb{R}^n$.

Likewise, a set of probability measures (x, X) can be denoted by (x) , and two elements within $M(x)$ can be denoted A and B, where $p \in [0,1]$. A convex combination $pA + (1-p)B \in M(x)$, will then be given by $(pA + (1-p)B)(Y) = pA(Y) + (1-p)B(Y)$ for all Y in the σ-algebra (x, X), and $M(x)$ will be a convex subset of $\mathbb{R}$ on (x, X) (Karni, 2014).

The point of delineating the mathematical foundation of expected utility, is to show that utility, u, has to be a unique point within the choice set $\mathbb{C}$, and by extension, has to exist in the linear space $\mathbb{R}^n$ where different utilities increase in accordance with a real value function of u along a finite dimensional linear space. The *independence* axiom explicitly states that this must be so, in order for expected utility to be expressed as a mathematical measure.

Furthermore, the intersects and unions between different sets, specify that utilities with higher probabilities must be further along linear space than utilities with lower probabilities, all else equal. However, the convexity of the choice set also specifies the importance of indifference, since equalities between choices can be found all along the convex set. These equalities can be thought of as mirror images of each other, where for all elements of a set, a specific probability combination can always be found that equates the utility of risky elements with the utility of non-risky elements. The *continuity* axiom explicitly states that this must be so, not as a behavioural assumption, but as a mathematical necessity.

The $\geq$ symbol implies a binary relationship on the choice set $\mathbb{C}$, and since this choice set can be expressed as $\Delta(X)$, the *independence* and *continuity* axioms will be satisfied with any binary relationship on $\mathbb{C}$, if a real value function u on X exists. This can be expressed as:

$$p \ \mapsto \sum_{x \in \{x \in X \mid p(x) > 0\}} u(x)p(x)$$

Equation 3

Such that a choice criterion could be:

$$\sum_{i=1}^{n} u(x_i)p_i \geq \sum_{i=1}^{n} u(x_i)p_i$$

Equation 4

(Karni,2014)

2.4. Subjective Utility

In 1954, the mathematician Leonard Jimmy Savage proposed an alternative to expected utility, based on a desire to extend the concept of Bayesian priors. In Savage's framework, decisions can be expressed mathematically when defining a set S for the states of nature, and a set of consequences C, where subsets of S are events. Mapping the states of nature to the consequences is done by describing a choice set, denoted F, and defining the elements within F as acts or actions that facilitate the set of consequences C. We can therefore write that $c(f,s) \in C$.

In order for this to work, is must be permitted that an individual's evaluation of consequences be represented numerically in the form of utility functions. Likewise, people's degree of beliefs in the likelihood of events, also has to be finitely additive and able to be represented numerically as well. In the model, there is a strict assumption of independence between consequences and events, stating that events that help to shape the perceive state of nature, does not influence the perceived utility of the consequences. Nor does the probabilities of the events influence the acts by which the consequences are achieved.

If an event E has acts f and g, then $f_E g$ is the act that follow from $(f_E g)(s) = f(s)$ when $s \in E$ and $(f_E g)(s) = g(s)$. If an event is null, meaning that the event has no occurrence in the probability space, then $f_E g \sim f'_E g$ for all acts f and f'. Non-null events are events that contain acts that assign the same consequences to all events, and as such can be thought of as constant acts. Yet some constant acts have to be acts of which the individual is not indifferent so that $f > f'$, which is why acts have to be independent of the probabilities of the events if both null and non-null events exist. Unfortunately there is no discernible way to determine whether or not a null-event is one that the agent deems impossible, or one in which the agent is indifferent between the outcomes of the events.

Although there are some slight differences in the axiomatic foundations between expected utility under risk, and subjective utility as proposed by Savage, for the sake of simplicity, we will not go into detail with these differences. For all intents and purposes the differences can be ignored without loss of understanding. There is however one axiom that requires some elucidation and this is *small-event continuity* (Karni, 2014) .

Small-event continuity: If a probability measure is non-atomic, meaning that the probability measure has an infinite number of distinct values such that $\mu(E) = \mu(e_1) > \mu(e_2) > \mu(e_3) > \mu(e_4) > \cdots > 0$, then π can represent this non-atomic measure. If for every event E in the state set, and a number $0 < \alpha < 1$, then $e \subset E$ will be an event such that $\pi(e) = \alpha\pi(E)$. This means that there are an infinite amount of world states, and that these states can have probability measures attached, representing an individual's beliefs about the different world states that are finitely additive. For every event E we can formally express this as $0 \leq \pi(E) \leq 1$, where $1 = \pi(S)$, and for every two disjoint events e and E, it can be expressed as $\pi(E \cup e) = \pi(E) + \pi(e)$.

Savage's theory of subjective utility states that if a non-atomic, finitely additive, unique probability measure π on S exist, together with a bounded, real-value function for u on C, then preference relations can be expressed as:

$$f \;\mapsto \int_S u\big(f(s)\big)d\pi(s)$$

Equation 4

Defining states s as x_i and taking the integrated form then gives us:

$$\sum_{i=1}^{n} u(x_i)\pi\big(f^{-1}(x_i)\big)$$

where f^{-1} is the inverse image of the function of x_i (Karni,2014).

By taking the invers image of the choice function that maps s to u, the conditions that result in a specific utility of consequence is preserved through f^{-1} instead of potentially "disappearing" into the codomain.

In concluding the introduction to the expected utility hypothesis, it will be pertinent to likewise introduce some of the problems that will be discussed more in depth later on. First of all, utility and risk is given ex ante by $p \in [0,1]$, that is, prior to any action initiation, and not discovered through action as would be more correct. This means that the configuration of the world is known to all agents, giving them the possibility of maximizing utility by choosing the correct policy in accordance with their preferences. This does not change if we add the subjective component π into the mix, as in equation 2, since all this does is to allow each agent to disagree on the correct configuration of the world, of which there are an infinite amount in accordance with small-event continuity, all the while still holding that whatever constitutes the subjective beliefs of the agent is given, and therefore correct from his or her perspective. Again, there is no discovery process that, through action, reveals hidden states of the world, allowing the agent to update beliefs and adapt action during the process of trying to maximize utility. As such, there is no temporal dimension connecting expectation with utility as must be the case when describing a process that works over time, from expectation to action to discovery to adaptation to reward. This missing temporal dimension is a direct consequence of the axioms explained previously, where a static mathematical formulation forces the economist to adopt a world view in which mutually exclusive actions are undertaken simultaneously as opposed to having a discrete and organic cause and effect process.

There is an argument to be made, that expected utility does indeed operate with a discrete time conception by choosing an action in step one, that maximizes the probability of ending up in a preferred world state in step two. There is however no real cause and effect structure in operation, since both actions and world states are parameterized and given to the agent by the axioms. We can define a set S for the states of nature, and we can also define a subset of S as events. We can even define a set of consequences and denote them as C. We cannot however say that $c(f, s) \in C$ and expect this to be a meaningful statement. All we are describing here are the elements of a given world state, quite arbitrarily it might be added, yet what we want to describe is the dynamic process that yields a specific world state as experienced by acting individuals, something of which expected utility has very little to say. When we define elements in a set, we are saying that the function of a particular element in the set produces another element within the codomain of C. This is the image of the function. The domain of the function therefore has to be uniquely defined along with the image, as seen when we take the inverse of the image and arrive back at the domain. Action and consequence, however it is defined, will therefore always be mutually determined by each other like two sides of the same coin, which helps to exemplify why the system must always be in equilibrium. With no stable real world relationship between human action and consequence, this mathematical conception of action is simply fallacious.

While the expected utility hypothesis can provide the economist with a useful hypothetical construct for analysing the direction in which the system is tending, it cannot in any way be used to describe actual economic behaviour or phenomenon, as these are products of complex agents operating within a complex system.

3. The Bayesian Brain

In the following section, Bayesian probability will be introduced as part of the building blocks of the thesis. This will be followed by the Bayesian brain hypothesis and the free energy principle. It should be mentioned that standard expected utility theory, as it applies to subjective utility in particular, also employ Bayesian methods as a way to update agent behaviour when confronted with changing world states as seen in equation 3. Often this will be in response to changes in relative risk ratios or to price changes, since a price can be conceived as a blanket term that includes practically every characteristic connected to a given good. This is not particularly helpful however. In the expected utility hypothesis we know that agents have priors, that is, we know that they have preferences. We do not know what these preferences are, but once they have been stated by the agent, they will have been revealed to us. Now these preference cannot change unless relative risk ratios or relative prices change, with due consideration for the price elasticity of demand. Since we cannot specify a point in which preferences have been revealed to us, preferences must be given ex ante, rendering the entire point of having agents with priors moot. If preferences are given, so will the rules be given that govern these preferences. Agents are therefore not updating priors by collecting evidence as much as they are blindly following rules in which the entire Bayesian updating and learning procedure can be ignored without much loss of clarity. This does not mean that these rules cannot be helpful for the economist, in fact, many times they are. But they never the less conceal the true nature of human agency, and may very well lead us to conclude that prices govern action, and not the other way around. This is not the same as saying that agents do not respond to incentives, but a way of saying that this really isn’t all that interesting, if responding to incentives is the same as reacting to stimuli, or as in this case, reacting to a price change. If this is all there is to it, then we could just as easily say that houseplants respond to incentives, but a collection of houseplants does not an economy make.

3.1. Bayesian Probability

According to Bayes' rule, joint probabilities can be expressed as $P(A,B) = P(A \mid B)P(B) = P(B \mid A)P(A)$. If we then wanted to test a hypothesis H in the light of the available evidence E, then the result will be given by the formula $P(H \mid E) = \frac{P(E \mid H)P(H)}{P(E)}$.

The first term in the equation, $P(E \mid H)$, is the likelihood function. If a hypothesis were believed to be true, the likelihood function would then express what the evidence for the hypothesis would look like, meaning the probability of the observed evidence, given the hypothesis.

The second term $P(H)$, is the prior and refers to the level of knowledge that is present before the evidence is evaluated. It is in this term that the subjectivity of the agent is present, and likewise the term that has caused the most problems for mathematicians, statisticians, economists and neurologists alike. It is indeed no wonder that the prior poses such a challenge for researchers and scientists, since it effectively has the capacity to encompass anything, everything and nothing, depending on the specific knowledge any specific agent is believed or required to have.

$P(E)$ is often referred to as the marginal likelihood or the model evidence. This term will be the same for all conceivable hypotheses, and can thus be thought of as a constant.

The $P(H \mid E)$ term is the posterior, and denotes the probability of the hypothesis after the evidence has been taken into consideration. Bayesian probability theory thus introduces the concept of learning into probability theory, since the theorem reflects a process by which individuals can update their beliefs based on available evidence.

Bayesian probability has numerous applications and can be expressed in many forms. The form already presented is called the simple form. The extended form goes as follows.

Using the law of total probability, $P(B) = \sum_i P(B \mid A_i)P(A_i)$, we can write Bayesian probability as $P(A_i \mid B) = \frac{P(B \mid A_i)P(A_i)}{\sum_i P(B \mid A_i)P(A_i)}$, where A_i is a partition of the sample space.

For binary variables, where the estimation results in a dichotomous statement, the equation can be written as $P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B \mid A)P(A)+P(B \mid \neg A)P(\neg A)}$, where the term $\neg A$ is the compliment of A.

For Bayesian probability with random variables the form can be written as $(x \mid Y = y) = \frac{P(Y=y \mid X=x)fX(x)}{P(Y=y)}$, where X is continuous and Y is discrete.

When Y is continuous and X is discrete the form will be $P(X = x \mid Y = y) = \frac{fY(Y \mid X=x)P(X=x)}{fY(y)}$, and $fX(x \mid Y = y) = \frac{P(Y=y \mid X=x)fX(x)}{P(Y=y)}$ if both X and Y is continuous.

Alternatively, Bayesian probability with random variables can be written as:

$$fY(y) \int_{-\infty}^{\infty} fY(y \mid X = S) fX(s)\, dS$$

Equation 5

Where S is some random set. In this version, the denominator has been eliminated due to the law of total probability. Consequently the function of $Y(y)$ becomes an integral, and somewhat similar to the form expressed in both expected utility under risk, and subjective utility.

A smart thing about the Bayesian theorem is its very intuitive applicability to set theory, and by extension, the notion of conditional probability. The probable validity of a hypothesis given some data is the same as the union between the hypothesis and some data, taking due account for the weight of the data. Formally this can be expressed as $P(H \mid E) = \frac{P(H \cup E)}{P(E)}$, provided that $P(E)$ is not zero (Olshausen, 2004).

The strengths of Bayesian probability, purely on a technical level, are due to the notion of conditional probability, meaning that the probability of some outcome is always conditional on other underlying factors. For instance, suppose you display some symptoms that could be associated with a serious illness. The question then becomes; what is the probability of actually having this illness given a test that can be performed with some measure of accuracy, say 95 % with a 4 % chance of receiving a false positive. If you take the test and it turns up positive, many individuals would then be inclined to conclude that there now is a 95 % probability of having the illness. This would however be incorrect. By taking into account priors and likelihoods, the probabilities often change dramatically. If for instance the base rate of the illness in the wider public is a mere 1 %, then the real probability of being sick given a positive test result will only be (0,95*0,01)/((0,95*0,01)+(0,04*0,99)=0,19, meaning 19 %.

Moreover, Bayesian statistics introduces the concept of updating beliefs, which implies learning, a central defining quality of biological systems that acts based on perception, that is, information.

In psychology and neurology there has typically been two ways to conceptualize the nature of perception and object identification. The first way is the bottom up approach first popularized in 1966 by J.J. Gibson. This approach suggests that perception is achieved exclusively from sensory inputs, meaning that the environment is reflected as it is via different wave lengths of light hitting the retina and nothing more (Gibson, 1978). Inspired by Helmholtz, R. Gregory proposed the top down approach in 1970, suggesting that perception primarily is inferred, and subsequently re-created from memory based on past experiences and exposures (Gregory, 1997), but Because both approaches has substantial strengths and weaknesses on their own terms, more recent models based on Bayesian inference and machine learning, employ both a top-down and a bottom up process.

One of the problems with the bottom-up approach is that it fails to explain the prevalence of optical illusions, but more importantly, the theory states that there is enough information present in the environment in order for the brain to be able to impute the structure of the environment without having to infer missing inputs (Gibson, 1978). Strictly speaking this is also correct, but the problem is that there is too much information present in the environment and with a pure bottom-up approach, there is no way of filtering and discriminating between relevant and irrelevant information. As it turns out, it actually becomes necessary to choose what information to focus on before the information can be consciously registered, as partially demonstrated in the popular invisible gorilla experiment (Simons & Chabris, 1999). In order to filter information we therefore need a module that pre-selects direction, attention and categories of interest in accordance with a top-down approach. However, a pure top-down process cannot learn and cannot explain how infants perceive, considering that they have no memories or experiences to infer from.

3.2. The Bayesian Brain Hypothesis

The link between Bayesian probability and action is the idea that the brain assigns probabilities in a very similar manner to real world scenarios by collecting evidence that combines to form a likelihood density, which then has to be compared with a prior density in order to yield a posterior density (Chater, Tenenbaum & Yulle, 2006; Körding & Wolpert, 2006; (Knill & Pouget, 2004; Griffiths & Tennenbaum, 2006). This however can be viewed in two different ways, the methodological and the theoretical. The theoretical way views the assigned probability distributions as actual physiological properties of the brain, where individual and clusters of neurons carry numerical information about the likelihood of any particular hypothesis. From this perspective, individuals are Bayes optimal in all aspects of life ranging from all cognitive, motor and perceptual functions, and experimental results that might suggest otherwise is often disregarded due to poorly defined priors (Bowers & Davis, 2012). This however, does not mean that individuals cannot err, but that errors in judgment or inference is solely a property of the probability distributions between two or more likely outcomes.

Suppose two distributions, A and B, are Gaussian, then the decision boundary will be the midpoint between the two means of the two distributions in the absence of priors. In distributions that have the same variance, the point where A and B intersect with each other (figure 1) will also be the decision boundary. In this simple example the distributions are known to produce a 50 per cent chance of either outcome. The area beneath the intersection is the union between A and B, and is due to sensory noise. Exactly where the decision falls between A and B, is determined by the prior so that the decision boundary shifts inside the union to reflect the value of the prior term. When Individuals thus err, it has to be consistent with a statistical error term that is subject to the prior knowledge and subjective evaluations of the individual.

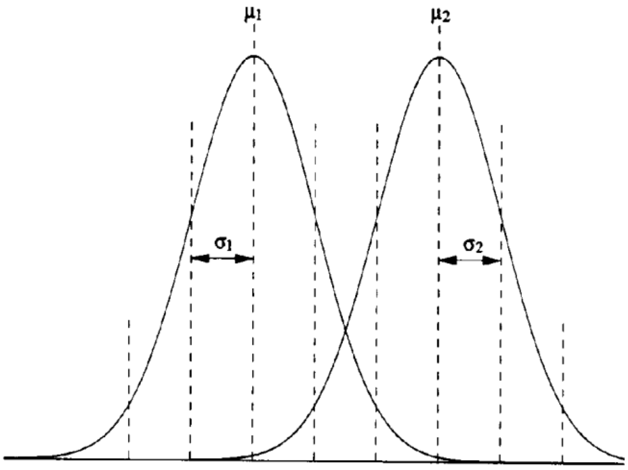

Fig. 1

In figure 2 are three examples of how priors determine the posterior given the likelihood function. Scenario A depicts a situation where the prior distribution shifts the posterior away from the likelihood function, but because of the great divide between beliefs and evidence, the posterior probability density is lower than the likelihood density, denoting a higher degree of uncertainty had the prior and posterior been closer together. In scenario B the prior and the likelihood are indeed closer, resulting in a higher probability density of the posterior, and a lower degree of uncertainty. However, given the overlap between the likelihood and the prior, there would be a great deal of ambiguity in the final decision, but since the probability density is greater in the likelihood, the decision boundary would be biased towards the mean of the likelihood function. Scenario C depicts a situation where no priors exist. Here the posterior and the likelihood would be an identity.

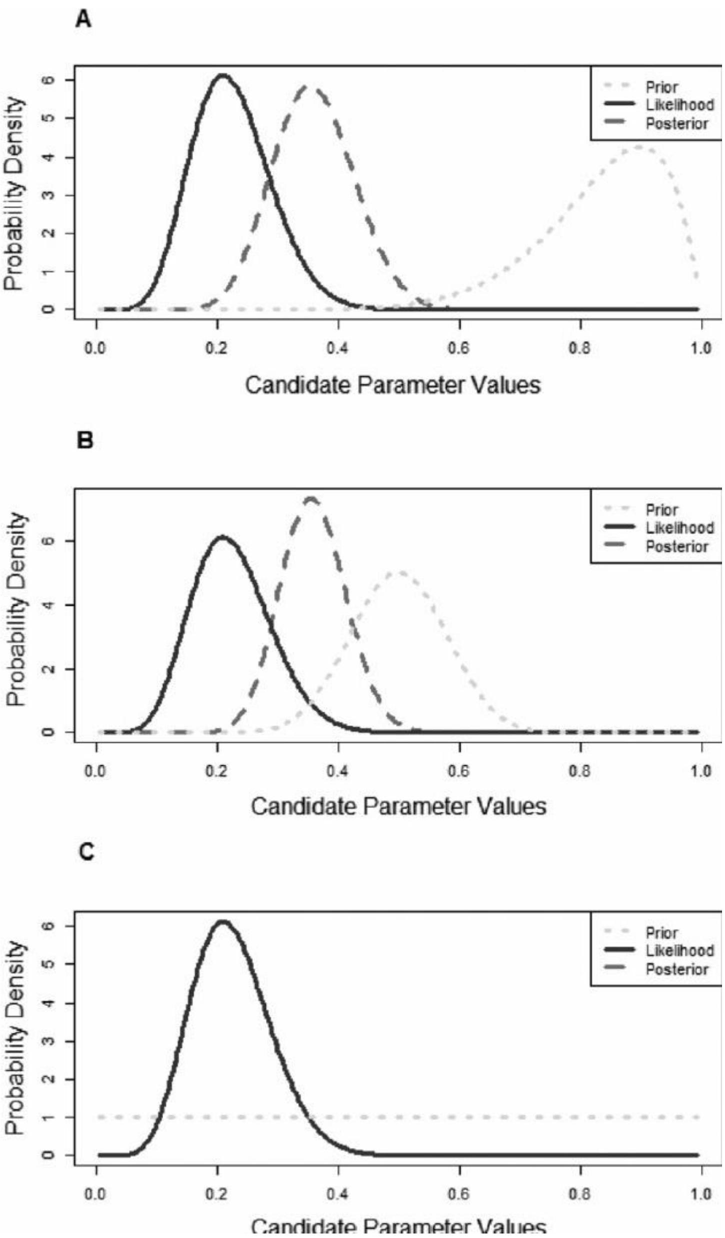

Fig. 2

The methodological approach simply asserts that the brain will act as if trying to maximize Bayesian model evidence, meaning that the brain is engaging in, what for all intents and purposes, is approximately equal to Bayesian inference. In this approach there is little to no consideration as to how the brain accomplishes Bayes optimality, only that Bayesian statistics is currently the best tool at our disposal for describing how the brain might operate (Geisler & Ringach, 2009). However, this difference in views does not influence the fundamental implication of the Bayesian brain, and produces the same results regardless of approach.

Far from simply being a mathematical construct of optimal behaviour, Bayesian decision theory might have an actual neurological basis according to many researchers (Chater, Tenenbaum & Yulle, 2006; Körding & Wolpert, 2006; Knill & Pouget, 2004; Griffiths & Tennenbaum, 2006; Jacobs, 1999; Knill & Saunders, 2003; van Beers et al, 1999; Battaglia et al, 2003; Burr & Alais, 2005). One major problem that has confounded psychologists, neurologists and philosophers alike, for many years, is the problem of perception and information integration. How to look at the world and extract meaning from it, in order to act in a manner that does not compromise ones further existence. This problem is not as straightforward as it might appear at first glance, for there is quite simply too much information present, at any given point it time, to effectively navigate the world, without some prefixed notion of what would constitute relevant and irrelevant information. But how do we know how to discriminate between relevant and irrelevant information? How is it possible for sensory inputs to convey relatively precise knowledge about any particular phenomenon, when there is so much information to contend with? Even in very simple situations, such as identifying the location of a sound, the manner in which the brain conceptually accomplishes this is fraught with white noise and huge margins of uncertainty.

Just to illustrate how complex this simple scenario really is, imagine being inside your head for a moment, and think of all the well-defined abstractions needed in order to solve the aforementioned problem. One thing is the physiological mechanisms by which sound reaches the auditory centre, quite another is the manner in which sensory queues get translated into understandable themes and concepts. First the brain needs an idea of self, something that separates it from all else. Then it needs an idea of what everything else is on multiple levels of resolution, so that concepts might make sense, and objects can get identified and distinguished from groups or clusters. The space between objects and concepts then has to be understood as denoting distance, where mapping this distance onto a three dimensional space creates the notion of direction.

In other words, you would need a working model of the world, a prior set of beliefs that map out an arena in which action and uncertainty resolution is possible. To detect the source of a sound, this arena must have some rules for conduct, so that it can be traversed and investigated for anomalies. It needs a set of hypothesis' about how this mental construct relates to its outside world, and another hypothesis about how sounds work and move, something that adds direction and focus to your investigation. Without this hypothesis, uncertainty resolution would be random, and instead of shifting your head towards the hypothesized direction, you might start clapping or jumping up and down. In fact you would need something very reminiscent of a likelihood function (Battaglia et al, 2003; Burr & Alais, 2005).

Testing your hypothesis by shifting your head, then gathers evidence, resolves some uncertainty and updates your model by integrating it into a framework expressed as a posterior, where the probabilities are assigned to the initial hypothesis so that it can be rejected or momentarily confirmed until new evidence presents itself (Körding & Wolpert, 2006; Friston, 2009).

It might be thought at this point, that assigning probabilities is an unnecessary step because of the binary outcome “reject” or “confirm”. But the state space where a potential solution exists has to be one of degree, if the point at which evidence searching or uncertainty resolving behaviour is to continue or conclude. This is due to the noise of the sensory inputs, and the fact that three dimensional information has to be decoded into a two dimensional environment. In far more complex examples than simple sound detection, it is furthermore evident that weights must be given to all potential solutions in order for one solution to be tentatively superior. Suppose that an additional piece of information presents itself in the form of a visual cue. If this information causes ambiguity as to the location of the sound, then the location has to be expressed as a degree of belief given the evidence (visual cue), and the prior (auditory cue) (Battaglia et al, 2003; Burr & Alais, 2005).

The bulk of the evidence for the Bayesian brain hypothesis is primarily derived from its mathematical implications. Using Bayesian algorithms, it has thus become far easier to simulate organic behaviour, a fact that has caused enormous progress in the fields of AI research and machine learning. There is however some neural imaging studies and behavioural experiments that likewise gives credence to the Bayesian hypothesis. Several studies have found that the uncertainty derived from various real world cues, vary according to one’s position relative to the cues. Cues are thus weighted differently on different trials as predicted by associated Bayesian computations (Jacobs, 1999; Knill &Saunders, 2003; Alais & Burr, 2005; van Beers et al, 1999; Battaglia et al, 2003). Bayesian algorithms have also been very successful in simulating the electrical pulses that transfers signals from one neuron to another. These electrical pulses, called neuronal spikes or spike trains, are the intermediary between sensory inputs and motor outputs, and are thus a crucial mechanism for understanding the brain. Since isolated spikes are all practically identical, the form of the spike does not carry any information. It is the sequence of spikes, number and timing, that code messages to be interpreted by the brain and these codes are the focal point of probabilistic neural models such as neural coding and neural decoding (Friston, 2010).

3.3. Critique of Bayesian Methods

Not all are equally enthusiastic about the Bayesian brain hypothesis. It goes without saying that the entire behavioural branch of economics was more or less conceived as a counter to the notion of optimal decisions, instead focussing on mental shortcuts know as heuristics in order to explain the decision making process (Kahneman et. al., 1982). Another critique was made by Jeffrey S. Bowers and Colin J. Davis in a paper from 2012 called "*Bayesian Just-So Stories in Psychology and Neuroscience".* In their paper, Bowers and Davis argue that Bayesian methods used to model both neurological processes and actual decisions, for the most part, are not falsifiable due to their flexibility. It will therefore always be possible to adjust priors, likelihoods and utility functions in order to make Bayesian theories account for any pattern of results (Bowers & Davis, 2012). As one of the examples of this, Bowers and Davis site a 2009 paper Bar-Eli, Azar and Lurie where they concluded that the alleged irrational behaviour of football players engaging in a penalty situation, was in fact not irrational if one changed their utility function. The players therefore did not want to maximize the probability of scoring a goal or saving the ball, but were instead trying to maximize some social utility function. This example strikes me as quite curious however. Bar-Eli et al argues that statistical models suggest that a football player would have a higher chance of scoring if he or she aimed at the top left or right corner of the goal. Instead, football players usually aim for the lower left or right corner, thus behaving irrationally (Bar-Eli et al, 2009). This line of reasoning makes little to no sense, and would be like arguing that statistical models suggest that individuals who receives the highest amount of points in a ring toss game, are the most successful ring toss players. Not being the most successful ring toss player is therefore irrational. The top left and right corner is the furthest distance from where the goal keeper usually stands, but also the most difficult place for the shooter to actually hit. This is because of the added probability of not even hitting the target to begin with, which is why the shooter will want to maximize this probability first. Going for maximum points in a ring toss game also means consistently aiming for the most difficult place to hit, and unless one is exceedingly skilled at tossing rings, this would be a terrible strategy because of the added probability of not getting any points at all.

Bowers and Davis also references a 1994 study from Oaksford and Chater in which they conclude that the Wason card sorting task actually proves Bayes optimal reasoning despite that fact that this test is consistently used to demonstrate how error prone human reasoning is. In Wason's card sorting task, individuals are shown four cards, an A, K, 2 and a 7. They are then asked to test the hypothesis "if there is an A on one side of a card, there is a 2 on the other" (Bowers & Davis, 2012). The logical thing to do is to flip over the 7 and/or the K card, but participants consistently select the 2 and/or the A card which is a logical fallacy called *confirming the consequent.* Oaksford and Charter showed that by using conditional probability, and what is known as the *rarity assumption*, then confirming the consequent could be the optimal solution. The rarity assumption says that any particular thing will always be rare compared to everything else, we have therefore developed a tendency to confirm or validate statements or the presence of objects, as opposed to using more logical disconfirmation or falsification methods. Arguing that choices in the card task are conditional on this rarity assumption then makes the assumed erroneous choices in Wason's card task Bayes optimal (Oaksford & Chater, 2001). If however the card test had shown that individuals consistently picked the most logical cards, Bowers and Davis argues, then this would also have been Bayes optimal (Bowers & Davis, 2012), meaning, heads I win, tails you lose, which is a fair assessment.

In accordance with rational choice, it is customary to state that individuals choose actions that maximize the probability of a preferred word state. If we then describe a preferred world state in accordance with a utility function, and state that actions chosen are conditional on priors, then it will be possible to make practically all patterns of behaviour conform to Bayesian methods, by adjusting the likelihood, in this case by changing the utility function, or by adjusting the prior, which is the conditional probability. But if we consider that individuals seek to optimize beliefs about world states by optimizing policies, instead of choosing single acts that maximizes a preferred world state, we could get a different perspective on both examples presented by Bowers and Davis.

In the two examples we actually get a good initial introduction to the strengths of the variational free energy principle compared to more traditional Bayesian methods. In the first example, Bar-Eli et. al. are constrained by the need to infer a motive to which a utility function can be attached. This is a consequence of trying to maximize the probability of a preferred world state by optimizing an action. If we however realized that individuals seek to optimize policies or action sequences, then each problem will have to be framed as a step by step optimization problem or a path integral. Before we can think about scoring a goal in a penalty situation, certain steps will have to be taken first, which will be part of the optimization problem. One of these steps will have to be trying to maximize the probability of hitting the target, before trying to maximize the probability of beating the goal keeper.

In the second example, we will see why optimizing beliefs, and not world states, is a better way of framing optimization problems. A hint to this can be found in a paper from 1992 by Leda Cosmides and John Tooby, where they found that people have a much easier time selecting the logically correct cards from Wason's card sorting task, when the hypothesis was reformulated to reflect social situations. Instead of an A, K, 2 and a 7, individuals were shown four cards with ""drinking beer", "drinking coke", "25 years old" and "16 years old", written on them. They were then asked to test the following hypothesis. "If a person is drinking beer, then he must be over 20 years old". Here test participants were picking the logically correct "drinking beer" and "16 years old" cards in a much more consistent manner (Cosmides & Tooby, 1992). Cosmides and Tooby argued that this was because we humans have had a long time developing in, and adapting to, a social context, and are therefore not well equipped to handle and understand overly abstract and academic problem formulations (Cosmides & Tooby, 1992).

Using the variational free principle, I will argue in a similar manner, but instead of adopting the social context explanation, simply suggest an action schema in lieu of priors or beliefs. This means that we have adapted to thinking in terms of actions that can be carried out and constantly seek to update and optimize some function representing a collection of action sequences that we can draw from when engaging the world, both directly and indirectly. This is basically what is meant by the process of optimizing beliefs about world states, by optimizing policies or action sequences.

We have now introduced the expected utility hypothesis, Bayesian statistics and the Bayesian brain hypothesis. This will lead us the variational free energy principle, followed by the principle operationalized, so as to generate an economic agent.

4. Variational Free Energy

The theory states that all organic systems are characterized by a common feature; the ability to combat the dispersal forces of entropy on the cellular level, and counterbalance entropy by exploiting energy from the surrounding environment. How organic systems do this is by minimizing the property known as variational free energy. An interesting aspect of this principle is its logical extension from a biological/thermodynamical principle, to a technical description of neurological processes. It is however necessary to point out, that variational free energy is not to be confused with thermodynamic free energy. Under the free energy principle, variational free energy is the upper bound on entropic surprise, the surprise element being unforeseeable or atypical states and events confronting the organism. This is known as the negative logarithm of model evidence in information theory, where the average surprise over time is entropy, that is, the law that all things move to a less and less ordered state. In order to minimize variational free energy, which can be seen as akin to maintaining a homeostatic environment, the brain must generate a probabilistic model of the environment and all the events typically encountered within it (Friston et al, 2006; Buckley et al, 2017).

This model is comprised of a recognition density that gives the likelihood of the probability of a sensory input given some environmental data, and a generative density that models what causes the sensory data ought to have. This will be akin to a prior describing the beliefs of the probability distribution of the environmental input. The Kullback-Leibler (KL) divergence between the recognition density and the generative density is a non-measurable quantity referred to as variational free energy, and by minimizing this quantity, the KL divergence approaches zero. Instead of having a model of the environment generated from the ground up by sensory input alone, the free energy principle suggests that the sensory inputs from the environment are inferred initially from beliefs, and subsequently matched with sensory input in accordance with a

Kalman-filtering process by active sampling, often referred to as active inference. This active inference generates prediction errors by matching predicted with actual signals, and through iterative sampling, i.e. iterative acts, the prediction errors diminish, minimizing the KL divergence along with variational free energy (Friston et al, 2006, 2009, 2010; Buckley et al, 2017).

In the theory of optimal control learning, the goal is to select an action that maximizes some value function such that the preferred state of the world would manifest given an action. This can take the form: $u_t^* = \arg max\, V\,(s_{t+1}|u_t) = \pi(s_t)$, where u is an action, s is a world state and π is an optimized action, an equation very much comparable to Bellman's optimality principle believed to be informed by Morgenstern and Neumann's book *Theory of Games and Economic Behaviour* from 1944 (Friston et al, 2010).

In Bellman's optimality principle, problems and objectives are more or less the same, meaning that problems are framed in terms of reaching an objective or a goal. This could be how to minimize costs or maximize profits or utility, and from this optimality principle, expected utility theory follows quite naturally.

In the free energy principle the goal is not to maximize the value of a world state, but to optimize beliefs about world states formally represented by: $u_t^* = \arg min\, F\,(Q(s_{t+1})|u_t)$ through subsequent actions $\pi^* = \arg\min \sum_\tau F(Q(s_\tau)|\pi)$ where $u_\tau = \pi(\tau)$. Here the first task will be to resolve uncertainty about the consequences of subsequent actions, which means that action cannot be a function of the states in the world, but must be a function of beliefs about states in the world. This generates an intensity measure, or more precisely an energy functional, which in this case is the free energy function F of a function Q which describes an approximate posterior distribution indicating beliefs given a policy π. It is important to note that the order in which actions are undertaken matters, which means that actions now can be framed in sequences of actions or policies. We are not just trying to optimize the next best action, but the best sequences of actions in line with a path integral, or a time average $\sum_\tau$, of an energy functional of beliefs about future world states s_τ. This is basically the same as invoking Himilton's principle of least action when framing policies in terms of good or bad behaviour as conceptualized by the economist, read, rational or irrational (Friston et al, 2012).

By using a Markovian decision process, behaviour can hereby be described by the following equation:

$$\ln P\,(\pi|\gamma) = -\gamma \cdot \mathbf{G}\colon \mathbf{G} = \sum\nolimits_{\tau} F(\pi,\tau)$$

Equation 6

Where $\ln P\,(\pi|\gamma)$ is the probability distribution of a policy π given a precision parameter or an inverse temperature of beliefs γ that is inversely proportional to expected free energy $\mathbf{G}$ equalling the average free energy associated with a policy over time. What we are saying here is simply that the selection of a policy has a cost, and that this policy cost is equal to expected free energy. With this equation we can even go as far as stating that the mere act of thinking about acting is costly, something that is reflected in behavioural economics and likewise one of the central arguments in Daniel Kahneman's book from 2011 "*Thinking Fast and Slow*" (Kahneman, 2011).

How good a policy turns out to be, is a negative free energy described as

$$-F(\pi,\tau) = E_{Q(O_\tau,S_\tau|\pi)}[\ln P\,(O_\tau,S_\tau|\pi)] + H[Q(S_\tau|\pi)]$$

Equation 7

where $[\ln P\,(O_\tau,S_\tau|\pi)]$ is an energy term that says that hidden states in the world S_τ, cause observable outcomes O_t given a particular policy π, and P is the probability distribution of outcomes that are generated by hidden states. The $H[Q(S_\tau|\pi)]$ term is an entropy term where Q is an approximate posterior distribution indicating beliefs, and H is a vector encoding entropy or ambiguity over outcomes for each hidden state.

$$= \; E_{Q(O_\tau,S_\tau|\pi)}[\ln Q\,(S_t|O_\tau,\pi) + \ln P\,(O_\tau|m) - \ln Q\,(S_\tau|\pi)]$$

Equation 8

If we remove any uncertainty or ambiguity with regards to observations $\ln Q\,(S_t|O_\tau,\pi)$, then what is left will be the KL divergence, here indicating risk sensitivity $E_{Q(O_\tau,S_\tau|\pi)}[\ln P\,(O_\tau|m) - \ln Q\,(S_\tau|\pi)]$. Without ambiguity, this term says that states are no longer hidden, which means that states and observations are the same. Here the KL divergence scores the difference between what we believe will happen given a particular policy $\ln Q\,(S_\tau|\pi)$, and what we want to have happen $\ln P\,(O_\tau|m)$ given that we are autonomous individuals m, which can also be described as our prior preferences about long term outcomes. This difference is therefore the same as a score for the objective risk we are willing to accept.

This also means that if we remove risk $\ln Q\ (S_\tau|\pi)$, what we are left with is a term describing subjective utility $E_{Q(O_\tau, S_\tau|\pi)}[\ln P\ (O_\tau|m)]$, where we expect our actions to maximize the probability of ending up in a preferred world state based exclusively on our prior beliefs.

$$= E_{Q(O_\tau|\pi)}[\ln P\ (O_\tau|m)] + E_{Q(O_\tau|\pi)}\left[D_{KL}\left[Q(S_\tau|O_\tau, \pi)\ \middle|\ Q(S_\tau|\pi)\right]\right]$$

Equation 9

The last term $E_{Q(O_\tau|\pi)}\left[D_{KL}\left[Q(S_\tau|O_\tau, \pi)\ \middle|\ Q(S_\tau|\pi)\right]\right]$ describes an information gain that resolves uncertainty by scoring the degree to which uncertainty is reduced by pursuing a given policy.

To get the expected free energy we simply remove the entropy vector H in equation 7 and take the average over time.

$$\mathbf{G} = \sum_\tau E_{Q(O_\tau, S_\tau|\pi)}[\ln P(O_\tau, S_\tau|\pi) - \ln Q\ (S_\tau|\pi)]$$

Equation 10

(Parr & Friston, 2018; Friston et.al., 2017, 2016)

If we then suggest that $\mathbf{G}$ can be viewed as an action schema from which any particular action sequence expected to work can be drawn, and call this expected action sequence g, we can then start to model an agent based on the variational free energy principle.

4.1. The Free Energy Agent

Acting in time necessarily introduces the problem of acting without reference to a tangible reward, which as a function of time, must be abstract or imaginary in the present. Expectations can therefore be seen as being analogous to strategies, where expectations born from experience places a reward in continuation of a set of specific acts (Covey & Cheer, 2019; Schultz, 2016; Coddington & Dudman, 2017), effectively bypassing the temporal "void", and equating the reward with the initial act. How to escape a mere subsistence level of existence is then predicated on the ability to extract out a unifying time structure from action, a structure that effectively places more remote ends into the immediate consideration. Here, time itself exists as an abstract linear function of a measure of sequences of actions expected to facilitate the transformation of means into ends. In other words, action sequences and patterns amalgamate trough neural interconnectedness into a linear time structure, the active perception of which indicates a similar perception of cause and effect structures. This evolves initially through observations of the changing environment, where motion is registered relative to static, and substantiates time perception in relation to an increasing individual action schema. The schema is fundamentally analogous to a system of beliefs about the structure of the world, and what actions to take in order to influence this structure so that means and ends become linked. This suggestion is merely an extension of what is already known about how the brain "counts" in service of motor coordination for instance, and to this regard, it formally invokes the principle of complexity where simple initial conditions results in highly complex and unpredictable patterns through continuous iteration (Turner & Baker, 2019; Friston et al, 2016).

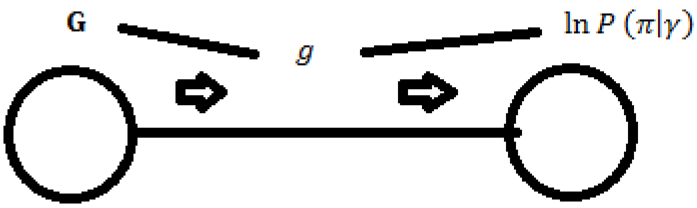

Figure 3: An expected action sequence g is drawn from the action schema **G,** in order to propose an initial solution that reduces variational free energy in accordance with $\ln P\ (\pi|\gamma)$, and linking rewards with initial acts.

Equating distant rewards with initial acts means that all rewards, however distant, has a present value, the value of which is dependent upon a rough measure of the acts needed in order to procure the reward. This action sequence is interchangeable with a time sequence, subjectively experienced as a strain on the mental faculties actively engaged in a simulation of the sequences of events that has to transpire, in order for the reward to finally be obtained. The higher the cognitive strain, the lower the present value of more and more remote rewards all else equal and conversely, the higher the discount rate on the future, even though this future is technically nothing but an abstract measure of effort given an end.

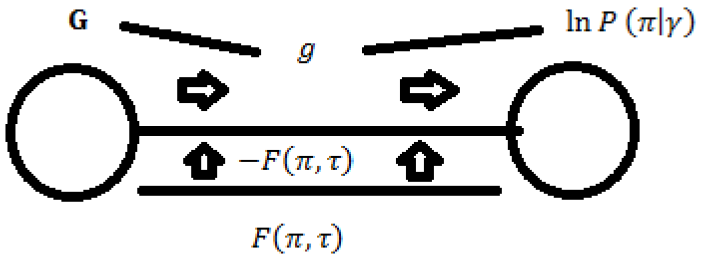

Figure 4: The expected action sequence is matched with observations and modified based on positive or negative free energy, that is, observations that confirm or reject predictions about states in the world given a policy.

Interestingly enough, this means that individual time perception must vary in accordance with the particular action sequence selected from the wider action schema, and again when trying to resolve uncertainty. If the action sequence of the individual changes over time, or from moment to moment in response to the changing environment, so to must the individuals' perception of the flow of time change, influencing the intensity of felt wants and the urgency with which actions should be undertaken towards desired ends.

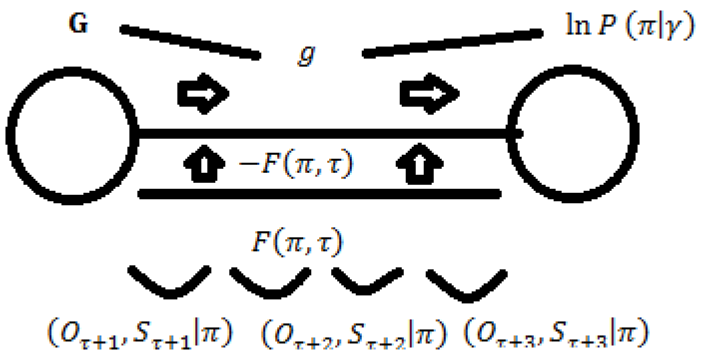

Figure 5: Along the action sequence, observations based on an initial policy reveal hidden states in a step by step process. This updates the action sequence, and the agent can take appropriate measures in order to either update beliefs, or manipulate states so as to conform to beliefs hereby minimizing free energy. The process is analogous to a policy that seeks to resolve uncertainty.

Any action is initiated by a call for alleviation or a rectification of a psychological or biological state. The action schema then posits a set of causally related action sequences or policies expected to work, weighted by the strength and the speed of the involved neuronal-connections. The most often used successful strategies in the past, conveys the highest probability for success, resulting in a strongly weighted expectation following the principle of least resistance or

Hamilton's principle of least action. The first step in any action sequence will be to search out the information that is predicted to indicate an appropriate policy. Based on the available information in the environment, each sequential step towards the desired end state is then motivated by a sub-strategy that seeks to align predicted with actual events mediated by the dopaminergic system. By this measure, each step in the action sequence is trying to resolve uncertainty about the next course of action by revealing hidden states in the world, drawing upon the wider action schema as an adaptive strategy when perfect alignment of predictions and events are unsuccessful. The perceived time flow is the felt strain of the action sequence from beginning to end, plus the added strain of necessary adaptive behaviour. The perceived time flow can therefore fluctuate in response to this adaptation, influencing intensity, urgency and the potential point of resignation along the action sequence, in so far as the action sequences has been frustrated by the necessity for adaptive actions.

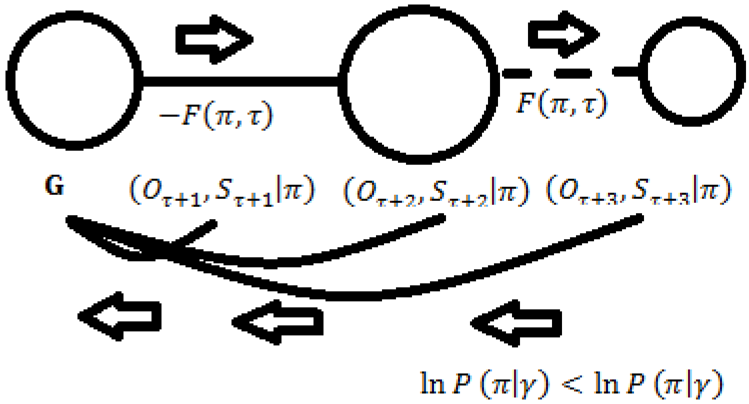

Figure 6: Each new observation reveals hidden states that do not conform to initial beliefs, i.e. the expected action sequence. This surprise increases free energy. At each revealed state, a new sequence is drawn from the action schema in order to reduce free energy, but the process influences the expected value of the end goal as seen in the $\ln P\,(\pi|\gamma) \; < \; \ln P\,(\pi|\gamma)$ term. Recall equation 6 $\ln P\,(\pi|\gamma) = -\gamma \cdot \mathbf{G}$: $\mathbf{G} = \sum_\tau F(\pi,\tau)$. Basically we are saying that the energy in the system equals the energy in the system at all times. If we increase energy in one part of the system, we must decrease it somewhere else. When moving to decrease free energy at on step in the sequence, this adaptation therefore increases free energy in the entire sequence, influencing the efficacy of the sequence and the value of the end goal. This is why beliefs/expectations/predictions have to be updated. "Better" beliefs (more sophisticated action schema), leads to better formulated expectations that leads to better predictions, which in turn reduces free energy and the cost of adaptation over time. Thus by supressing variational free energy, the organism learns.

This of course violates one of the fundamental laws of expected utility theory, namely the law of transitivity, since subjective value, and therefore preference, no longer can be said to be a stable property of any individual for any duration of time. Moreover, it is also the case that the vast majority of ends sought are themselves hidden states to be discovered through action. What most individuals think of as ends in themselves are in fact just the final means that best satisfies a felt biological or psychological want akin to a satisficing mean as opposed to a direct reward. Few would therefore view a car as nothing but a machine that transports one from A to B, or food as a nothing more than simple sustenance. These things are imbued with properties beyond their immediate function, which means that satisfying the same want can come in a myriad of different shapes, sizes, colours, textures and flavours, leaving plenty of room for adaptive actions to solve the problem of how to satisfy a felt want.

In spite of the fact that a wide range of different goods has the ability to satisfy the same want, it cannot be said however that these goods therefore are of equal subjective value. Different action sequences (policies) imply different evaluations based on the different time strains they incur. This means that even two identical goods can be valued differently at different points in time depending on the amount of adaptation that is necessary in order to obtain or use the goods. Thus, even the truism of the reflexivity axiom cannot be said to hold when factoring in any length of time. Having to adapt policies during action sequences must therefore influence subjective valuations in numerous obvious and subtle ways, meaning that one of the staples of microeconomic consumer analysis, the indifference curve, is likewise not commensurate with any notion of choice in reference to human action and behaviour.

This off course also violates the continuity and the revealed preference axiom. Even though the revealed preference axiom states that preference ordering is fixed unless relative prices change, changing prices alone cannot encompass the entirety of information relevant for adaptive actions. Therefore the axiom will not hold in its strictest sense, but will be correct if prices are substituted for information, a concept so vast and all-encompassing, that it in essence is meaningless. We can see this more clearly when considering that the suppression of variational free energy can be interpreted as a proxy for subjective value, where the factors that help to minimize the former inform the latter. Thus any attempt to measure or express numerically the latter, will likewise encounter the problem of measuring or expressing numerically the former, bearing in mind that variational free energy is an energy term and not a quantity, a fact that seriously problematizes the very idea of positive affine transformations essential for the expected utility hypothesis.

The indifference curve describes a situation where an agent chooses between various configurations or bundles of two different goods, illustrated by a downward sloping convex curve. The convexity of the indifference curve is due to a diminishing marginal rate of substitution between the two goods, indicating that agents would prefer a mix of goods instead of a higher quantity of just one good. Any point on the curve is thus equally as preferential as any other point, allowing for the aforementioned indifference to occur. But because the present value of all goods depend on the perceived time flow incurred from specific action sequences, urgency, and therefore preference, along with derived utility, shifts in response to any and all changes along the action sequence. This includes having to choose between bundles X and Y. Concordantly there can be no indifference between one choice and another. Partly because this would imply having to perform two mutually exclusive actions at the same time, but more importantly, because adaptation during any action sequence is unavoidable, and that it would be correct to view this adaptation as either a gain or a loss function on the end goal in response to any and all dynamic updates to any action sequence. A gain when an unexpected opportunity presents itself, shortening time flow perception, heightening urgency and increasing present value of the good towards which the opportunity was afforded, and a loss when externalities frustrates expectations, lengthening time flow perception, lowering urgency and decreasing present value of the good towards which opportunity was frustrated. Of course this gain and loss function includes, but are in no way limited to, price changes.

The crucial part is that the initial act must be based on an established belief system (action schema) including a policy for resolving uncertainty. This expectation term selects an action sequence born from experience, and in so doing, posits a more or less articulated goal conveying a direction and a measure of time denoting urgency. Unless this expectation term is perfectly satisfied during the proposed action sequence, meaning that no new information however minute can present itself, there is no reason to assume that setting out to buy milk will in fact result in the purchase of milk, to say nothing of a specific brand of milk. That’s not the interesting part however. The interesting part is that all new information is weighted against the expectation term generating a modified action sequence based on the available responses/actions “stored” in the wider action schema.

Solving for preference is thus an “either/or”, not an “and/or” problem, meaning that there is only one solution and one solution alone to how the individual solves a given action sequence. Giving up some of good X in order to get some of good Y is framing the problem incorrectly, since good X and Y are linked to two different policies that in many cases just happens to utilize a lot of the same actions up to a point. In instances where an individual does in fact substitute one good for another, this has to represent a decrease in the present value of the initial formulation, due to either a new opportunity, or an unexpected frustration.

There is a very good argument to be made at this point, that what we are trying to do as a species, is to reduce the cost of adaptation so that the fewest number of actions yields the highest possible payoff, and from the perspective of Hamilton’s principle of least action this would make absolute sense. This notion is not much different from the transaction cost theory laid fourth in Ronald Coase’s *Social Cost* paper from 1960 and later expanded by Oliver Williamson (Coase, 1960; Williamson, 1981, 2007). Adaptive actions are a response to unexpected information in the world, and since unexpected information always presents itself despite our best efforts to select uncertainty resolving policies, the cost of any biological system of handling the resulting added uncertainty will create a pressure for homogenization or harmonization by reducing the degrees of freedom in the internal system, the environment or both.

When an individual acts in the world every action is geared towards increasing his or her probability of success however this success is conceptualized. Only by positioning ourselves in an environment that is familiar or by manipulating the environment such that it becomes familiar, can we reduce the disparity between predicted an actual event in the world. By so doing, the cost of adaptation is minimized along with uncertainty and the probability of failure. Everything that helps to minimize the cost of adaption is valued in proportion to the cost reduction it facilitates, factoring in availability or ease of procurement. So while present value conceptualizations motivate action initiation, the sub strategy that seeks to reduce uncertainty generates additional present value conceptualizations, the sole purpose of which is to facilitate the expected action sequence of the origin, that is, the procurement of means in order to procure other means of a more primary nature.

4.2. Perception and the Problem of Choosing

The introduction of secondary means or facilitatory goods complicates things substantially and constitutes a problem at first examination. All perceivable goods are intricately related to a specific action sequence, and while many pursuits of various goods utilize a lot of the same actions up to a point, sooner or later a specific action segment must determine a point of differentiation in felt intensity and urgency between viable alternatives. The final choice herby has the ability to ratify the entire action sequence, creating the emergence of a strategy or patterned behaviour, in so far as the cost of adaptation is not influenced by unexpected frustrations or opportunities beyond a certain level of significance determined by the composition of the wider action schema. While the preference for one good over another is based on the preferred action sequence related to its acquisition, it is also the case that this relation is a self-contained process that does not need the introduction of comparative policies in order to explain itself.

However, the need to minimize adaptive behaviour, and hereby reduce uncertainty, introduces additional action sequences with related present values that are themselves self-contained, yet at the same time, conditional on a more primary action sequence. The problem then becomes how assign a present value to secondary means that are not, in and of itself, part of the original goal formulation or expectation term. It is not as easy as just stating that fishing is far more efficient with tools than without, since this would be begging the question. We actually need a process that explains how the direct pursuit of a felt want gets subverted by more roundabout means, while simultaneously recognizing that the goods that best satisfy a felt want and the goods that, at an earlier point in time, facilitate this satisfaction are based on two distinct action sequences.

The Problem of perception and choice are as follows: An individual has a felt biological or psychological want. The action schema posits a suggestion in the form of X containing a set of means that can satisfy the felt want. This set is contingent on the specific environment that has informed the action schema and a nature given component that signifies the individuals’ own ability to expand the action schema and to integrate new action sequences during adaptation. The size of the environmental component, which can be seen as a subset of X, determines the speed at which the initial action is initiated, since a greater set of means reduces the need for uncertainty resolution during action initiation. This is due to the fact that fewer potential actions facilitate a greater array of satisficing means, like going to the local super market, as opposed to having to choose between going north to pick berries or going south to hunt for game. It is of course also the case, that the need for uncertainty resolution increases once viable alternative satisficing means present themselves. With X being defined at the onset of an action, and a present value formulation imputed, what then will be the present value of a specific element in Y, if Y is the set of all facilitatory goods? Bear in mind that both goods are interrelated in the same action sequence, but must still be consequences of unique and specific action sequences before they can be expressed with an attached present value.

In order to answer this question, it is necessary to take a little detour where the nature of the relationship between perception and objects is taken fully into account. While the reader may become somewhat confused by the following segment, its inclusion is actually of great importance, although I fear that it will not be seen as such. How we effortlessly exploit the surrounding environment, and use tools to further our wellbeing, may blind us to the relevance of actually understanding how fundamental cognitive mechanisms create the emergence of higher systems of cooperation. Systems that is dependent upon such things as capital, savings, investment, and interest. The following segment is therefore an attempt to investigate the relationship between perception and objects in order to establish a link between cognition and higher economic concepts. By doing so we will be able to place facilitatory and satisficing means in a clear and concise context so as to not confuse the one with the other. In addition, we will become able to formulate universal rules of action that can come to serve as a foundation for economic theory.

4.3. Categories, Borders and Time Preference

Consider the strange category described by cognitive psychologist Lawrence Barsalou. "Things to take with you when your house is on fire" (Barsalou, 1983). Here is a category of objects that on the surface does not share any similarities with each other, and therefore might raise some questions as to what the category is comprised of, and if it even can be described as a category at all. When contemplating how seemingly unrelated objects, within a very short time, can be arranged into a meaningful set, it is important first to realize the appropriate level of analysis. In the case of the burning home, what objects one might save is dependent on the particulars of what constitutes the home and the self in the first place. Here, many things that on first glace might seem important for this conceptualization, is really of no importance at all. The walls are just walls, and things are just things, especially under potentially life threatening situations. But the actions that were performed in the space between the walls, and with various specific objects within the home or outside, are an intrinsic part of one's neurological makeup. Thus space and objects are an extension of oneself, insofar as space, object and action, through various arrangements, has solicited a neurochemical response with neuro-permutational consequences denoting the "weight" of the objects.

A simple little ornament perched on the mantel piece can thus play an important role for maintaining an internal state of neuro-chemical equilibrium, linking the internal with the external, and grounding characterological idiosyncrasies firmly in an objective and tangible environment from the perspective of the individual. For this purpose, no one would doubt the usefulness of a picture for instance, and here a picture is indeed the lowest hanging fruit for correctly conceptualizing the nature of our attachments to inanimate objects, attachments that might even seem irrational from certain perspectives (Epley, Waytz, & Cacioppo, 2007).

"Things to take with you when your house is on fire", can thus be seen as a category, not of what to save from one's home, but more of what to save of one self (Belk, 1988), in the sense of minimizing entropic surprise and conversely maximizing system stability under the circumstances. This is likewise linked to what some neurologists would call the maximization of model-evidence. The important thing is that the relationship between individual and objects is not one easily demarcated. It is a relationship better understood as an emerging property rooted in the same cognitive category (Epley, Waytz, & Cacioppo, 2007).

There are things and the space between things fundamentally. Things appear to have different qualities. These qualities fractionate in proportion to the level of sophistication in the observers perceptual systems, not only from a sensory perspective, but from a conceptual perspective as well. A very specific set of qualities categorises some things as trees, some things as plants, some things as animals and some things as human beings. But as a first approximation, all these objects are not subject to differentiation at all, as they are perceived by the sensory systems first and foremost as the absence of nothing with no discernible actions attached. Within milliseconds however, higher cognitive systems, as well as additional sensory cues from smell and hearing, categorizes and specifies the particular object as being one or the other based on an availability heuristic. The heuristic is referencing past experiences containing categories of problems and solutions, in turn based on present (and/or potential) desires and how to most effectively achieve them. Within these categories, the temporal unification of problem and solution is mediated by the object, where the solution now contains the object and the action that the object can potentially facilitate.

One might now object to the ease at which objects and the individual here are placed in the same category, but as stated previously, the demarcation between object and individual is in no way self-evident. Object recognition is performed by the inferior temporal cortex. When the visual field is stimulated, the information is passed on along the retina to the occipital lobe containing the visual cortex. Situated in the back of the brain, the visual cortex is responsible for generating the sensation of sight, but do not have the ability to actually understand the images that it generates. Information is then passed along to other areas of the brain, all specialized in performing a specific task that in concert encompasses the human visual system. The two most important systems in this process is commonly held to be the ventral and the dorsal system, often referred to as the "what pathway" and the "where pathway" respectively. While the dorsal pathway, (running from the occipitoparietal cortex to the posterior parietal lobe and extending to the dorsolateral prefrontal cortex), is responsible for spatial recognition, the ventral pathway, is responsible for object identification (Squire et al 2008). Identifying objects appears to be a gradual process, as signals are passed along the occipitotemporal cortex, through the anterior inferior temporal gyrus to the ventrolateral prefrontal cortex, adding complexity and reducing uncertainty along the way (Efremova & Inui, 2015).

The reduction in uncertainty can also be seen as a process that seeks to supress variational free energy in sequential feed-back loops, herby allowing complex objects to retain their foundational aspects, in line with some form of logical composition. This allows us to fractionate, analyse and understand each component of the object, by revisiting what the object was before it presented itself in its entirety. This type of analysis of course requires its own process, allowing for the representations of objects within the brain to be situated in very different parts of the inferior temporal cortex from individual to individual, in line with the perceived complexity of the object (Leibo et al, 2011).

What this means, is that objects do not just present themselves to us as they are. Instead they are more or less built from the ground up, retaining their more simple representations as part of the end product, something that is supported in the research by Donald Hoffman and more recently, James J. DiCarlo and Rishi Rajalingham (Hoffman & Prakash, 2014; Rajalingham et al, 2018). Never the less, this particular arrangement of the brain is what allows for the ease at which objects and animals can come to be imbued with very human traits, and likewise the ease at which we humans can imbue each other with inanimate and animalistic characteristics (Epley, Waytz, & Cacioppo, 2007).

The identification of animals, things and people are all made possible by the same brain region. As such, there is a great deal of interconnectedness between objects in the world, as they are represented in the brain. Just take a moment to think about how much of our language is saturated with figurative phrases and highly abstract representations. These abstractions are highly relevant and meaningful to us, exactly because they describe complex concepts at a low enough resolution, so that communication proceeds in reference to more universal principles. Instead of describing objects and concepts at their most complex state, we describe them by referencing the more simple components of which they are comprised. In this way, we communicate on a more instinctive basis, as if staying within the realm of the ability of our filo-genetic brain structure to represent the world.

The ability of humans, (and some animals), to abstract out present world states into possible future world states, allows for more complex goal formulations. This can basically be described as the understanding of the flow of time, and the impact that this flow can have on the surrounding environment. The knowledge that things change automatically, but that this change can be mitigated, accelerated, or in other ways exploited to one's advantage, or ignored to one's detriment, follows directly from this understanding. Confronting the world with a pre-set recognition of a changing nature thus codifies the aforementioned transmutability of objects, as being part of the same changing landscape.

If the landscape changes, then the reciprocal relationship between landscape and object must change as well, rendering the object malleable for whatever purpose is deemed worthy of pursuit. It is indeed this transmutability, nested within a changing landscape, that allows us to alter sensory inputs and change the classification of objects as needed, experiencing a log of wood as a chair in one instance, and fuel for a fire in the next. Again, it is not as such the shape of an object that categorizes it, but its function and its relationship with comparable functions, all in service to a more or less specific goal. We perceive the apparent similarities in shape between various objects, and herby infer probable usage. However, probable usage is a prerequisite for observing any object. Why? Because objects exist within a context, and this context needs to be understood before the perceptual systems can differentiate one object from all others, which again is a prerequisite for understanding the similarities between objects, be it shape, size or colour. As such, function shapes the object and not the other way around where shape implies function.

Before we can actually see an object, say an apple, we first have to need it in order for it to be visible. Irrelevant data is simply irrelevant, and not worth expending precious energy towards perceiving. This does not mean that we cannot perceive of an apple if we are not hungry, but that we need an acute understanding of the relationship between apple and hunger, in order for the apple to be relevant for the perceptual systems as an apple with all its known characteristics. In the absence of a relationship between apple and hunger it would be very difficult, if not impossible, to imagine how the apple would be perceived by the perceptual system and thus conceptualized by the brain. We therefore do not see an apple as such, but the actions afforded to us by the apple, exemplified in a set of characteristics that present themselves to us as an apple (Hoffman, 2014). Furthermore, apples simply floating around disembodied are meaningless, and to that extend, absurd. It is the solid grounding of the apple in a contextual framework that brings it into existence, both from a perceptual and a conceptual standpoint. As such, an apple found in the middle of the dessert will immediately bring one to either doubt the veracity of what one is experiencing, or to concoct speculative stories that might explain the mystery of the apple. The apple is out of its elements, and must be so for a reason, casting the object itself into question until the context that surrounds it is restored to a satisfactory degree.

When we see or experience things in the world, we do not see the things in themselves, for these things do not exist in a vacuum. More accurately we see stories or mental constructs that involve everything we know about both the object, and the landscape or framework that surrounds it, for the purpose of manipulation. Our perceptual system is called on to perform a specific function, scans the environment and samples the information presented in the visual field as relevant or irrelevant. Through the interaction between perception and action, a logical structure then emerges, creating the foundation for cognition itself where branch is to tree, what arm is to body, one concept a consequence of the other. By the same measure causality emerges, though not as a consequence of physical reality per se, but as context dependent categories layering perceptual inputs on top of each other in a structure conducive of action or manipulation, space, object and action respectively. So even though the inferior temporal cortex is great at categorizing and classifying objects, animals and individuals, it is the breakdown and analysis of their constituent elements, and the subsequent removal of borders that enables the construction of categories of action in which sub-elements of a goal can be identified in the environment and used for the requisition of the goal. This breakdown between borders is the same concept that Dr. Donald W. Pfaff, makes use of in “The neuroscience of Fair Play” where he tries to explain the neurological foundation for the *golden rule* (Pfaff, 2007).

So what is an apple? Aside from its direct characteristics and the emotive markers it communicates, it is the branch it hangs from, the tree crown, the trunk and the soil. Alternatively it is the package it is wrapped in, the table it rests upon, the room, the house and the city. This means that objects in the world is in a natural relationship with each other allowing objects to be “hidden” within other objects placed higher in some form of conceptual hierarchy.

To finally answer the question of the present value of facilitatory goods, it becomes clear that facilitatory goods have no present value, in so far as they are completely integrated into a given expected action sequence, something that is made possible by the breakdown of borders between categories. They do however have a cost associated with any deviation from this expectation, and a cost derived from their discovery or acquisition. This cost is expressed as entropic surprise, the minimization of which expends energy that in turn can be expressed as a loss function. Initially this must increase time perception, lower urgency and present value of the satisficing means. Since this would move satisficing means further into the future, Hamilton's principle of least action is violated, unless we posit that time flow perception is independent of direct action, but dependant on expectations. We therefore have two time flow perception terms, one of which is dependent on expected actions, and one is dependent on actual actions, but because expected time flow perception is a rather clumsy term of phrase, we shall simple call this time preference hence forth.

If a process of uncertainty resolution has consistently yielded the same result, then the result will be integrated into the time preference term, consequently moving satisfaction closer in time. Continuing to resolve uncertainty after this point will be inconsistent with the variational free energy principle, since beliefs will not have been updated through active inference. Facilitatory means will become integrated into an action sequence in the same manner, as part of expectations or beliefs. Only if facilitatory means in one way or the other deviate from expectations will they require deliberative effort. This will move satisficing means further away in time, imposing a loss connected to expectations, plus a cost associated with the need to again resolve uncertainty.

Time preference is generated implicitly by the composition of the action schema, and becomes explicit through the selection of a particular action sequence aimed at a particular goal that communicates an expected present value formulation. Embedded within time preference is already all the relevant facilitatory means extracted from the action schema, and as such, there is no loss function associated with roundabout methods of acquiring satisficing means. But since the loss function is connected to the cost of surprise, it is only for the purpose of minimizing a specific surprise that a tool or facilitatory mean can come into existence. This means that the act of discovering a stick that can reach previously unattainable fruit for instance, or a rock that can open a nut as necessity dictates, is perfectly commensurate with Hamilton's principle of least action, but the invention of a stick for the purpose of acquiring a greater quantity of an easily attainable fruit is not. If a perfectly good action sequence exists, then modifying this action sequence is impossible based on time preference alone.

Time preference only dictates initial urgency based on expectations, and therefore does not factor in to adaptive actions. Inventing a basket in order to carry fruit is therefore highly improbable if a perfectly good solution already exist at a less complex conceptual level. There has to be intermediary "evolutionary" steps that in the last instance might result in a basket for the purpose of carrying fruit, but not before the basket previously was a specially shaped rock, and that the rock likewise previously was a leaf for the purpose of carrying water to which hands are far less suited than carrying fruit. If for instance an essential resource has become scarce, this would be reflected in a higher time preference, increasing the present value of satisficing means. The relative cost of adaptive actions must therefore decrease since uncertainty resolution becomes more important when minimizing entropic surprise along with variational free energy. It is thus an imperative that actual time flow perception is lengthened, slightly decreasing the present value of satisficing means, as this allows for energy to be expended towards facilitatory means. The individual therefore moves from an exploitation state, to an exploration state.

If an organism is successful in this endeavour, it will possibly have integrated a facilitatory mean into its action sequence, which will function as an extension of the body's own capabilities or an extension to which the body's capabilities has already been augmented, in order to "mend the gap" between time preference and actual time flow perception.

We can now begin to see a proto-structure emerging between different orders of facilitatory means, in which a discovery process connects a duration of time with some primitive notion of productivity. It therefore becomes clear that the primary cost associated with the procurement of a facilitatory mean is equal to the cost of adaptation expressed as time preference plus actual time flow perception. When over time, fewer and fewer actions has the ability to yield a greater array of facilitatory, as well as satisficing means, up to a point, then actual time flow perception will approach zero, which is the same as saying that the KL divergence approaches zero, leaving only time preference as the real cost of adaptation.

The order of facilitatory means, while originally explained as the transmutation of a single mean, is therefore intricately connected to how time preference is expressed when influenced by real factors. From this it can be seen that two conditions must be present if the order of facilitatory means is to be extended. One: Time preference is lengthened and actual time flow perception shortened. Two: Time preference is shortened and actual time flow perception is lengthened. In modern times, this state of affairs can actually co-exist within the same individual depending on how the individual conceptualizes his or her specific endeavour. This is because that modern society supports the ability of few actions to yield a large array of facilitatory and satisficing means, meaning, that the opportunity cost between exploration and exploitation is equally low, so low in fact, that it is possible to switch effortlessly between one state and the other. Yet for the sake of simplicity and clarity, the two states will be treated as defining different individuals, whose states are contingent upon the exchange of resources between them. Obviously there will be a time lag between this resource exchange, and it will be here that the concept of interest will emerge.

At this point it might be pertinent to interject, that a low time preference individual will not act to postpone satisfaction simply for the sake of postponing satisfaction, but will have preferences geared towards facilitatory and satisficing means not immediately obtainable, that is, facilitatory and satisficing means that generally takes a longer time to bring into existence compared to other means potentially available. If however these means could become available at an earlier time than expected, so much the better, anything else would be violating Hamilton's principle of least action.

Low time preference individuals will supply resources to high time preference individuals in exchange for future facilitatory or satisficing means that match the urgency, or lack thereof, of their preferences. However, because of the uncertainty involved in operating within a complex system, where the topographic landscape constantly changes, simply supplying resources in exchange for future resources will violate Hamilton's principle of least action, as low time preferences individuals will have increased variational free energy based on an unnecessary increase in uncertainty that cannot be reduced through any action sequence, save for one. This policy will have to be generated by the expected future yield from the actions of the high time preference individuals, and therefore targets a specific future want in the form of a more or less specific facilitatory or satisficing mean. The policy is therefore constructed for the purpose of bringing the future want closer in time, either by directly shortening the production process, or by demanding a fixed rate of return with a present value higher than the present value of the resources supplied. This exchange pattern will therefore result in the formation of interest, which in essence can be viewed as the cost of time given uncertainty, and will function as a rough gauge of future demand. A high general interest rate will thus be commensurate with a low future demand, and a low general interest rate commensurate with a high future demand. If this interest rate is not generated by the exchange patterns of individuals with different time preferences, it will cease to function as an interest rate that provides a rough gauge of future demand. In the same manner, if a price for a particular good is not generated by the exchange patterns of individuals with different present preferences, then prices will cease to function as a rough gauge of present demand.

4.4. Principles of Action

It has been suggested that organic systems act by supressing variational free energy. This involves the process of active inference, where organisms initiate actions based on beliefs about states in the world, and subsequently search for corroborating information in order to reduce uncertainty about following actions. Organisms do not seek to optimize an action that leads to a preferred world state, but seeks to optimize sequences of actions or policies in which order and temporal considerations are paramount. Once this temporal dimension is added, the efficacy of a policy, and the actions that the policy is comprised of, can be expressed as a time average or a path integral, suggesting that time and action are mutually determined by each other. Every action sequence therefore has an expected cost expressed as a time duration denoting urgency, also termed time preference, and an actual time flow perception connected to uncertainty resolution, which can be called an adaptation term. It can therefore be stated that satisficing means will always be preferred sooner rather than later, all things being equal.

The more urgent wants, the less time must be spent on satisficing the want, and therefore a shorter action sequence is selected for the task. This however affects uncertainty resolution, resulting in a wider array of potentially satisficing wants, instilling the agent with a more exploratory behaviour. This exploratory behaviour is made possible by other agents whose wants are less urgent, allowing them to formulate wants in other ways, or to re-structure actual preferences with respect to a lower time preference.

Being less discriminatory in defining wants, does not in any way, shape or form imply indifference. Subjective valuations, or preference, are connected to the actions needed in order to satisfy them. It is the actions or policies that are being optimized, and as such, it is the actions that carry the valuations fourth in accordance with expectations, and the need for adaptation, until satisfaction is reached. Since no two mutually exclusive actions can be performed at the same time, no two satisficing means can be valued identically at the same time.

This is equally true for facilitatory means, but unlike satisficing means, facilitatory means cannot be expressed with an attached present value. What determines the present value of satisficing means, are the time preference and actual time flow perception that encompasses the expected and the adaptive parts of the action sequence of which facilitatory means are a part of. If we were to assign a present value to facilitatory means, we would thus be confusing means and ends with each other, and implicitly state that we valued non-consumables because of their taste for instance, all the while trying to impute the present value of not performing an action. What we mean to say is that facilitatory means influences the present value of satisficing means because of their ability to shorten action sequences, and hereby move satisficing means closer in time towards present satisfaction. This does not mean that we cannot put a price on a facilitatory mean in the final analysis however, but only that this price has to be a derivative of its expected yield.

Given that subjective valuations are derived from action, it follows that upon completing an action sequence, where a satisficing mean has been obtained, the next action sequences must be further away in time. This implies an order of preference where by each subsequent action sequences must be geared towards lesser and lesser valued means as expressed by the cost of acting. The difference between this principle and the revealed preference axiom is that preference is only substantiated when acted out. In no other way can beliefs and expectations be updated such that uncertainty is resolved. Without action we therefore have to state that preference are ordered in a certain way unless anything however insignificant changes, which is a useless statement.

In a similar manner, if in possession of a set of facilitatory means, even if these facilitatory means can be said to be interchangeable, then each mean will be integrate into a specific action sequences for a specific purpose. As a consequence of the order of preference established above, then each additional facilitatory mean must then be assigned a less and less “important” action sequence, geared towards satisficing means further and further away in time. It is therefore obvious that each additional facilitatory mean must likewise have a diminishing marginal value or utility.

Because of the cost associated with action, leisure will always be valued higher than work. There can however be some confusion as to what actually constitutes work and leisure. A simple explanation would be that work is remunerated activity and leisure is not, but this is purely conventional. Leisure refers to activity that has the lowest possible cost of action, relative to other activities one could engage in. It just so happens that low cost activities is highly correlated with activities that are under complete or partial control of the individual.

If organisms act by minimizing variational free energy, the KL divergence will be at its lowest in stable and predictable environments, most commonly associated with environments in which the individual has the ability to voluntarily avoid the need for costly adaptations. It is however perfectly possible to not be working while working, and to work while not working, as it is also perfectly possible to consider the complete restoration of an old sports car leisure, but emptying the dishwasher as work. Emptying the dishwasher will in this example itself be a costly adaptation when an expected action sequence, not involving dishwashers, is in process. Leisure is therefore best associated with voluntary action sequences in which the need for adaptive actions, for one reason or another, is minimized to its highest possible degree. Without external forces or pressures from hitherto not considered factors, leisure activities would encompass the entire action schema, excluding actions driven by biological necessity.

Lastly it can be concluded, that given the cost associated with action, there is a natural tendency towards risk aversion. When behaviour could be construed as risk seeking, it has to be based on the composition of the expected action sequence, and will therefore not be perceived as risky by the individual. Only an ex post examination of the relevant objective factors involved can possibly reveal risk seeking behaviour as such, and only if the risk seeking behaviour failed to facilitate the expected outcome. While risk aversion is a natural consequence of uncertainty resolution during action, risk seeking behaviour can only manifest when uncertainty resolution is either impossible, or the means with which uncertainty could be minimized is unknown. It is therefore action referencing time preference only, in a one-step action sequence.

Many forms of games of chance involving money can therefore not be considered risk seeking behaviour, in so far as the individual takes steps to reduce uncertainty in order to influence the odds of success. Pure gambling is a different matter however. In pure gambling no uncertainty resolution is possible and has to be based on expectations alone. The expectation, or time preference, is referencing a single instance of success, or an easily imagined prospect of success and the action sequences is a direct copy of the original action sequence involved with that success, real or imagined. Pure gambling therefore has to be based on simple repetitive acts that have little or no need of adaptive actions. This in turn makes the act of gambling practically costless, and whatever money that might be involved is easily converted to valueless tokens, conceptually speaking, given that subjective value is derived from the cost of action. This becomes even easier when money is literally converted to tokens, as is the case with many one-armed slot machines for instance.

We can see this more clearly if we consider equation 8, and set all things equal to zero leaving only the subjective utility component $E_{Q(O_\tau, S_\tau|\pi)}[\ln P\,(O_\tau|m)]$. It is important to note that when we are talking about risk, we are not referencing probabilities in the traditional sense, but probability densities as they relate to physics and thermodynamics. As such, the KL divergence $\boldsymbol{E_{Q(O_\tau, S_\tau|\pi)}}[\ln Q\,(S_t|O_\tau, \pi) + \mathbf{ln}\,\boldsymbol{P}\,(\boldsymbol{O_\tau}|\boldsymbol{m}) - \mathbf{ln}\,\boldsymbol{Q}\,(\boldsymbol{S_\tau}|\boldsymbol{\pi})]$ in equation 8 is an energy term that scores a measure of deviation between two densities P and Q. The reason for this is that we have to be able to formulate a theory of choice that does not rely on the ability of individuals to recognize the numerical probabilities of risky prospects, just as we have to be able to explain how individuals chose a larger over a smaller quantity, without the ability to count.

Perhaps a more intuitive way to demonstrate this point could be to envision a closed system that organizes its internal activities based on relative temperature changes. It just so happens that the brain is such a system, where most processes that govern neural activity are temperature dependent (Kiyatkin, 2010). When dealing with such things as energy and temperature, we are dealing with what is called intensive properties as opposed to extensive properties. Extensive properties are defined by solid matter and discrete quantities that expand or extend in direct proportion to what is added. Intensive properties do not. If one adds 30 degrees Celsius to a room that is already 30 degrees Celsius, the temperature does not rise to 60 degrees Celsius.

Based on what has been presented op until this point, some principles of action can be stated. These principles are meant as first principles from which further economic statements can be derived.

1st principle: Facilitatory and Satisficing means will always be preferred sooner rather than later, all things being equal.

2nd principle: The more urgent the want, the less selective the definition of the want and visa-versa. The less selective the definition of the want (high time preference), the more exploratory the behaviour, if, and only if, adaptation is high. The more selective the definition of the want (low time preference), the more the exploitative the behaviour, if, and only if, adaptation is high.

3rd principle: No two satisficing means can be valued identically at the same time.

4th principle: Action implies an order of preference where by each subsequent action sequence must be geared towards lesser and lesser valued means as expressed by the cost of acting.

5th principle: All facilitatory and satisficing means have a diminishing marginal value.

6th principle: Leisure will always be valued higher than work.

7th principle: Agents have a natural tendency towards risk aversion.

8th principle: Facilitatory means do not have a direct present value, but must ultimately be expressed as a derivative of the present value of the satisficing means they help to facilitate.

9th principle: Exchanges between individuals with different time preferences must result in the formation of interest. This interest will be a gauge of future demand, and be expressed as the cost of time given uncertainty.

5. Behavioural Economics and the Variational Free Energy Principle

In the following section the variational free energy approach will be compared to behavioural economics by applying the concepts and ideas that have been formulated up until now.

5.1. The Allais Paradox and Prospect Theory

The Allais paradox was formulated in 1953 by the French physicist an economist Maurice Allais. Allais was able to demonstrate that the independence axiom was not consistent with actual human behaviour. The experiment that Allais devised for this purpose was as follows:

1. Common Consequence:
   Situation A: Certainty of receiving 100 million.
   Situation B: 10% chance of 500 million; 89% chance of 100 million; 1% chance of nothing.

   Situation A´: 11% chance of 100 million; 89% chance of nothing.
   Situation B´: 10% chance of 500 million; 90% chance of nothing.

2. Common Ratio:
   Situation C: Certainty of receiving 100 million.
   Situation D: 98% chance of 500 million; 2% chance of nothing.

   Situation C´: 1% chance of 100 million; 99% chance of nothing.
   Situation D´: 0.98% chance of 500 million; 99.02% chance of nothing.

In situation A and B there is a common consequence of receiving 100 million with 89% chance, where situation A can be viewed as yielding 11% chance of receiving 100 million, and 89% chance of also receiving 100 million. In situation A´ and B´ this common consequence has been removed, but should not influence how agents decide what bet to pick if the independence axiom is correct. If individuals prefer situation A, they should therefore also prefer situation A´, and the same goes for B and B´. Allais however demonstrated that this is in fact not correct. Individuals had no trouble in choosing A and B´ thereby violating the axiom (Andreoni & Sprenger, 2010). In situation C´ and D´, the probabilities involved in situation C and D have simply been divided by 100, meaning that the ratios involved in both situations are the same. Therefore, if C is preferred to D, the independence axiom states that C´ should also be preferred to D´. Again, Allais demonstrated that this did not match the preferences of actual individuals, including some distinguished economists such as Samuelson and Savage (Andreoni & Sprenger, 2010).

Since 1953, similar experiments have consistently demonstrated the Allais paradox, and thus the apparent refutation of the independence axiom (Kahneman & Tversky, 1979). More than simply refuting the independence axiom, Kahneman and Tversky extended Allais' paradox in 1979 by introducing a broader set of violations of the expected utility hypothesis, including the reflexivity, revealed preference and the continuity axiom.

In order to explain these apparent anomalies, Kahneman and Tversky proposed prospect theory, wherein probabilities were not treated extensively, but multiplied by a decision weight. Furthermore, value had to be viewed as a reference point rather that a final state (Kahneman & Tversky, 1979). Prospect theory distinguishes two phases in the decision process, where the first one is the editing phase, and the second is the evaluation phase. In the editing phase, individuals first code their prospects as either a gain or a loss relative to a reference point. This reference point will usually be the individuals' normal state of wealth all things considered, and will be influenced by framing effects as well as the expectations of the decision maker. Secondly, individuals simplify prospects by combining the probabilities of identical outcomes whenever possible. Thirdly, and if possible, riskless and risky prospects are segregated so as to convey the riskless component separately from the risky component. Finally, individuals will disregard information that is common to two or more prospects, and proceed to evaluate them by referencing their relative dissimilarities. It is also worth mentioning, that according to prospect theory, individuals will round probabilities to numbers that are easier to handle mentally, as well as disregard extremely unlikely outcomes (Kahneman & Tversky, 1979).

In the evaluation phase individuals compute a utility function based on the probability of an outcome, weighted or modified in accordance with the editing phase, and a subjective value scale that sets gains and loses relative to a subjective reference point.

In general form, prospect theory can be written as

$$V = \sum_{i=1}^{n} \pi(p_i) v(x_i)$$

Equation 10

where π is the weight on probabilities, and v measures the value of the deviation from the reference point.

Prospect theory thus states that the value functions of the individual will generally be S-shaped, defined as a gain or a loss from a reference point, convex for losses and concave for gains and steeper for loses than for gains, reflecting that individuals generally experience loses as greater than comparable gains (Kahneman & Tversky, 1979).

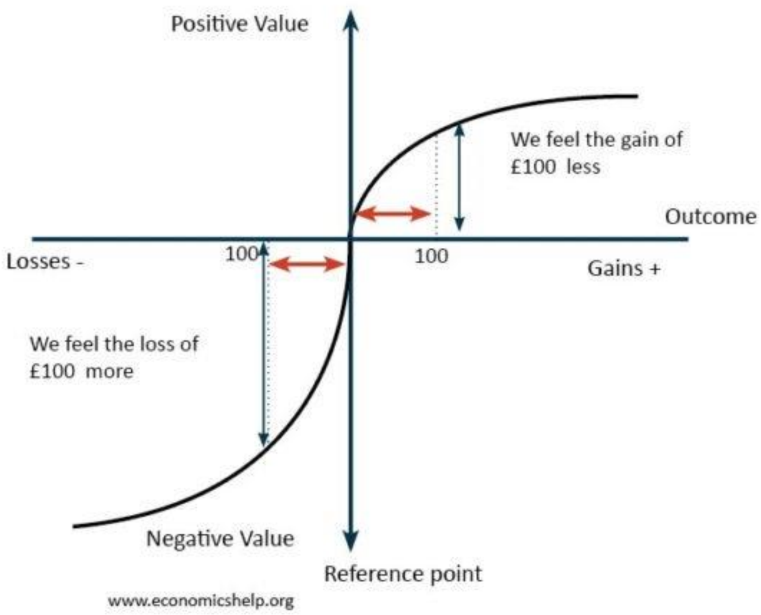

Figure 7

5.1.1. Variational Free Energy Approach

The idea that agents have a reference point from which gains and losses are compared is in perfect alignment with the variational free energy principle. Agents thus minimize variational free energy by selecting action sequences that optimizes beliefs about world states. This is the same as stating that beliefs are the references point to which gains and losses are compared. Kahneman and Tversky also posited something they called a certainty effect, where the prospect of a bet with a certain outcome will be valued higher than its expected value would suggest when compared to other bets (Kaheman & Tversky, 1979). This is likewise a property of the variational free energy principle. It must however be stated that absolute certainty does not exists in the real world, but if it did, these outcomes would be highly favoured as the agent would not need to adapt actions during their related action sequences. Variational free energy is the upper bound on surprise, and if agents did not experience surprise, the KL divergence would be at a minimum. By using a reductio ad absurdum, we will be able to see this more clearly. If an agent were confronted with a certain payoff, and in addition did not have to wait any amount of time for its manifestation, then the impossible situation would manifest where the KL divergence is zero, signifying that the present value, or “expected value”, of the payoff would be infinite. It should however be emphasized that this is an absurd situation, as it would only apply to a world in which time was not a factor.

In the variational free energy approach, subjective value is initially derived from an expected action sequence containing an energy term that communicates the “length” of, or the effort associated with, the policy. This is then positively or negatively modified by the need, or lack of need, for adaptation in real time. When confronted with numerical probabilities linked to risky prospects, to which we have not adapted in evolutionary terms, and to this day are exceedingly rare phenomenon, the normal way that agent’s experience subjective value then has to be translated to a “different language”. It is expected that something will get lost in translation, and the question then becomes what?

The short answer would be precision, something that is also covered by prospect theory, where agents tend to round probabilities, and likewise disregard extremely unlikely outcomes. The long answer would involve a transformation algorithm that changes abstract themes and concepts to actionable events, that is, action sequences that can actually be carried out. Such an algorithm would however require a considerable amount of data collection and research into individual choice patterns when confronted with abstract decision scenarios, and is unfortunately beyond the scope of this thesis. The main problem with probability distributions, as opposed to probability densities involving energy, can nonetheless be re-stated in the form of a question. How is any individual expected to carry out a 78% prospect of receiving a brand new car? And how exactly is this question to be understood in terms of updatable action sequences, when the question does not involve any explicit actions? It is however clear, that in the absence of the need for adaptation, as when agents are confronted with certain prospects, then certain prospects must be favourably weighted, since the mere act of conceptualizing a risky prospect is a form of adaptation, though imaginary.

Finally, prospect theory suggests that losses are felt more intensely than comparable gains. This loss aversion is not to be confused with risk aversion that relates to the tendency of individuals to avoid or reduce uncertainty. Loss aversion on the other hand, relates to losses compared to an expected outcome. While the free energy principle states that individuals indeed employ expected action sequences to which gains and losses are compared, there is nothing in the theory that would suggest that losses should be felt more intensely that comparable gains. The obvious question is how much more losses are felt, and here one study has stated that losses are felt twice as intensely as gains (Kahneman & Tversky, 1992).

This cuts very close to some form of cardinal utility, and introduces the notion of treating gains in the same manner. We should therefore be able to say that an individual feels the gain of drinking milk, twice as intensely as eating yogurt, that watching television is 1.87 times as enjoyable as playing scrabble, or that a child is loved 1.15 units more than another child. When suggesting that something is felt twice as much as something else, then one has to be able to answer the question; twice as much as what? 4, 16 or perhaps 9 2/3?

In the variational free energy principle, gains and losses from a reference point are meant to simplify the adaptive actions that move a satisficing mean closer to, or further away from, the agent in time, communicating the need to exude more or less energy in his or her pursuit, thus affecting the efficacy of the action sequence, and therefore the "value" of the end goal. These gains and losses are surprise factors that either reduce, but mostly increase, variational free energy, forcing the agent to adapt behaviour by selecting a new and appropriate action sequence from his or her action schema. While a negative surprise may initially frustrate the agent, it is ultimately the adaptive response that determines if the entire action sequence was a loss or even a gain, insofar as a better action sequences was found. Here gains will be associated with pursuits that renders the outcome better than expected, while the payoff is associated with the outcome that goes more or less as expected, with due consideration for the composition of the individual action schema. A loss will be an outcome that exceeds some lower bound of expected payoff, regardless of whether or not this loss can be said to have been a gain by someone else based on some objective standard. We can however in no way get any more precise in our specifications without introducing the absurdities of cardinal utility.

5.2. Hyperbolic and Quasi-Hyperbolic Discounting

In standard rational choice theory, agents are said to discount the future at a constant rate, thus expressing intertemporal preferences in accordance with an exponential discount function. This means that the only factor that can influence the individuals' marginal rate of substitution is time. As such, and agent will discount an expected future reward with the same rate throughout the entire waiting period. An implication of this is the *consumption independence* assumption, which states that consumption in time t is independent of consumption in time t+1, or t-1. Having pizza yesterday should therefore not influence ones preference for pizza today or tomorrow. Furthermore, preferences are not expected to change over time in accordance with what is termed *stationary instantaneous utility*, and since the discount function is a single variable, the same discount rate must apply to all forms of consumption. This is called *Independence of discounting from consumption* (Ainslie, 2003).

A considerable body of literature however, suggests that intertemporal preferences are dynamically inconsistent and that the individuals' discount function is approximately hyperbolic, not exponential (Laibson, 1994, 1997; Ainslie, 2003; Loewenstein et. al., 2002). Hyperbolic discounting gives agents the ability to discount future rewards at different rates in different points in time during the waiting period, and manifests in an individuals' tendency to prefer smaller rewards sooner, over larger rewards later (Ainslie, 2003). This has many behavioural implications, and helps to elucidate on addictive behaviour as well as explain the many forms of lapses in willpower that exist under the common rubric of "bounded willpower". For instance, why do individuals feel the need to commit to a future preference formulation by trying to safeguard themselves against their own potential future impulsivity (Loewenstein et. al., 2002)? If individuals where exponential discounters, then this behaviour would not happen since the exponential discounter would anticipate his or her lack in willpower and consequently adopt a discount rate to reflect this (Ainslie, 2003). Both David Laibson and George Ainslie have likened this phenomenon to a form of split personality in constant battle with each other over control. One personality rational, far sighted and thrifty, the other irrational, impulsive and profligate (Laibson, 1997; Ainslie, 2003).

In quasi-hyperbolic discounting, the discount rate is modified so as to accommodate an exponential discount function once the smaller sooner effect has been taken into account. The quasi-hyperbolic discount rate is however not quite as good at modelling real discounting behaviour as the pure hyperbolic discount rate, but still superior to the exponential rate (Ainslie, 2003).

The exponential and the hyperbolic discount functions are often depicted in the following manner.

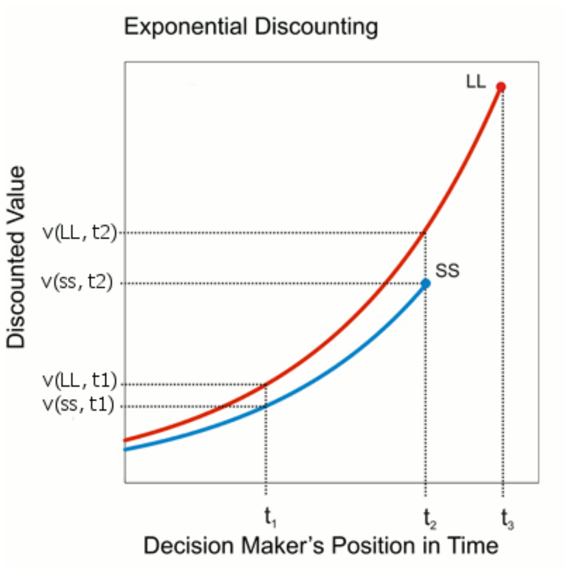

Figure 8

In figure 8 we can see two exponential discount rates, one for a small payoff and one for a large payoff. If an agent communicates a preference for one over the other, the relative discount rates will have to reflect this initial preference order for all periods until maturation. Furthermore, the slopes of the exponential discount rates reflect that the relative payoffs are discounted at a constant rate. One can interpret the discount rate as communicating how long an agent will want to wait for a specific payoff recognizing that waiting incurs a cost on the agent, or alternatively, the rate at which an agent will want to be remunerated for deferring present consumption. The longer the wait, the more a payoff will be discounted, commensurate with a higher rate of remuneration.

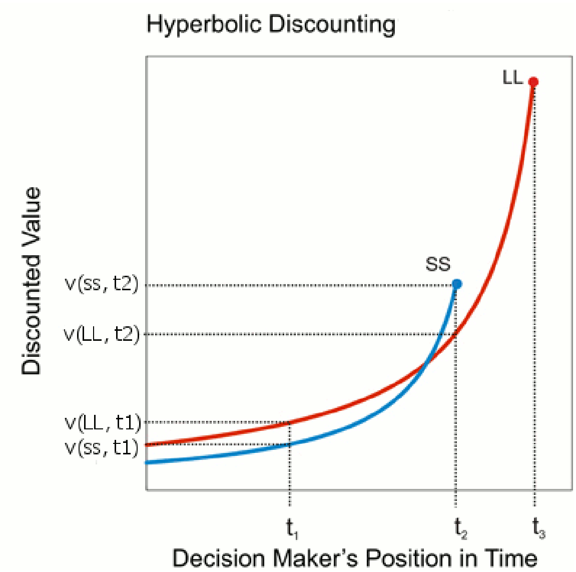

Figure 9

In hyperbolic discounting, depicted in figure 9, the rates have been modified so as to reflect the tendency of agents to reverse their preference ordering in accordance with the smaller sooner effect. As an example, agents could have a tendency to prefer a payoff of 100 now over a payoff of 150 three months from now, but not a payoff of 100 three months from now, over a payoff of 150 six months from now. The hyperbolic slope also reflects the reported tendency of individuals to discount a future payoff in a dynamically inconsistent way (Thaler, 1985; Lowenstein et.al., 2002; Ainslie, 2003). However, these results often come about by asking respondents at what points in time they will be indifferent between one payoff and a higher payoff, for instance, a payoff of 20 now or a payoff of 50 three months from now and so on (Thaler, 1985; Lowenstein et.al., 2002). While indifference is a term that practically every individual understands, it is nevertheless a term that is impossible to demonstrate in practice. Discount rates established from questionnaires involving points of indifference should therefore be met with a fair amount of scepticism.

5.2.1. Variational Free Energy Approach

In the variational free energy principle, it is impossible to separate present choices from intertemporal choices, as all choices are based on time preference and adaptation. The difference between Kahneman and Tversky's certainty effect and the smaller sooner effect is therefore purely academic. An important thing to note is that, since time and action combine to establish subjective valuations, the strain that is associated with effort carries over into the perception of the passing of time. Even though the majority of deliberations on possible future rewards involve actions to be undertaken, this need not necessarily be the case. The strain associated with action is implicit within the passing of time, even when no specific or discernible actions are required. The expected action sequence will therefore not necessarily contain any actions, but simply a time sequence generated by the composition of the action schema. In addition, it is also worth considering the weight of the possible actions and opportunities the future reward is deemed to facilitate, and therefore the cognitive strain that the waiting period therefore must produce. It is however clear that a direct conflation between what could be called an active and a passive time deliberation, is not a tenable proposition. Even though the general principle remains the same, this cannot be said about the explicit present value of any remote rewards, when the means to obtain the rewards differs. We can therefore not expect individuals to have a single discount rate for all possible goods, but will adopt rates in general accordance with their time preferences.

It is customary to view discount rates in relation to the diminishing value of future outcomes given time preference (Loewenstein et. al., 2002), that is, the fact that individuals will prefer satisfaction sooner rather than later, all things being equal. While there is absolutely nothing wrong with this view, it is quite informative to take the concept of time preference and apply it to deferred satisfaction rather than expected future satisfaction. When doing this, the discount rate becomes a rate of remuneration for deferred consumption, that is, the rate of compensation at which an individual will supply his or her own resources to another party.

From the first perspective we see that a discount rate will always be positive, because anything else would imply that an individual would not prefer the delivery of a good today over the same good delivered after a week, a month, a year or even twenty years (Mises, 2007). From the second perspective we see the same argument in work, but in reverse. Because individuals prefer satisfaction sooner rather than later, not consuming something that is readily available or expected to be available given an action sequence, must incur a cost. This cost has to be remunerated in some way, and will manifest itself as an increasing sequence of remuneration given time preference.

Since the variational free energy principle states that individuals seek to optimize action sequences in order to optimize beliefs about world states, we are therefore in a unique position to recognize, that the preference for improving sequences is not at odds with positive discounting (Loewenstein et. al., 2002), but that they are in fact one and same thing. That individuals have a positive discount rate, that is, a rate that over time increases the total amount discounted, signifies that waiting for a payoff diminishes the present value of the payoff by a larger and larger amount the longer one has to wait. This implies that utility is maximized if the payoff is received or consumed immediately, but appears to be contradicted by the fact that individuals in general prefer improving sequences, where the payoff that should convey the greatest utility is postponed till at the end of this sequence (Loewenstein et. al., 2002). If individuals simply select an action that maximizes a preferred world state, then this analysis would be a correct. When only given one shot, better make it count. However, there are no single actions in real life, save for acts of pure gambling, but always sequences of actions that lead to either the confirmation or the updating of beliefs about world states. When engaging in an action sequence, we therefore have to consider the cost of waiting until the sequence has concluded, insofar as the sequence has a natural conclusion. This waiting cost must be remunerated in accordance the individuals' time preference or discount rate.

Acting in time necessarily introduces the problem of acting without reference to a tangible reward, which as a function of time, must be abstract or imaginary in the present. Forming expectations therefore implies that whatever is expected to come into one's possession already is in one's possession psychologically speaking, thereby bridging the temporal "void" between initial action and consequence. If expecting a three course meal, then the three courses is already being discounted, and the individual must therefore wait until the entire dinner sequences has concluded. Waiting for the entrée is discounted less than waiting for the main course, which in turn is discounted less than waiting for dessert. This positive discount rate will then have to be compensated by a proportional increase in "utility", thereby optimizing the entire action sequence step by step, and reducing variational free energy over the span of the waiting period. This also serves to explain why individuals generally prefer an increasing wage profile over a decreasing profile, even when the decreasing profile yields more money (Loewenstein et. al., 2002). Here the time preference of the individual is discounting the future, which in turn must yield a higher and higher rate of remuneration for the losses incurred by the sheer passing of time.

The *peak-end rule* is a mental heuristic first discovered by Daniel Kahneman and Barbara Fredrickson, where individuals generally experience the end of a sequence as defining the entire span of the sequence (Kahenman, 2011). This rule is explained perfectly by the concept of step by step optimization informed by the variational free energy principle. Instances where individuals under physical duress report that a declining pain sequence is more preferable than a flat or increasing pain sequence, even when the declining sequence is longer or more intense (Kahenman, 2011), is thus the inverse of step by step optimization and likewise fully explained by this theory.

In hyperbolic discounting, discount rates are said to be dynamically inconsistent. One way this is evident, is that individuals tend to discount high amounts of money at a lower rate than low amounts. Moreover, there is some experimental evidence suggesting that the same payoff will be discounted at a lower and lower rate during the discounting period (Kaheman & Tversky, 1979, 1992; Thaler, 1985; Lowenstein et.al., 2002). While this may or may not be reflective of how real individuals discount the future, the variational free energy principle does not offer any insight as to why this may be the case. The reason for this is that value or utility is an intensive property controlled by an energy functional and relative temperature changes within the brain, and are therefore not subject to traditional methods of measurement. Subjective utility or individual discount rates does not have a corollary numerical representation to be found anywhere in the real world, but exist solely as intensities subject to ordinal ranking by more or less energy and hotter or colder temperatures.

The smaller sooner effect is on the other hand fully explained by the variational free energy principle, as this effect is a function of the order by which actions must be undertaken. It does however open up a discussion of the passions, and the extent to which we are in control of our emotions and our baser drives. This has fostered a number of *multiple-self* models, such as Ainslie's battling personalities (Ainslie, 2003), and Thaler and Shefrin's *planner-doer* model (Thaler & Shefrin, 1981). While the discussion itself has a long and prestigious pedigree, it strikes me as an interesting philosophical dilemma, but not an economic dilemma in the last analysis. Ultimately it is a discussion of resource allocation between different time preferences within the same individual, and therefore normative in nature.

While many economists may offer prescriptions for optimal resource allocation, these prescriptions are, and will remain, prescriptive. The open-ended nature of most prescriptive solutions means that the discussion can never conclude, as there will always be a "better" way to allocate resources, and there will always be more than enough draconian suggestions on how this may be done. The approach laid forth in this thesis is however purely descriptive and not appropriate for making value judgments. The only thing we can say is that, whether the individual is or is not in control of his or her emotions and impulses, the variational free energy principle will apply to whoever is acting in the moment, regardless of character and mental fortitude.

6. Conclusion

In this thesis I have sought to answer the following question:

Can the principle of variational free energy be used to explain how agents optimize, while simultaneously allowing for the presence of cognitive biases?

To this end, I have critiqued the expected utility hypothesis and proposed an approach that purports to describe how organic systems optimize in a complex system. I have then attempted to operationalize the underlying principle without reference to equilibrium analysis, and from this perspective, I have proposed some general principles of economics from which further implications can be drawn. These principles can be viewed as axioms of action, but may be incomplete. Lastly, I have attempted to show the extent to which my approach conforms to behavioural economic findings and theories amongst which is *reference dependence, framing effects, the certainty effect* and *the smaller sooner effect*. Furthermore, I have suggested that the apparent problem existing between positive discounting and preferences for improving sequences, is not in fact a problem, but a natural consequences of each other, when action is viewed solely in terms of sequences. This was also used to explain the *peak-end rule*.

The question I sought to answer has therefor been answered to a satisfactory degree. The implication that agents do in fact optimize, and therefore do not exclusively employ heuristics in the decision process, serves to restrain more extraordinary behavioural criticisms of standard economic theory. Furthermore, that utility cannot be represented explicitly according to the variational free energy principle, implies that any attempt to precisely determine the agents' level of rationality or irrationality must be in vain.

Ressources:

Books:

- Hayek, Freidrich A.: The Fatal Conceit – The Errors of Socialism; University of Chicago Press, 1989 (1988).
- Blaug, Mark: The Methodology of Economics – Or How Economists Explain; Second Edition, Cambridge University Press. 1992 (1980).
- Samuelson, Poul A.: Foundations of Economic Analysis; Sixth Printing, Oxford University Press, 1947 (1946).
- Ariely, Dan: Predictably Irrational – The Hidden Forces That Shape Our Decisions; HarperCollins Publishers, 2008.
- Wilkinson, Nick & Klaes, Matthias: An Introduction to Behavioural Economics; Second Edition, Palgrave Macmillan, 2012
- Thaler, Richard H.: Misbehaving – The Making of Behavioural Economics; Penguin Books Ltd, 2015
- Kahneman, Daniel: Thinking Fast and Slow; Penguin Books Ltd, 2011
- Mises, Ludvig von: Human Action – A Treaties on Economics Vol. 1. 2. 3. 4, liberty fund edition, Liberty Fund Inc., 2007.
- Kahneman, Daniel & Slovic, Paul & Tversky, Amos: Judgment Under Uncertainty - Heuristics and Biases; 1st Edition, Cambridge University Press, 1982.
- Squire, Larry & Berg, Darwin & Bloom, Floyd & du Lac, Sascha & Ghosh, Anirvan & Spitzer, Nicholas: Fundamental Neuroscience; Third Edition, Elsevier Inc., 2008.
- Pfaff, Donald W.: The Neuroscience of Fair Play – Why We (Usually) Follow the Golden Rule; Dana Press, The Dana Foundation, 2007.

Articles:

- Rabin, Matthew: A Perspective on Psychology and Economics; European Economic Review 46, pp. 657 – 685, 2002. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.578.6408&rep=rep1&type=pdf

- Camerer, Colin F. & Loewenstein, George: Behavioural Economics – Past, Present, Future; 2002. https://www.researchgate.net/publication/228359684_Behavioral_Economics_Past_Present_Future
- Angner, Erik & Loewenstein, George: Behavioural Economics; 2006 https://www.cmu.edu/dietrich/sds/docs/loewenstein/BehavioralEconomics.pdf
- Posner, Richard A.: Behavioural Law and Economics – A Critique, American Institute for Economic Research, Vol. XLII No. 8, 2002. https://www.aier.org/sites/default/files/Files/Documents/Research/3731/EEB200208.pdf
- Posner, Richard A.: Rational Choice, Behavioural Economics, and the Law; Stanford Law Review Vol. 50:1551, 1998 (1997). https://pdfs.semanticscholar.org/6a45/89ed6bdf74f01bdae91a95d55e17dd81625b.pdf
- Turner, John R. & Baker, Rose M.: Complexity Theory – An Overview with Potential Applications for the Social Sciences; MDPI, 2019 https://www.researchgate.net/publication/272353040_Complexity_Theory
- Neumann, John von & Morgenstern, Oskar: Theory of Games and Economic Behaviour; Princeton University Press, 1947. https://pdfs.semanticscholar.org/0375/379194a6f34b818962ea947bff153adf621c.pdf
- Øvlisen, Frederik R.: Essays in Bounded Rationality and Strategic Interaction; University of Copenhagen, PhD Series No. 137, 2009. https://curis.ku.dk/ws/files/16834576/PDF
- Levin, Jonathan: Choice Under Uncertainty; 2006. https://web.stanford.edu/~jdlevin/Econ%20202/Uncertainty.pdf
- Karni, Edi: Economics of Risk and Uncertainty – Axiomatic Foundations of Expected Utility and Subjective Probability; Vol. 1 Oxford, New Holland, Elsevier B.V., 2014. https://www.econ2.jhu.edu/people/karni/chapter.pdf
- Olshausen, Bruno A.: Bayesian Probability Theory; 2004. http://www.rctn.org/bruno/npb163/bayes.pdf

- Gibson, James J.: The Ecological Approach to the Visual Perception of Pictures; Leonardo Vol. 11 pp. 227 – 235, Pergamon Press Ltd, 1978. http://classes.matthewjbrown.net/teaching-files/ccc/gibson.pdf
- Gregory, Richard L.: Knowledge in Perception and Illusion; 1997. http://www.richardgregory.org/papers/knowl_illusion/knowledge-in-perception.pdf
- Simons, Daniel J. & Chabris, Christopher F.: Gorillas in Our Midst - Sustained Inattentional Blindness for Dynamic Events; Perception Vol. 28, pp. 1059 – 1079, 1999. http://www.chabris.com/Simons1999.pdf
- Charter, Nick & Tenenbaum, Joshua B. & Yuille, Alan: Probabilistic Models of Cognition - Conceptual Foundations; Trends in Cognitive Sciences Vol. 10, No. 7, 2006. https://web.mit.edu/cocosci/Papers/ChaterTenenbaumYuille-TICS06.pdf
- Körding, Konrad P. & Wolpert, Daniel M.: Bayesian Decision Theory in Sensorimotor Control; Trends in Cognitive Sciences Vol. 10, No. 7, 2006. http://cbl.eng.cam.ac.uk/pub/Public/Wolpert/Publications/KorWol_TICS_06.pdf
- Knill, David C. & Pouget, Alexandre: The Bayesian Brain – The Role of Uncertainty in Neural Coding and Computation; Trends in Cognitive Sciences Vol. 27, No. 12, 2004. http://www.wcas.northwestern.edu/nescan/knill.pdf
- Griffiths, Thomas L. & Tenenbaum, Joshua B.: Optimal Predictions in Everyday Cognition; Association for Psychological Science Vol. 17, No. 9, 2006. https://web.mit.edu/cocosci/Papers/Griffiths-Tenenbaum-PsychSci06.pdf
- Bowers, Jeffrey S. & Davis, Colin J.: Bayesian Just-So Stories in Psychology and Neuroscience; Psychological Bulletin Vol. 138, No. 3, pp. 389–414, 2012. https://pdfs.semanticscholar.org/5179/29fddde5b3ddde7ade54b9767b890414fb78.pdf?_ga=2.146891895.1547173656.1559481665-1520320642.1558544498
- Geisler, Wilson S. & Ringach, Dario: Natural Systems Analysis – Introduction; Visual Neuroscience Vol. 26, pp. 1 – 3, 2009. https://www.researchgate.net/publication/24259053_Natural_systems_analysis_Introduction

- Jacobs, Robert A.: Experience-Dependent Integration of Texture and Motion Cues to Depth; Vision Research Vol. 39, pp. 4062 – 4075, 1999. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.212.1920&rep=rep1&type=pdf
- Knill, David C. & Saunders, Jeffrey A.: Do Humans Optimally Integrate Stereo and Texture Information for Judgments of Surface Slant?; Vision Research Vol. 43, pp. 2539 – 2558, 2003. https://www.researchgate.net/publication/5249460_Do_humans_optimally_integrate_stereo_and_texture_information_for_judgments_of_surface_slant
- Battaglia, P.W. & Jacobs, Robert & Aslin, R.N.: Bayesian Integration of Visual and Auditory Signals for Spatial Localization; Journal of the Optical Society of America Vol. 7, pp. 1391 – 1397, 2003. https://www.ncbi.nlm.nih.gov/pubmed/12868643
- Burr, David & Alais, David: Combining Visual and Auditory Information; 2005. http://www.allpsych.uni-giessen.de/rauisch12/readings/Burr.Alais.Review_final.pdf
- van Beers, R.J. & Sitting, A.C. & Gon, J.J.: Integration of Proprioceptive and Visual Position-Information - An Experimentally Supported Model; Journal of Neurophysiology Vol. 3, pp. 1355 – 1364, 1999. https://www.ncbi.nlm.nih.gov/pubmed/10085361
- Friston, Karl J.: The Free-Energy Principle - A Rough Guide to the Brain?; Trends in Cognitive Sciences Vol.13 No.7, 2009. https://www.fil.ion.ucl.ac.uk/~karl/The%20free-energy%20principle%20-%20a%20rough%20guide%20to%20the%20brain.pdf
- Friston, Karl J.: The Free-Energy Principle – A Unified Brain Theory?; Nature Reviews Neuroscience Vol. 11, pp. 127 – 138, 2010. https://www.researchgate.net/publication/41001209_Friston_KJ_The_free-energy_principle_a_unified_brain_theory_Nat_Rev_Neurosci_11_127-138
- Bar-Eli, Michael & Azar, Ofer H.: Penalty Kicks in Soccer - An Empirical Analysis of Shooting Strategies and Goalkeepers’ Preferences; Soccer & Society Vol. 10, No. 2, pp. 183 – 191, 2009. https://www.researchgate.net/publication/263266318_Penalty_kicks_in_soccer_An_empirical_analysis_of_shooting_strategies_and_goalkeepers'_preferences

- Oaksford, Mike & Chater, Nick: The Probabilistic Approach to Human Reasoning; Trends in Cognitive Sciences Vol. 5, No. 8, pp. 349 – 357, 2001. http://www.cogsci.ucsd.edu/~coulson/203/oaksford.pdf
- Cosmides, Leda & Tooby, John: Cognitive Adaptations for Social Exchange; Oxford University Press, 1992. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.140.7186&rep=rep1&type=pdf
- Friston, Karl J. & Kilner, James & Harrison, Lee: A Free Energy Principle for the Brain; Journal of Physiology Vol. 100, pp. 70 – 87, 2006. https://www.fil.ion.ucl.ac.uk/~karl/A%20free%20energy%20principle%20for%20the%20brain.pdf
- Friston, Karl J. & Daunizeau, Jean & Kiebel, Stefan J.: Reinforcement Learning or Active Inference?; PLoS ONE Vol. 4, 2009
- Friston, Karl J. & Daunizeau, Jean & Kiebel, Stefan J. & Kilner, James: Action and Behavior - A Free-Energy Formulation; Biological Cybernetics Vol. 102, pp. 227 – 260, 2010. https://www.fil.ion.ucl.ac.uk/~karl/Action%20and%20behavior%20A%20free-energy%20formulation.pdf
- Friston, Karl J. & Samothrakis, Spyridon & Montague, Read: Active Inference and Agency – Optimal Control Without Cost Functions; Biological Cybernetics, 2012. https://www.fil.ion.ucl.ac.uk/~karl/Active%20inference%20and%20agency%20optimal%20control%20without%20cost%20functions.pdf
- Friston, Karl J. & FitzGerald, Thomas & Rigoli, Francesco & Schwartenbeck, Philipp & Pezzulo, Giovanni: Active Inference – A Process Theory; Neural Computation Vol. 29, 2017. https://www.fil.ion.ucl.ac.uk/~karl/Active%20Inference%20A%20Process%20Theory.pdf
- Friston, Karl J. & FitzGerald, Thomas & Rigoli, Francesco & Schwartenbeck, Philipp & O'Doherty, John & Pezzulo, Giovanni: Active Inference and Learning; Neuroscience and Biobehavioral Reviews Vol. 68, pp. 862 – 879, 2016. https://www.fil.ion.ucl.ac.uk/~karl/Active%20inference%20and%20learning.pdf
- Friston, Karl J. & Buzsáki, Gyorgy: The Functional Anatomy of Time – What and When in the Brain; Trends in Cognitive Sciences Vol. 20, No. 7, 2016. https://www.fil.ion.ucl.ac.uk/~karl/The%20functional%20anatomy%20of%20time.pdf

- Parr, Thomas & Friston, Karl J.: Generalised Free Energy and Active Inference - Can the Future Cause the Past?; Welcome Trust Center for Neuroimaging, 2018. https://www.biorxiv.org/content/biorxiv/early/2018/04/23/304782.full.pdf
- Covey, Dan P. & Cheer, Joseph F.: Accumbal Dopamine Release Tracks the Expectation of Dopamine Neuron-Mediated Reinforcement; Cell Reports Vol. 27, pp. 481 – 490, 2019. https://www.sciencedirect.com/science/article/pii/S2211124719303857
- Schultz, W.: Dopamine Reward Prediction-Error Signalling - A Two-Component Response; Nature Reviews Neuroscience Vol. 17, No. 3, pp. 183 – 195, 2016. https://www.ncbi.nlm.nih.gov/pubmed/26865020
- Coddington, Luke T. & Dudman, Joshua T.: Emergence of Reward Expectation Signals in Identified Dopamine Neurons; BioRxiv, 2017. https://www.biorxiv.org/content/10.1101/238881v1.full
- Coase, Ronald H.: The Problem of Social Cost; The Journal of Law and Economics Vol. 3, 1960. https://www.law.uchicago.edu/files/file/coase-problem.pdf
- Williamson, Oliver: The Economics of Organization – The Transaction Cost Approach; American Journal of Sociology Vol. 87, No. 3, 1981. https://www.researchgate.net/publication/235356934_The_Economics_of_Organization_The_Transaction_Cost_Approach
- Williamson, Oliver: Transaction Cost Economics - An Introduction; Discussion Paper No. 2007-3, 2007. http://www.economics-ejournal.org/economics/discussionpapers/2007-3
- Barsalou, Lawrence W.: Ad Hoc Categories; Memory & Cognition Vol. 11, No. 3, pp. 211 – 227, 1983. http://eprints.gla.ac.uk/110657/
- Epley, N. & Waytz, A. & Cacioppo, J.T.: On Seeing Human - A Three-Factor Theory of Anthropomorphism; Psychological Review Vol. 114, No. 4, pp. 864 – 886, 2007. https://www.ncbi.nlm.nih.gov/pubmed/17907867
- Efremova, N.A. & Inui, T.: An Inferior Temporal Cortex Model for Object Recognition and Classification; Scientific and Technical Information Processing Vol. 41, No. 6, pp. 362 – 369, 2014. https://www.researchgate.net/publication/272370833_An_inferior_temporal_cortex_model_for_object_recognition_and_classification

- Leibo, Joel Z. & Mutch, Jim & Poggio, Tomaso: Why The Brain Separates Face Recognition FromObject Recognition; 2011. http://cbcl.mit.edu/publications/ps/Leibo_Mutch_Poggio_face_invar_v08_cam_rdy_letter_Dec2011.pdf
- Hoffman, Donald D. & Prakash, Chetan: Objects of Consciousness; Frontiers in Psychology, 2014. https://www.researchgate.net/publication/263704213_Objects_of_consciousness
- Rajalingham, Rishi & Issa, Elias B. & Bashivan, Pouya & Kar, Kohotij & Schmidt, Kailyn & DiCarlo, James J.: Large-scale, high-resolution comparison of the core visual object recognition behavior of humans, monkeys, and state-of-the-art deep artificial neural networks; The Journal of Neuroscience Vol. 38, No. 33, pp. 7255 – 7269, 2018. https://www.biorxiv.org/content/10.1101/240614v1
- Kiyatkin, E.A.: Brain Temperature Homeostasis - Physiological Fluctuations and Pathological Shifts; Frontiers in Bioscience Vol. 15, pp. 73 – 92, 2010. https://www.ncbi.nlm.nih.gov/pubmed/20036808
- Andreoni, James & Sprenger, Charles: Certain and Uncertain Utility - The Allais Paradox and Five Decision Theory Phenomena; 2010. http://pubmail.dklevine.com/archive/refs4814577000000000447.pdf
- Kahneman, Daniel & Tversky, Amos: Prospect Theory – An Analysis of Decision Under Risk; Econometrica Vol. 47, No. 2, 1979. https://www.uzh.ch/cmsssl/suz/dam/jcr:00000000-64a0-5b1c-0000-00003b7ec704/10.05-kahneman-tversky-79.pdf
- Kahneman, Daniel & Tversky, Amos: Advances in Prospect Theory - Cumulative Representation of Uncertainty; Journal of Risk and Uncertainty Vol. 5, pp. 297 – 323, 1992. http://cemi.ehess.fr/docannexe/file/2780/tversjy_kahneman_advances.pdf
- Ainslie, George: The Effect of Hyperbolic Discounting on Personal Choices; 2003. https://www.researchgate.net/publication/266574294_The_Effect_of_Hyperbolic_Discounting_on_Personal_Choices
- Laibson, David I.: Hyperbolic Discounting and Consumption; 1994. https://dspace.mit.edu/handle/1721.1/11966
- Laibson, David I.: Golden Eggs and Hyperbolic Discounting; Quarterly Journal of Economics Vol. 112, No. 2, pp. 443 – 477, 1997.

https://dash.harvard.edu/bitstream/handle/1/4481499/Laibson_GoldenEggs.pdf?sequence%253D2

- Frederick, Shane & Loewenstein, George & O'Donoghue, Ted: Time Discounting and Time Preference - A Critical Review; Journal of Economic Literature Vol. XL, pp. 351 – 401, 2002. https://www.researchgate.net/publication/4981445_Time_Discounting_and_Time_Preference_A_Critical_Review
- Thaler, Richard: Mental Accounting and Consumer Choice; Marketing Science Vol. 4, No. 3, 1985. http://bear.warrington.ufl.edu/brenner/mar7588/Papers/thaler-mktsci1985.pdf
- Thaler, Richard & Shefrin, Hersh: An Economic Theory of Self-Control; Journal of Political Economy Vol. 89, No. 2, pp. 392 – 406, 1981. https://www.researchgate.net/publication/24108084_An_Economic_Theory_of_Self-Control