This is the last lecture in this series — so please fill in a feedback form. Thank you. I hope you've enjoyed them.

Overview

We'll examine the basic concepts of behavioural economics and how they can be used to describe and predict behaviour, look at the ways in which humans are irrational, and apply these ideas to the problem of addiction.

'Top down': behavioural economics in the analysis of choice

Last time, we examined motivated behaviour and choice in a 'bottom-up' fashion, looking at several psychological processes that promote action. But we can also analyse it in a 'top-down' manner, ignoring the details of what's going on within the mind and looking at what affects overall behaviour. This is the province of behavioural economics — a merging of traditional economic theory with psychological studies of choice (Rachlin *et al.*, 1976; Allison, 1979).

The rational agent: cornerstone of traditional economic theory

When we consider a subject who must choose between different alternatives, it is natural to attempt to quantify the *value* of the alternatives to the subject. *Utility theory* is a formal way to do this. It starts by assuming attributes of preference that perfectly rational 'agents' should possess (von Neumann & Morgenstern, 1947; Russell & Norvig, 1995). For example, one is *transitivity*: if an agent prefers *A* to *B* and *B* to *C*, then it must prefer *A* to *C*. Suppose some agent violates this principle: it prefers apples to bananas, bananas to cherries, and cherries to apples ($A>B>C>A$). If it starts off with apples, then somebody more intelligent than it could offer it cherries in return for apples plus a small sum of money, and it'd accept. The intelligent soul could then offer bananas in return for cherries plus a further small sum of money — and then apples again in return for bananas plus yet more money. By this time the agent is back in its original state (with apples) but with less money. Assuming money is desirable, this is irrational.

> *'Economics is that way of understanding behaviour that starts from the assumption that people have objectives and tend to choose the correct way to achieve them.' (Friedman, 1990)*

If agents are rational, utility theory tells us that there must exist a *utility function U* that assigns a real number (utility) to every outcome *O* such that $U(O_1) > U(O_2)$ if the agent prefers O_1 to O_2, and $U(O_1) = U(O_2)$ if the agent doesn't care which of them it has. (Economists sometimes refer to 'utiles' as hypothetical units of utility, but you can never measure and quantify utility directly — you can just compare the utility of two different alternatives.) This idea of utility is logically unavoidable, because to choose between two goals that differ in nature, such as food versus money, they must be compared on a single dimension. So if you are able to make a decision between saving one 70-year-old's life and performing 6000 hip replacements, you must be able to value lives and hip replacements in some common way. Utility functions achieve this by converting complicated alternatives to numbers, one number per alternative. Psychologically and neurally, a similar process must also happen (though it might not be a simple process) — if at no earlier stage of processing, incompatible behaviours must compete for access to motor output.

Goal-directed action in an uncertain world

How can a goal-directed being act optimally? Actions produce outcomes. To act in order to obtain one's goals, one must know the value of different outcomes — goals (e.g. obtaining food) have high value, and things to avoid (e.g. being shot) have low value. One must also know the consequences of one's actions (e.g. shopping in a supermarket is likely to obtain food; assaulting a drug dealer is likely to get one shot).

To allow for the fact that actions may not always have totally predictable consequences (is the supermarket out of bread? is the drug dealer armed?), the agent's knowledge about the causal nature of the world may be probabilistic. We could write $p(action \rightarrow outcome_n \mid evidence)$ to denote the probability, given the available evidence, that *action* causes $outcome_n$.

The *expected utility* of an action is therefore given by $EU(\text{action} \mid \text{evidence}) = \sum_n p(action \rightarrow outcome_n \mid evidence) \cdot U(outcome_n)$. A rational agent should select the action with the maximum expected utility; this is known as the MEU principle. This theory doesn't say what your value system (utility functions) should be — anything can be valued, so if you wanted to get shot, you'd assault the drug dealer. Nor does it specify the way that the decision is arrived at. 'Rationality' doesn't require complex logical thought processes — merely that the net behaviour (made up of habits, goal-directed actions, etc., that we talked about last time) obeys rational principles.

The philosophy of economics and behavioural economics

We've assumed that people are rational — that when they are faced with any decision, they make the choice that advances their own purposes most effectively. That is, people maximize their own net benefits. We can turn that idea around: if this is true, then by observing people's choices, we can infer their purposes and value system. If an absolutely rational person chooses A over B, we can infer that their value system values A more than B. This is called the **principle of revealed preference.** It applies to what people do, not what they say. If a smoker tells you he's just had a heart attack, is desperate to give up and is in mortal fear for his life as a result of his smoking, and then smokes a cigarette, an economist would say that the value of smoking was higher to him than the value of not smoking, whatever he says.

> *Two economists walked past a Porsche showroom. One pointed at a shiny car in the window and said, 'I want that.' 'Obviously not,' replied the other.*

This idea — that we should assume that organisms behave optimally (to maximize their net benefit) and that our job is to work out their value system — can also be found in the animal behaviour literature (e.g. Williams, 1994, p. 94).

Furthermore, we must assume that people have *reasonably simple objectives.* With no idea of people's objectives, it is impossible to make any prediction about what they will do. Any behaviour, however peculiar, can be explained by assuming that the behaviour itself was the person's objective. (Why did I stand on my head on the table while holding a burning £50 note between my toes? I *wanted* to stand on my head on the table while holding a burning £50 note between my toes.) (Friedman, 1990.)

Behavioural economics in the analysis of choice

Behavioural economists suggest that animals have a 'bliss point' — an ideal distribution of behaviours that would occur in the absence of restrictions. There may be environmental constraints, such as contingencies imposed by experimenters or the outside world, but within those constraints, animals try to get as close as possible to

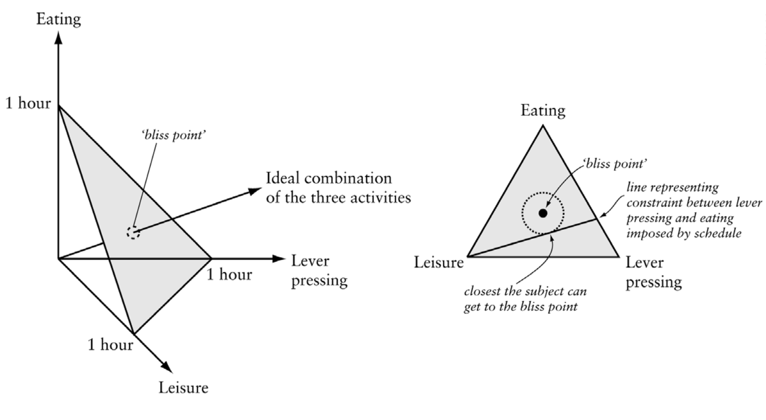

Left: *One hour divided among eating, lever pressing, and leisure. The rat is unconstrained and may choose any point on the grey triangle. Its 'ideal' allocation is shown.* ***Right:*** *Constraints. We take the same triangle but now impose a schedule such that if the subjects wants to eat, it must spend a certain amount of time pressing a lever to earn food. Its behaviour is therefore constrained to the line shown. It can pick any point along that line; the one it chooses depends on its value system, but it will maximize its value (e.g. get as close as possible to the 'bliss point') (Staddon, 1983, ch.7; Rachlin, 2003).*

their bliss point (Rachlin & Burkhard, 1978; Rachlin, 2003) — see figure. This might describe some of the odder results we talked about last time. For example, it could describe Hundt & Premack's (1953) finding that running in a motorized wheel was positively reinforcing when rats weren't running, and negatively reinforcing when they were — perhaps the rats' 'bliss point' was to run at a speed somewhat lower than the motor was set at, and the rats oscillated about this bliss point. And if we can calculate the combination of behaviours that the animal's aiming for, we should be able to predict its behaviour if the constraints change (see Rachlin, 2003).

Basics of economic theory

Economics provides us with more useful tools to understand behaviour. Let's visit some basic concepts first. Economic theory is the theory of trade, or the exchange of goods. Things are termed 'goods' if you prefer more of them to less, all other things being equal (if you like ice cream, then ice cream is a good to you), 'bads' if you prefer less to more, all other things being equal (the smell of warm bloody placenta; eggs), and neutral if you don't care.

The barter economy: price has meaning only in terms of other goods

In the most primitive economy, the barter economy, goods are traded for each other without the use of money (see Friedman, 1990; Rachlin, 2003; Vuchinich & Heather, 2003).

Suppose that Alice has 10 apples and no oranges, and would rather have a *mixture* of fruit. Somebody might ask Alice, 'How many of your apples would you trade for one orange?' Because she has plenty of apples, she might be willing to trade as many as 3 or 4 apples for an orange. Suppose she trades 3 apples for an orange; she is just as happy with her new collection of 7 apples and 1 orange as she was with 10 apples and 0 oranges. She was willing to trade because although apples and oranges may be equally valuable, they are *imperfectly substitutable*. Because Alice prefers to have a mixture of fruit than all one kind, she is willing to have less total fruit if she can have a mixture.

Now suppose that Alice is asked again how many of her remaining apples she is willing to trade for a second orange. Since she already has one orange (she already has a mixture), she would be less willing to part with her apples. Suppose she trades two apples for the second orange. Now she only has 7 pieces of fruit (5 apples + 2 oranges) but she is just as happy as before. As she keeps trading apples for oranges, she becomes less and less willing to part with her diminishing stock of apples; at some point she begins to demand several oranges for each of the (now scarce) stock of apples. Finally, she ends up with 17 oranges and no apples. We can sketch the series of trades in two ways:

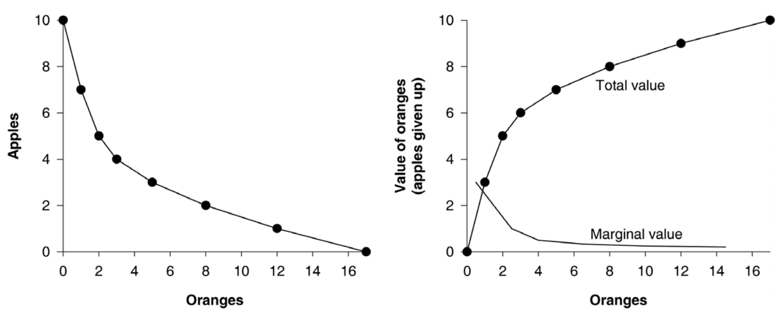

Left: *each point represents a particular combination ('bundle') of apples and oranges. The subject is equally happy with any of these points; the line is called an* ***indifference contour. Right:*** *the same data, expressed as the value of oranges (in terms of apples).* ***Total value*** *is the value of the set of oranges (the number of apples given up to obtain it) — more oranges are worth more.* ***Marginal value*** *(the slope of the total value curve) is the value of the* next *orange — it gets smaller the more oranges you already have.*

We could have plotted many such indifference contours. If Alice has started off with 11 apples, we could have plotted another contour above this one; if she started with 9 apples, we could plot another below it. The shape of the indifference contour depends on the *substitutability* of the commodities. The more substitutable they are, the flatter the indifference contours (if the commodities are completely substitutable, the indifference contour is a flat line). The less substitutable the two commodities, the more bent the indifference contour. The contours for oranges *v*. grapefruit would be flatter than those for oranges *v*. apples; the contours for oranges *v*. chocolate would be more bent (steeper) than those for oranges *v*. apples. We would say that oranges and grapefruits are *substitutes* while oranges and chocolate are *complements*.

In a barter economy, the concept of price has meaning only in terms of the particular goods being traded for each other. The right-hand figure illustrates 'price' in our imaginary series of trades.

A monetary economy

If there are many possible goods we can trade, we might invent money. Money is useful because it is *fungible*; it is a single commodity that is substitutable for almost all others. (When you buy something with money, you are essentially trading what you *could* have bought with that money for what you *do* buy.)

Own-price elasticity

Own-price elasticity measures the change in consumption of a good as its price changes. Suppose that cookies cost £0.10, and I eat 100 cookies per week; this costs me £10 per week. If the price doubles to £0.20, I could do one of several things. I

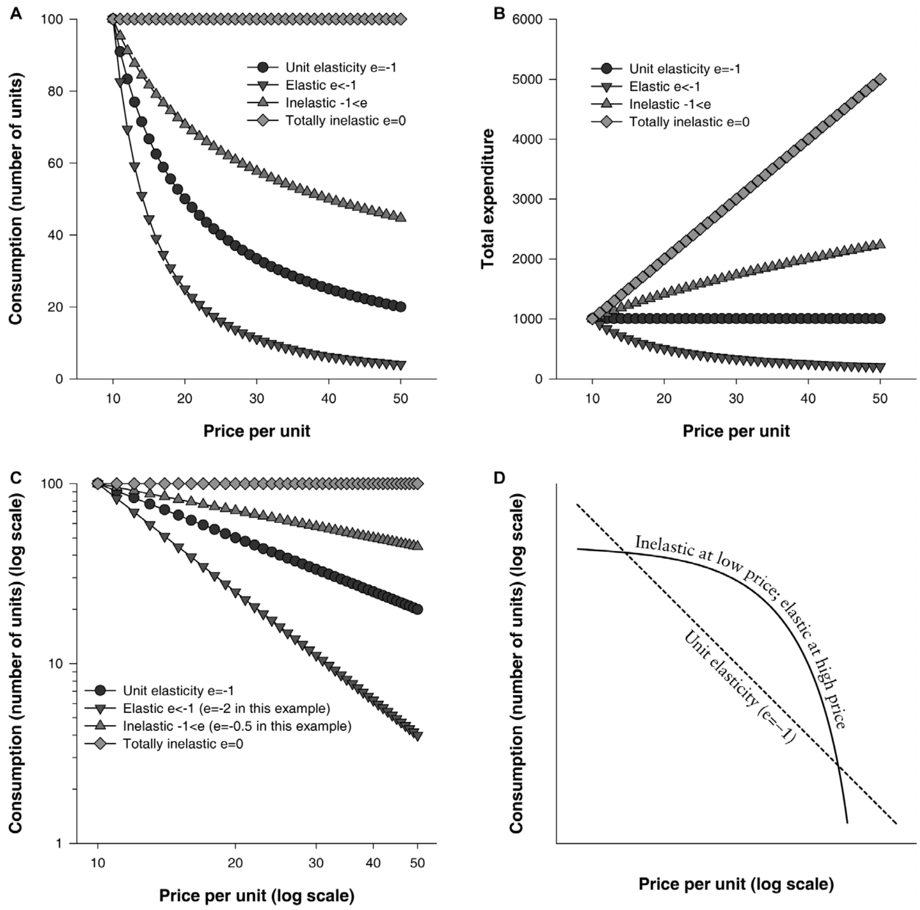

Own-price elasticity of demand. ***A:*** *as the price of cookies goes up, I will probably eat fewer cookies. But how much does the cookie price influence my consumption? Not at all (totally inelastic demand)? Loads (very elastic demand)? Exactly enough to keep the total amount of money I spend the same (unit elasticity)?* ***B*** *shows how my total expenditure would vary.* ***C*** *shows the same data on a log–log plot, in which elasticity is the gradient of the line.* ***D*** *shows a more realistic situation — mixed elasticity. If cookies cost 1p each, how many I eat is unlikely to be affected much by small price changes (inelastic demand — flat line — at low price). If cookies cost £10 each, my consumption will be strongly affected by small price changes (elastic demand — steep line — at high price).*

Aside — only for those who are interested! Economists define elasticity, ε, as $\varepsilon = \frac{\text{\% small change in quantity}}{\text{\% small change in price}} = \frac{\Delta q}{Q} \div \frac{\Delta p}{P} = \frac{\Delta q}{\Delta p} \cdot \frac{P}{Q} = \left(\frac{P}{Q}\right)\frac{\mathrm{d}Q}{\mathrm{d}P}$ where P is price and Q is quantity consumed. Since an increase in price generally makes us consume less, not more (the 'law of demand'), own-price elasticity is typically negative. (Some people reverse the sign of ε so it's usually positive.) But written like this, $\varepsilon = -1$ represents unit elasticity, $\varepsilon < -1$ is elastic demand, $-1 < \varepsilon$ is inelastic demand, $\varepsilon = 0$ is totally inelastic demand ('I eat 100 cookies per week no matter how much I have to pay'), and $\varepsilon > 0$ means that consumption goes *up* when prices goes up. You may recognize that $\frac{\mathrm{d}Q}{\mathrm{d}P} = \frac{\varepsilon Q}{P}$ and therefore $Q = ce^{\varepsilon \ln P}$ where c is a constant (Q when P = 1). Economic textbooks sometimes demonstrate how to calculate elasticities simply by taking the ratio of percentage changes — e.g. if price increases from £9 to £10 and demand decreases from 150 units to 110 units, they'll tell you that the change in price is (10–9) / 9 = 0.1111, and the change in demand is (110–150) / 150 = –0.2667, so elasticity is –0.2667 ÷ 0.1111 = –2.4. However, if we analysed the same change backwards (price £10 → £9, demand 110 → 150), we'd find that elasticity was [(150–110) / 110] ÷ [(9–10) / 10] = –3.636. So this method is simple but not the exact answer. To find the *arc elasticity*, the thing we're actually after, we can approximate — with two points, we'd take the mean P and the mean Q as denominators, giving us ε = [(150–110) / 130] ÷ [(9–10) / 9.5] = –2.92 for this example — or use calculus.

could halve my weekly consumption to 50 cookies, so I continue to spend £10 per week. This would be called *unit elasticity*. I could reduce my consumption by more than this, thinking 'there's no way I'm paying 20p per cookie', and saving money. This would be *elastic* demand — my demand is very sensitive to price. Or I could think 'ho hum', and not cut back very much, say to 90 cookies per week — spending more money in total. This would be *inelastic* demand.

Cross-price elasticity

Cross-price elasticity of demand is the ratio of proportional changes in consumption of one commodity to proportional changes in the price of *another* commodity. It tells us how the commodities are related. Some commodities are *substitutes*, like butter / margarine. If the price of butter goes up, we may buy more margarine (instead) ($\varepsilon > 0$). Some commodities are *complements*, like gin / tonic. If the price of gin goes up, gin-drinkers may buy less tonic (because they buy less gin) ($\varepsilon < 0$). Some commodities are *independent*, like butter / computers, where the price of one doesn't affect consumption of the other ($\varepsilon = 0$).

So in a barter economy, in which the price of goods is measured in terms of other goods, price depends not only on the mutual value of the two goods, but also on their degree of mutual substitutability. There is no such thing as 'the' elasticity of demand for a commodity. The elasticity of demand for butter would depend on whether you measured its price in margarine or computers. In a monetary economy it does make sense to speak of 'the' elasticity of demand for a given commodity — meaning the elasticity of demand for that commodity in terms of money. When you trade a commodity for money, you are trading the commodity for all the other goods that money can buy.

Whew. That's all the economics we need.

Problems with rationality

> *'Economists assume rationality not because it is true but because it is useful; people are in part rational, and it is their rationality that provides the predictable element in their behaviour. This implies that irrationality is not very useful, since it is unpredictable, but not that it does not exist… Unless we can predict not just* that *someone will be irrational, but* which *irrational thing that person will do, we are better off assuming he is rational.' (Friedman, 1990, paraphrased.)*

The rational decision-making approach we've described suffers from two particular *theoretical* deficiencies. Firstly, computing the expected utilities of a set of actions takes time. It may often be better to make an imperfect decision quickly than *eventually* to make what *would* have been the perfect decision — a difficult problem (Russell & Norvig, 1995). Secondly, the principle of maximizing expected utility means that our agent will always do the same thing in the same situation (it is a 'pure' strategy). Once again, this may not be wise in the real world. Game theory (von Neumann & Morgenstern, 1947) has shown that in many situations in which one must make choices in uncertain situations, the best strategy is to assign probabilities to making different choices but to let the actual decision be governed by chance (a 'mixed' strategy). How animals use randomness in decision-making is poorly understood (see Mérö, 1998).

Not only does 'rational choice' have a few *theoretical* problems, it clearly isn't a good description of actual human behaviour! People systematically deviate from the optimum when making decisions (Kahneman *et al.*, 1982; Heckerman *et al.*, 1992; see also Chase *et al.*, 1998). They often seem not to reason rationally; instead, they use heuristics (rough and ready problem-solving rules based on trial and error) and are biased. For example, most people believe themselves to be better-than-average drivers — after all, they drive often without accidents and accidents happen to others (see Slovic *et al.*, 1982)… yet clearly a large number of them must be wrong. This 'it won't happen to me' bias is just one of many (another is shown in the figure).

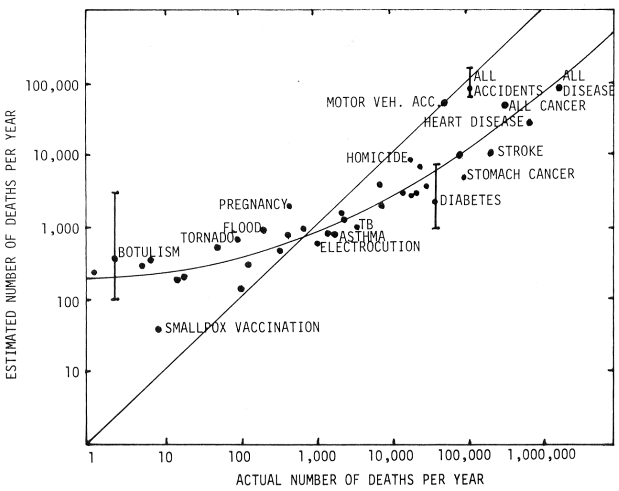

People misperceive risk (Slovic et al., 1982). American subjects were told the true death rate for motor vehicle accidents in the USA — 50,000 per year (Lichtenstein et al., 1978). They were then asked to judge the frequency of 40 other causes of death. Rare causes are overestimated; common causes are underestimated.

Human decision-making behaviour deviates from the ‘rational ideal’ in at least three major ways (Mullainathan, 2002). They exhibit *bounded rationality*: human cognitive abilities are limited. They exhibit *bounded self-interest*: they are frequently willing to sacrifice their own interests to help others. And they exhibit *bounded willpower*: people frequently make choices that aren’t in their long-term interest.

Delayed reinforcement: temporal discounting and impulsivity

Let’s examine the last of these; it’s very important. People often make choices that are advantageous in the short term but disadvantageous in the long term. Why? We can understand this in terms of *temporal discounting of future rewards*. Given a choice between £100 now or £100 in a year, we’d rather have £100 now — we value future rewards less than current rewards (we ‘discount’ the future). Animals and people vary in the extent to which they do this: impulsive individuals value future rewards less than self-controlled individuals (see figure).

One of the most striking experimental findings has been that people do not discount the future in the most sensible way (Ainslie, 1975; 2001). If temporal discounting were rational, the value of future rewards would decline according to an exponential function (never mind the details — the shape is shown in the figure). Exponential discounting generates *time-consistent* choices: if I prefer A to B when I’m choosing five minutes in advance, I’ll also prefer A to B if I choose five years in advance. But animals and people actually exhibit *hyperbolic* temporal discounting (Ainslie, 1975). This generates *time-inconsistent* choices — I might prefer A to B when I’m choosing five minutes in advance, but B to A when I’m choosing five years in advance.

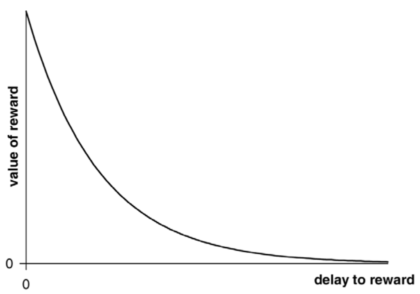

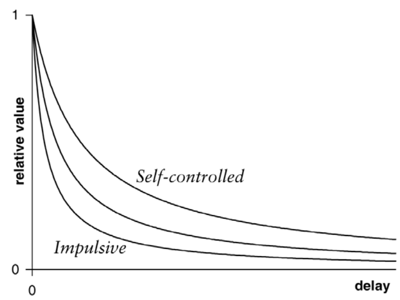

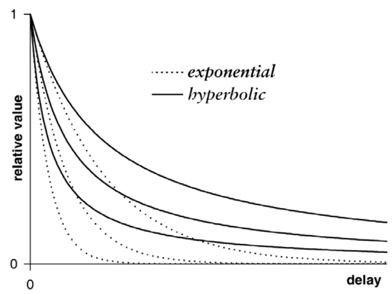

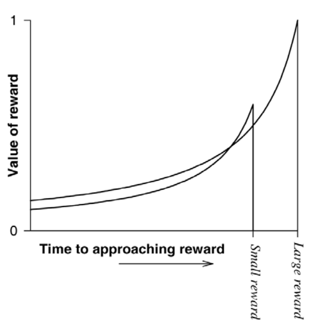

Temporal discounting. ***Top left:*** *the value of a reward declines the longer you have to wait for it — the more it is delayed.* ***Top middle:*** *some individuals don’t value future rewards very much (they discount steeply) and are impulsive. Others value future rewards more, and are self-controlled.* ***Top right:*** *it turns out that animals and people tend to discount the future in a ‘hyperbolic’ way, which isn’t the rational (exponential) way that they perhaps should.* ***Left:*** *this leads to an odd effect: preference reversal. If a subject chooses between a big reward and a small reward when both are a long way in the future, he’ll choose the big one. But as time passes and he gets closer in time to both, there may come a point at which preference reverses, he values the small reward more highly, and he chooses the smaller reward — he acts impulsively.*

This phenomenon is called *preference reversal* (see figure). It may be worth noting that we don't know how this discounting process works in 'bottom-up' terms, and it may be explicable in terms of low-level phenomena such as Pavlovian–instrumental transfer making things more salient and promoting their choice when they're immediately available (Cardinal *et al.*, 2003; Gjelsvik, 2003).

For example, people generally prefer £100 now to £200 in three years. But they also generally prefer £200 in nine years to £100 in six years — yet it's the same choice viewed at a different time. Humans on a diet make calm decisions to avoid desserts (dessert next week or losing weight over the next month?) but then reverse their preference when the dessert trolley is actually under their nose.

If you know that your preference may change in this way, you may be able to improve your happiness in the long run by using *self-control strategies* (Ainslie & Monterosso, 2003). Ainslie (2001) refers to this as bargaining with your future self. Odysseus (Ulysses) knew he would be tempted by the Sirens' song to sail onto the rocks, and knew he would be unable to resist; he had himself bound to his ship's mast in advance, with strict instructions to his crew not to release him no matter his later instructions — an example of pre-commitment (Homer, *Odyssey*). Alcoholics may take disulfiram; while they're taking it, alcohol causes severe illness. People may put social credibility at stake by announcing their diet to everyone, in an attempt to make it harder to have the dessert later — for choosing the dessert then means some loss of face as well. Even pigeons use such self-control strategies. When pigeons choose between smaller-sooner and larger-later rewards, they are often impulsive (Rachlin & Green, 1972). But they'll actually work to avoid being offered the smaller-sooner alternative (Ainslie, 1974; Ainslie & Herrnstein, 1981)!

How can we apply the ideas we've covered to a real problem — that of addiction?

Drug addiction — abnormal motivation?

What's addiction, anyway?

Addiction is often thought of as being involuntary. For example, the diagnostic criteria for alcohol or other drug dependence (APA, 2000) boil down to a *compulsion* to take the drugs. However, drug use can be voluntary, and one can debate whether even full-blown addicts take drugs involuntarily (Schaler, 2000; Skog, 2003): just because someone says they don't want to smoke and then later smokes doesn't mean they're smoking involuntarily — it might simply be that they're inconsistent. Furthermore, not everyone who smokes wants to give up. The figure shows one classification of addicts.

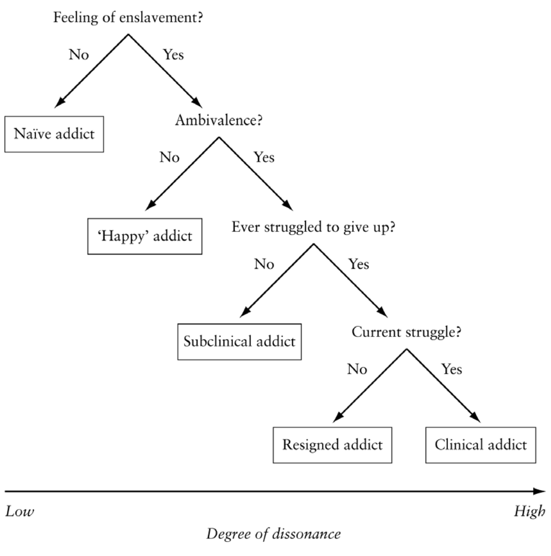

Skog's (2003) view of addiction. A person may be unaware that it is difficult for him or her to live without a drug. They are enslaved, but unaware; Skog calls them 'naïve' addicts. He offers the example of a heavy drinker in Paris in World War II, who had never realised that he was dependent on alcohol until rationing came along and he was limited to one litre of wine per week. Then there are those who know that life would be harder without, but are happy with this situation: 'happy' addicts, such as the 1950s smoker who thought that smoking was good for you (or at least, not bad). Those who aware smoking is bad for you but feel no particular motivation to cut back are called 'subclinical' addicts by Skog. Finally, there are those who have tried and failed but aren't trying at the moment, and those in an active struggle to quit.

Most drugs are not irresistible, contrary to some popular folklore; for all those 'proper' (stereotypical) addicts who use extremely large amounts of heroin, cocaine, and alcohol, there are many others who use infrequently — sometimes referred to as 'chipping' (MacCoun, 2003). So addiction is not an all-or-nothing phenomenon. Furthermore, most addicts recover (Warner *et al.*, 1995; Heyman, 2003).

Addiction as inelastic demand

An obvious way to think about addiction is that demand for drugs is *inelastic* compared to demand for other things. The more someone is addicted, the more inelastic their demand is — they will therefore sacrifice other commodities (work, money, social interaction) rather than sacrifice drug. Indeed, alcohol demand in rats can be more inelastic than demand for food (Heyman *et al.*, 1999; Heyman, 2000). While this is a useful way to think about addiction, it is not an all-or-nothing phenomenon. The elasticity of demand for food is of course not constant: animals are more inelastic when they're hungrier, and if there are no alternative ways of obtaining food (e.g. Hursh, 1978). The elasticity of demand for drugs can vary in the same way: demand for cigarettes is more inelastic when smokers have been abstinent (Madden & Bickel, 1999).

Is addiction rational?

It has been suggested that addiction is a completely rational choice made by people with consistent preferences (Stigler & Becker, 1977; Becker & Murphy, 1988). However, this doesn't seem very satisfactory. As Winston (1980) puts it:

> *'[T]he addict looks strange because he sits down… surveys future income, production technologies, investment / addiction functions, and consumption preferences over his lifetime… maximizes the discounted value of his expected utility, and decides to be an alcoholic. That's the way he will get the greatest satisfaction out of life. Alcoholics are alcoholics because they want to be alcoholics,* ex ante*, with full knowledge of [the] consequences.'*

In any case, we know that people's preferences are not consistent over time, and therefore they do not maximize their total expected reward — that is one of the central implications of hyperbolic temporal discounting (Ainslie & Monterosso, 2003). For example, cigarette taxes can make smokers happier (Gruber & Mullainathan, 2002). This implies that addiction is not 'rational' — addicts' preferences are not consistent over time, and so cigarette taxes make smokers happier because they serve as a valuable self-control device, helping them to avoid smoking.

A bottom-up view: drugs change your brain, and the way you think or choose

Many addiction theories focus on the way in which drugs change the brain. As Kelley & Berridge (2002) recently noted, drugs may activate the same circuits as natural rewards, perhaps in a more potent manner; they may create new states, such as the motivational state of withdrawal; and / or they may differentially affect the balance of processes that normally contribute to responding for natural rewards (such as habits, goal-directed actions, and cue-induced motivation; we talked about this last time). There may be other effects, too: for example, food makes you full and exercise makes you tired, but some drugs may not satiate you to the same extent (Heyman, 2003). Acute intoxication impairs decision-making, so the decision to have the sixth pint may not be made in exactly the same way as the decision to have the first. Chronic use of some drugs may alter the brain so as to impair your ability to make good choices (e.g. Rogers *et al.*, 1999). Some forms of brain damage may make you more impulsive and more likely to choose small, immediate rewards over larger, delayed rewards (e.g. Cardinal *et al.*, 2001).

A top-down view: drug use changes your value system; your preferences are inconsistent over time

We could also look at addiction from the top down: how drugs influence behaviour in a particular environmental context. A key concept in theories of addiction is that addictive behaviour, such as consumption of alcohol beyond a certain rate, *reduces the value of future activities*. For example, alcohol consumption may reduce the value of future alcohol consumption (e.g. tolerance), but may also reduce the value of other things, such as social activity. This is suggested to lead people down a 'primrose path' to addiction (see figure).

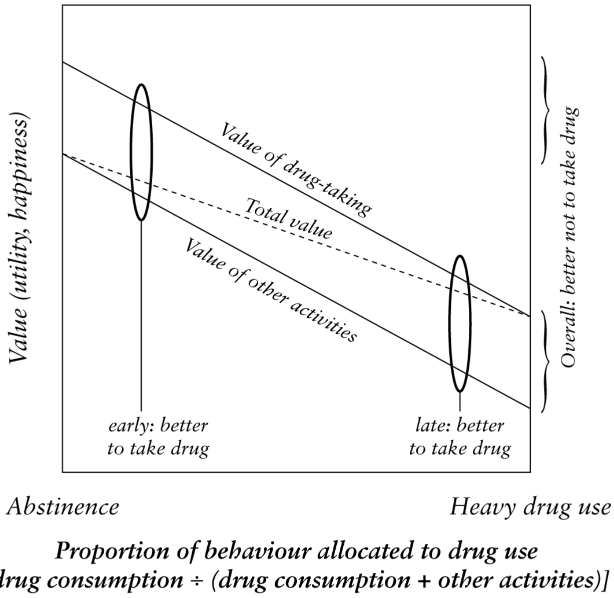

Good now, bad in the long run — the 'primrose path' to addiction. At any point, drug-taking has a higher value than other activities, so you take the drug. But drug-taking lowers both the value of future drug-taking (e.g. alcohol consumption → tolerance → future alcohol isn't worth as much) and the value of other activities (e.g. the more alcohol you consume, the less you socialize and the worse you are at socializing; the more heroin you take, the worse you are at your job). So as you drink more, your total happiness goes down — you'd be better off not being an alcoholic. But even when you are an alcoholic, drinking now is worth more than not drinking now — for you are sensitive to local, not global, utility. (After Herrnstein & Prelec, 1992; Rachlin, 1997; Rachlin, 2003; Vuchinich & Heather, 2003.) As Rachlin (2000) puts it: 'The alcoholic does not choose to be an alcoholic. Instead he chooses to drink now, and now, and now, and now. The pattern of alcoholism emerges in his behaviour… without ever having been chosen.'

For this to happen, people have to make decisions that are advantageous in the short term but disadvantageous in the long term; hyperbolic discounting supplies a mechanism. In fact, drug addicts may discount the future more steeply (and therefore be even more impulsive and short-termist) than non-addicts (Ainslie & Monterosso, 2003; Bickel & Johnson, 2003; Mitchell, 2003; Vuchinich & Heather, 2003). Similar short-termism can explain relapse (Heyman, 2003). Since one cigarette doesn't cause cancer and one shot of heroin doesn't condemn you to a junkie lifestyle, a person can correctly reason that since it's 'just for one last time', the drug is the better choice. But a series of 'one-last times' turns into a relapse.

Drug addiction — what can we do about it?

It's important to recognize that since addiction is not an all-or-nothing problem, we shouldn't necessarily search for a 'cure' for addiction — just methods to reduce consumption. If that reduces it to zero, fine; if not, we can pursue other methods to reduce it further should we want to.

Prohibition: the economics of illegal drugs

Many drug policy analyses focus on *prevalence* — the number of people who use a drug. MacCoun (2003) argues that the appropriate analysis is that of *total harm reduction*. For example, if total social harm = average harm per use × number of users × average amount used, then by focusing on prevalence we may neglect the strategies of reducing the average quantity used and reducing the amount of harm per use.

We spend a lot of money controlling illicit drug use. The USA federal drug budget is over $19 billion per year (USA, 2002). The UK spends about £1 billion per year on programmes to deal with illegal drugs (UK, 2000), of which £380 million is spent on reducing drug availability and £400 million on treatment. If demand is relatively constant, then restricting supply pushes up the price. Lab cocaine is about £90 per gram (Sigma-Aldrich, Feb 2004), and would undoubtedly be much cheaper if it were manufactured in larger quantities — it's not used much in medicine. Street cocaine is about £60 / g and 20–65% pure (e.g. Streatfeild, 2001), so cocaine itself costs £92–£300 / g on the street. Injectable medicinal heroin, which is manufactured in respectable amounts and known coyly as diamorphine in hospitals, is presently about £40 / g (www.bnf.org, Feb 2004); heroin in tablet form costs the NHS only £12 / g, and the production cost is less than £1 / g (Shaw, 2000). Street heroin was £60 / g in 2002 and is 20–90% pure, so heroin itself costs £67–£300 / g on the street, and it's never sterile. These kinds of profit margins explain why the world-wide illegal drug trade is worth some £400 billion per year (Keh, 1996). Of course, the 'price' of street drugs should also include the costs of criminal activity associated with drug use. A serious heroin addict has to find about £50 per day to pay for heroin. In the USA, addiction treatment programs save about $42,000 (~£22,000) per treated addict per year in the costs of crime and the criminal justice system (McCollister & French, 2003), so the criminal cost of drug use is presumably at least this. About 30% of those arrested in the UK are dependent on an illegal drug (Shaw, 2000).

Pushing up the price reduces demand for illicit drugs (Saffer & Chaloupka, 1995) — which can be seen as a valid goal of drug policy. However, we've seen that demand for drugs is often fairly inelastic ($|\varepsilon| < 1$) at current prices. If price goes up and demand is inelastic, the total amount spent on drugs simply goes up. For those willing to take the risks of supplying drugs, the profits can be very large (Streatfeild, 2001), and the UK intercepts only ~20% of imported drugs (Shaw, 2000). What would happen if drug policy were reversed — to legalize drugs such as heroin and cocaine? Well, we can indulge in some informed speculation. In scenario A, it becomes a totally free market: the costs of smuggling vanish, supply increases, the price plummets, companies vie to sell you cheap heroin (perhaps the state imposes some quality control). If you take the view that the state shouldn't interfere with the right of an individual to take whatever drugs he or she pleases as long as he / she doesn't harm others, you would support this. However, usage would increase substantially.

In scenario B, licensed suppliers are regulated and unlicensed supply remains criminal. Supply would increase; the price would drop. Legal sources would compete with the black market, so the final price of legal drugs would have to be near or below the lowest price that's profitable for the black marketeers. Policing illegal supply (catching smugglers) costs money but keeps the black market price up, which is a good thing. Why is a good thing? Well, assuming that the final legal price is higher than the cost of producing the drugs (which is likely in a state-sanctioned production environment!), the state could generate income by taxing the sale of the drugs, to increase the price again up to the maximum that's still competitive with the black market. So the final market price to addicts is lower than before, but not vastly lower. Because of the reduced economic incentive for black market dealers, drug supply-related crime would drop; because of the reduced price to addicts, drug purchasing-related crime would drop. Health costs from impure drugs would drop. The last two reasons are why the UK allows the controlled prescribing of heroin to addicts (which works; e.g. Uchtenhagen, 1997).

In either scenario, since price will drop, consumption will almost certainly go up — both through increased consumption in existing users and consumption in new users — so the direct health costs of drug use will increase. This is certainly a reason to think carefully about whether or not to legalize drugs. Saffer & Chaloupka (1995) estimate that legalizing and taxing illicit drugs would lead to a 100% increase in heroin consumption and a 50% increase in cocaine consumption. At the moment, opiates cause about 1,000 deaths per year in England and Wales, and cocaine causes about 80 deaths, while alcohol causes 5,000–40,000 (estimates vary widely) and tobacco causes about 120,000 across the UK (*Hansard*, 17 July 2002). When you treat addicts, the gain in health costs is small (about $2,000 per addict per year) compared to the gain in crime reduction (about $42,000 per addict per year) (McCollister & French, 2003) — and certainly a reasonable proportion of those health costs are to do with treating infections associated with drug use, which should be less of a problem if sterile drugs are supplied, rather than to do with direct effects of the drugs. And inelasticity works two ways — if drug demand is inelastic, a given proportional drop in price will produce a smaller proportional increase in consumption. It has been estimated that the USA could raise $3–17 billion per year in tax on drugs that are presently illegal; along with reduced policing and criminal justice costs, the net benefit would be to the tune of $24 billion per year (see Shaw, 2000). Once you've saved the money, you can use it for whatever you think most valuable — treatment strategies for unhappy addicts, hip replacements, whatever.

Taxation: the economics of legal drugs

Criminalizing legal drugs is another option, but the arguments above suggest that this may not be very sensible. Direct experience bears this out — the Prohibition of alcohol in 1920s USA certainly reduced overall consumption, but not by much and not for long; price soared, expenditure on alcohol soared, major crime soared, the death rate from impure alcohol shot up, there were few clear health benefits observed, and people disliked the policy intensely (Thornton, 1991).

Taxation is the obvious alternative. Each year, alcohol misuse is estimated to cost the UK perhaps £11 billion, of which up to £3 billion is spent by the NHS treating alcohol-related diseases (UK, 2003). In contrast, alcohol taxation generates about £11 billion (Smith, 1999). Tax on cigarettes currently accounts for 79% of the £4.59 pack price (April 2003), and this generates about £9.5 billion. The NHS spends £1.5 billion treating smoking-related diseases (Parrott *et al.*, 1998), so this is an example where taxation is appropriately high enough to cover the financial costs of addiction to the state — and higher. Do these taxes reduce consumption?

Yes. The elasticity of demand for cigarettes is typically about –0.4 (Gruber *et al.*, 2002; Chaloupka *et al.*, 2003) — that is, if the price goes up by 10%, consumption goes down by 4%. This is for two reasons. First, when price goes up, some people quit altogether (*participation elasticity*). Second, people who continue to smoke smoke less (*conditional elasticity of demand*, in the jargon). And as for most commodities, elasticity varies with price: smokers working for cigarette puffs in the laboratory are fairly inelastic when the price is low (ε = –0.56), but become more elastic when the price goes up (ε = –1.58) (Bickel *et al.*, 1995b; DeGrandpre & Bickel, 1995). Probably for this reason, elasticity is greater for poorer smokers, for whom cigarettes are proportionally more expensive (Gruber *et al.*, 2002). The UK elasticity of demand for alcohol varies from –1.69 for wine (pretty elastic), through –0.86 for spirits (somewhat inelastic), to –0.76 for beer (Smith, 1999).

Finally, if we tax tobacco and alcohol in order to reduce consumption, we must consider whether taxing one drug increases consumption of another. For this we need to know their cross-price elasticities. Wine and spirits are substitutes, so increasing the price of spirits makes people drink more wine, but beer and wine are complements, and so are beer and spirits (Smith, 1999). And cigarettes and alcohol are either independent or complements (Gruber *et al.*, 2002), so reducing the consumption of one either reduces the consumption of the other or doesn't affect it. Similarly, one can work out which other drugs of abuse substitute for each other, complement each other, or substitute for / complement non-drug reinforcers (Bickel *et al.*, 1995a).

Treating individual addicts

So there are perfectly respectable arguments to the effect that as long as there's demand for an addictive drug, society as a whole does best by legalizing and taxing it. (And if there isn't any demand, it doesn't matter what society does.) This is, of course, at the expense of individual addicts, who may increase in number but who pay for their health care and everybody else's hip replacements (or whatever else the surplus tax is spent on). And it's their free choice… certainly for our 'naïve', 'happy', or 'subclinical' addicts. But addicts who don't want to be addicted are a different matter. It's certainly often in their health interests to stop; smokers and alcoholics die early and it's often a very unpleasant death. How can we reduce demand? How can we change their value system so that they consume less drug?

Price, price, price

We've seen that one way is to increase the price — be it by increasing tax on cigarettes (financial), drug availability (practical), stigmatizing smokers (social), banning smoking in public places (legal). It works. Increasing the price of alcohol reduces alcohol consumption and the consequences of alcoholism, such as liver cirrhosis, deaths caused by drunk drivers, and violent crimes (see Chaloupka *et al.*, 2002). Increasing the price of cigarettes decreases consumption, both in controlled laboratory studies and the real world (e.g. Keeler *et al.*, 1993; Madden & Bickel, 1999; Chaloupka *et al.*, 2003). Conversely, if the price of a drug goes down — for example, if people get richer through economic growth but alcohol prices do not go up correspondingly — consumption increases. What else can we do to reduce it?

Money can't buy everything: providing alternatives

We said that money was useful because we can trade it for almost all other things. But money can't buy everything, and sometimes we don't have money. Rats in an operant chamber don't have money. When a rat works for food by pressing a lever, you could say that it trades leisure for food — the more leisure it gives up (the more it presses the lever), the more food it gets. You could say that the price of the food is leisure given up. Since this is a barter economy, the price depends on the relative values of food and leisure, *but also on their substitutability*. And since leisure is what the rat does when it's not doing other things, this will also depend on what else is on offer (such as drinking water, running in a wheel, sniffing sawdust, and so on).

This line of thought can be applied to the treatment of addiction (Rachlin, 2003). If you focus on the financial price of cigarettes, you are focusing on substituting *things that you can buy with money* for cigarettes. Smoke less, and you can buy more cinema tickets. But there are also many substitutes for addictive activities that cannot be immediately bought with the money you save by giving up your addiction. If you stop smoking, you can't go out and buy social support with the money you save. So making it easier for an addict to obtain substitutes for drugs may be as effective as making it harder for the addict to obtain drugs (McCollister & French, 2003; Rachlin, 2003). Indeed, this strategy works (Green & Fisher, 2000). You can also reward abstinence directly, with money or other tangible rewards, and this also promotes abstinence (see Heyman, 2003). Finally, you can use self-control techniques like pre-commitment to improve your sensitivity to the long term (Ainslie, 2001).

Reducing the value of drugs directly

There are a variety of pharmacological ways to reduce the value of specific drugs. For example, methadone treats opiate withdrawal symptoms, reducing the value of heroin, and it reduces the 'high' of any heroin that you take — again, reducing the value of heroin. Nicotine patches treat nicotine withdrawal, reducing the value of nicotine. Disulfiram alters the metabolism of ethanol so that if you drink while you're taking disulfiram, you feel very ill indeed; thus, disulfiram reduces the value of alcohol. Vaccination against cocaine is being tried at the moment; this reduces the 'high' and therefore the value of cocaine. All of these can be seen as self-control tactics: because the addict would prefer a drug-free lifestyle in the long term, he deliberately adopts a strategy (e.g. taking disulfiram) that reduces the future value of his drug. It may also be possible to target the brain's motivational systems directly; chemicals that reduce drug-seeking in animals (e.g. Pilla *et al.*, 1999) may be another line of therapy.

Better knowledge of the risks of drug-taking is another way that reduces the perceived value of drugs (Heyman, 2003). But we know that humans are poor at using risk information (Slovic et al., 1982). So effective advertising of risk is intended to take advantage of human reasoning biases — that your perception of the risk of getting atherosclerosis from smoking is enhanced by vivid images of atheroma being squeezed out of arteries. Misinformation may also help people avoid addiction: a personal theory that cocaine use inevitably leads to full-blown destructive addiction might not be true, but it might be a useful view to hold, because this view of cocaine is a self-control device that prevents you taking any cocaine (Ainslie, 2001).

Summary

Motivated behaviour is subject to economic influences, and can be analysed in economic terms, even though humans are not perfectly rational agents. Specific irrationalities — such as hyperbolic temporal discounting of future reward — lead to preferences that are inconsistent over time. These concepts provide an understanding of addiction, and tools with which to treat it, at an individual and a societal level.

All references cited in the handout

I'm not suggesting that you read these! They are here as pointers to the original literature; so if you are for some reason keen to read more, or if you disagree with something I've claimed, you can check for yourself.

Ainslie, G. (1974). Impulse control in pigeons. *Journal of the Experimental Analysis of Behavior* **21:** 485-489.

Ainslie, G. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological Bulletin* **82:** 463-496.

Ainslie, G. (2001). *Breakdown of Will*, Cambridge University Press, Cambridge, UK.

Ainslie, G. & Herrnstein, R. J. (1981). Preference reversal and delayed reinforcement. *Animal Learning and Behavior* **9:** 476-482.

Ainslie, G. & Monterosso, J. (2003). Hyperbolic discounting as a factor in addiction: a critical analysis. In *Choice, Behavioral Economics and Addiction* (Heather, N. & Vuchinich, R. E., eds.), pp. 35-61, 67-69. Elsevier, Oxford.

Allison, J. (1979). Demand economics and experimental psychology. *Behavioral Science* **24:** 403-415.

APA (2000). *Diagnostic and Statistical Manual of Mental Disorders, fourth edition, text revision (DSM-IV-TR)*, American Psychiatric Association, Washington DC.

Becker, G. S. & Murphy, K. M. (1988). A theory of rational addiction. *Journal of Political Economy* **96:** 675-700.

Bickel, W. K., DeGrandpre, R. J. & Higgins, S. T. (1995a). The behavioral economics of concurrent drug reinforcers: a review and reanalysis of drug self-administration research. *Psychopharmacology (Berl)* **118:** 250-259.

Bickel, W. K., DeGrandPre, R. J., Higgins, S. T., Hughers, J. R. & Badger, G. (1995b). Effects of simulated employment and recreation on drug taking: a behavioral economic analysis. *Experimental and Clinical Psychopharmacology* **3:** 467-476.

Bickel, W. K. & Johnson, M. W. (2003). Junk time: pathological behavior as the interaction of evolutionary and cultural forces. In *Choice, Behavioral Economics and Addiction* (Heather, N. & Vuchinich, R. E., eds.), pp. 249-271, 276-278. Elsevier, Oxford.

Cardinal, R. N., Pennicott, D. R., Sugathapala, C. L., Robbins, T. W. & Everitt, B. J. (2001). Impulsive choice induced in rats by lesions of the nucleus accumbens core. *Science* **292:** 2499-2501.

Cardinal, R. N., Robbins, T. W. & Everitt, B. J. (2003). Choosing delayed rewards: perspectives from learning theory, neurochemistry, and neuroanatomy. In *Choice, Behavioral Economics and Addiction* (Heather, N. & Vuchinich, R. E., eds.), pp. 183-213, 217-218. Elsevier, Oxford.

Chaloupka, F. J., Emery, S. & Liang, L. (2003). Evolving models of addictive behaviour: from neoclassical to behavioral economics. In *Choice, Behavioral Economics and Addiction* (Heather, N. & Vuchinich, R. E., eds.), pp. 71-89. Elsevier, Oxford.

Chaloupka, F. J., Grossman, M. & Saffer, H. (2002). The effects of price on alcohol consumption and alcohol-related problems. *Alcohol Res Health* **26:** 22-34.

Chase, V. M., Hertwig, R. & Gigerenzer, G. (1998). Visions of rationality. *Trends in Cognitive Sciences* **2:** 206-214.

DeGrandpre, R. J. & Bickel, W. K. (1995). Human drug self-administration in a medium of exchange. *Experimental and Clinical Psychopharmacology* **3:** 349-357.

Friedman, D. D. (1990). *Price Theory: An Intermediate Text*, South-Western.

Gjelsvik, O. (2003). Reason and addiction. In *Choice, Behavioral Economics and Addiction* (Heather, N. & Vuchinich, R. E., eds.), pp. 219-238, 245-247. Elsevier, Oxford.

Green, L. & Fisher, E. B. (2000). Economic substitutability: some implications for health behavior. In *Reframing health behavior change with behavioral economics* (Bickel, W. K. & Vuchinich, R. E., eds.), pp. 115-144. Erlbaum, Mahwah, NJ.

Gruber, J. & Mullainathan, S. (2002). Do cigarette taxes make smokers happier? [Working Paper 8872]. National Bureau of Economic Research.

Gruber, J., Sen, A. & Stabile, M. (2002). Estimating price elasticities when there is smuggling: the sensitivity of smoking to price in Canada [Working Paper 8962]. National Bureau of Economic Research.

Heckerman, D. E., Horvitz, E. J. & Nathwani, B. N. (1992). Toward normative expert systems: Part I. The Pathfinder project. *Methods of Information in Medicine* **31:** 90-105.

Herrnstein, R. J. & Prelec, D. (1992). A theory of addiction. In *Choice Over Time* (Loewenstein, G. & Elster, J., eds.), pp. 331-361. Russell Sage Press, New York.

Heyman, G. M. (2000). An economic approach to animal models of alcoholism. *Alcohol Res Health* **24:** 132-139.

Heyman, G. M. (2003). Consumption dependent changes in reward value: a framework for understanding addiction. In *Choice, Behavioral Economics and Addiction* (Heather, N. & Vuchinich, R. E., eds.), pp. 95-121. Elsevier, Oxford.

Heyman, G. M., Gendel, K. & Goodman, J. (1999). Inelastic demand for alcohol in rats. *Psychopharmacology (Berl)* **144:** 213-219.

Hundt, A. G. & Premack, D. (1953). Running as both a positive and negative reinforcer. *Science* **142:** 1087-1088.

Hursh, S. R. (1978). The economics of daily consumption controlling food- and water-reinforced responding. *Journal of the Experimental Analysis of Behavior* **29:** 475-491.

Kahneman, D., Slovic, P. & Tversky, A., Eds. (1982). *Judgement Under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press.

Keeler, T. E., Hu, T. W., Barnett, P. G. & Manning, W. G. (1993). Taxation, regulation, and addiction: a demand function for cigarettes based on time-series evidence. *J Health Econ* **12:** 1-18.

Keh, D. I. (1996). Drug money in a changing world: economic reform and criminal finance. United Nations Office on Drugs and Crime.

Kelley, A. E. & Berridge, K. C. (2002). The neuroscience of natural rewards: relevance to addictive drugs. *Journal of Neuroscience* **22:** 3306-3311.

Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M. & Combs, B. (1978). Judged frequency of lethal events. *Journal of Experimental Psychology: Human Learning and Memory* **4:** 551-578.

MacCoun, R. (2003). Is the addiction concept useful for drug policy? In *Choice, Behavioral Economics and Addiction* (Heather, N. & Vuchinich, R. E., eds.), pp. 383-401, 407. Elsevier, Oxford.

Madden, G. J. & Bickel, W. K. (1999). Abstinence and price effects on demand for cigarettes: a behavioral-economic analysis. *Addiction* **94:** 577-588.

McCollister, K. E. & French, M. T. (2003). The relative contribution of outcome domains in the total economic benefit of addiction interventions: a review of first findings. *Addiction* **98:** 1647-1659.

Mérö, L. (1998). *Moral Calculations: Game Theory, Logic, and Human Frailty*, Springer-Verlag, New York.

Mitchell, S. H. (2003). Discounting the value of commodities according to different types of cost. In *Choice, Behavioral Economics and Addiction* (Heather, N. & Vuchinich, R. E., eds.), pp. 339-357. Elsevier, Oxford.

Mullainathan, S. (2002). Behavioral economics. In *International Encyclopedia of the Social & Behavioral Sciences* (Baltes, P. B. & Smelser, N. J., eds.). Pergamon, Oxford.

Parrott, S., Godfrey, C., Raw, M., West, R. & McNeill, A. (1998). Guidance for commissioners on the cost effectiveness of smoking cessation interventions. Health Educational Authority. *Thorax* **53 Suppl 5 Pt 2:** S1-38.

Pilla, M., Perachon, S., Sautel, F., Garrido, F., Mann, A., Wermuth, C. G., Schwartz, J. C., Everitt, B. J. & Sokoloff, P. (1999). Selective inhibition of cocaine-seeking behaviour by a partial dopamine D3 receptor agonist. *Nature* **400:** 371-375.

Rachlin, H. (1997). Four teleological theories of addiction. *Psychonomic Bulletin & Review* **4:** 462-473.

Rachlin, H. (2000). *The science of self-control*, Harvard University Press, Cambridge, Massachusetts.

Rachlin, H. (2003). Economic concepts in the behavioral study of addiction. In *Choice, Behavioral Economics and Addiction* (Heather, N. & Vuchinich, R. E., eds.), pp. 129-149. Elsevier, Oxford.

Rachlin, H. & Burkhard, B. (1978). The temporal triangle: Response substitution in instrumental conditioning. *Psychological Review* **85:** 22-47.

Rachlin, H. & Green, L. (1972). Commitment, choice and self-control. *Journal of the Experimental Analysis of Behavior* **17:** 15-22.

Rachlin, H., Green, L., Kagel, J. & Battalio, R. (1976). Economic demand theory and psychological studies of choice. In *The Psychology of Learning and Motivation* (Bower, G., ed.), pp. 129-154. Academic Press, New York.

Rogers, R. D., Everitt, B. J., Baldacchino, A., Blackshaw, A. J., Swainson, R., Wynne, K., Baker, N. B., Hunter, J., Carthy, T., Booker, E., London, M., Deakin, J. F., Sahakian, B. J. & Robbins, T. W. (1999). Dissociable deficits in the decision-making cognition of chronic amphetamine abusers, opiate abusers, patients with focal damage to prefrontal cortex, and tryptophan-depleted normal volunteers: evidence for monoaminergic mechanisms. *Neuropsychopharmacology* **20:** 322-339.

Russell, S. J. & Norvig, P. N. (1995). *Artificial Intelligence: a modern approach*, Prentice-Hall, Upper Saddle River, New Jersey.

Saffer, H. & Chaloupka, F. J. (1995). The demand for illicit drugs [Working Paper 5238]. National Bureau of Economic Research.

Schaler, J. A. (2000). *Addiction is a choice*, Open Court Publishing, Chicago, Illinois.

Shaw, B. (2000). UK Drug Policy Report to the Cleveland Police Authority. Cleveland Police.

Skog, O.-J. (2003). Addiction: definitions and mechanisms. In *Choice, Behavioral Economics and Addiction* (Heather, N. & Vuchinich, R. E., eds.), pp. 157-175, 182. Elsevier, Oxford.

Slovic, P., Fischhoff, B. & Lichtenstein, S. (1982). Facts versus fears: understanding perceived risk. In *Judgement Under Uncertainty: Heuristics and Biases* (Kahneman, D., Slovic, P. & Tversky, A., eds.), pp. 463-489. Cambridge University Press, New York.

Smith, Z. (1999). The revenue effect of changing alcohol duties. Institute for Fiscal Studies.

Staddon, J. E. R. (1983). *Adaptive Behavior and Learning*, Cambridge University Press, Cambridge.

Stigler, G. & Becker, G. S. (1977). De gustibus non est disputandum. *American Economic Review* **67:** 76-90.

Streatfeild, D. (2001). *Cocaine*, Virgin, London.

Thornton, M. (1991). Alcohol prohibition was a failure [Cato Policy Analysis No. 157]. The Cato Institute.

Uchtenhagen, A. (1997). Summary of the Synthesis Report. Institute for Social and Preventive Medicine at the University of Zurich.

UK (2000). *Spending Review 2000: New Public Spending Plans 2001-2004*, HM Treasury, London.

UK (2003). Alcohol misuse: how much does it cost? Cabinet Office Strategy Unit.

USA (2002). *Federal Drug Control Programs: The Budget for Fiscal Year 2003*, Executive Office of the President, Washington, DC.

von Neumann, J. & Morgenstern, O. (1947). *Theory of games and economic behavior*, Princeton University Press, Princeton, New Jersey.

Vuchinich, R. E. & Heather, N. (2003). Introduction: overview of behavioural economic perspectives on substance use and addiction. In *Choice, Behavioral Economics and Addiction* (Heather, N. & Vuchinich, R. E., eds.), pp. 1-31. Elsevier, Oxford.

Warner, L. A., Kessler, R. C., Hughes, M., Anthony, J. C. & Nelson, C. B. (1995). Prevalence and correlates of drug use and dependence in the United States. Results from the National Comorbidity Survey. *Archives of General Psychiatry* **52:** 219-229.

Williams, B. A. (1994). Reinforcement and choice. In *Animal Learning and Cognition* (Mackintosh, N. J., ed.), pp. 81-108. Academic Press.

Winston, G. C. (1980). Addiction and backsliding: a theory of compulsive consumption. *Journal of Economic Behavior and Organization* **1:** 295-324.