THEORETICAL AND METHODOLOGICAL PERSPECTIVE

An improved demand curve for analysis of food or drug consumption in behavioral experiments

Mark Newman[1] · Carrie R. Ferrario[2]

Abstract

The incorporation of microeconomics concepts into studies using self-administration procedures has provided critical insights into the factors that influence consumption of a wide range of food and drug reinforcers. In particular, the fitting of demand curves to consumption data provides a powerful analytic tool for computing objective metrics of behavior that can be compared across a wide range of reward types in both human and animal experiments. The results of these analyses depend crucially on the mathematical form used to fit the data. The most common choice is an exponential form proposed by Hursh and Silberberg, which is widely used and has provided fundamental insights into relationships between cost and consumption, but it also has some disadvantages. In this paper, we first briefly review the use of demand curves to quantify the motivating effects of food and drugs, then we describe the current methodology and highlight some potential issues that arise in its application. To address these issues, we propose a new mathematical framework for the analysis of consumption data, including a new functional form for the demand curve. We show that this proposed form gives good fits to data for a range of different reinforcers and experimental protocols, while allowing for straightforward calculation of key metrics of demand, including preferred consumption level, maximum response, price at maximum response, and price elasticity of demand. We provide software implementing our entire analysis pipeline, including data fits, data visualization, and the calculation of demand metrics.

Keywords Self-administration · Motivation · Addiction · Reward · Behavioral economics · Reinforcement

Introduction

Behavioral economics is a branch of microeconomic theory and practice that aims to understand and quantify the economic choices made by individuals (Ayers and Collinge 2004; Kahneman 2011; Perloff 2016; Cartwright 2018). In recent years, methods adapted from behavioral economics have been fruitfully applied to improve our understanding of the motivating effects of food and drugs, both in humans and also in animals, particularly in the context of self-administration procedures (Hursh 1993; Winger et al. 2002; Hursh et al. 2005; Winger et al. 2006; Hursh and Silberberg 2008; Galuska et al. 2011; Bentzley et al. 2013; Kawa et al. 2016; Pantazis et al. 2019). Although animals do not have an economy in the sense of markets, currency, credit, and so forth, in many cases, they are willing to "pay" for goods such as food or drugs by performing work. By measuring the amount of work they are willing to do for a given return, one can develop notions of "price" and "demand" for goods, on which the tools of economics can then be brought to bear.

The utility of economic concepts in the study of food and drug motivation and addiction has been demonstrated extensively since their introduction to the field in the 1990s (Hursh 1993; Hursh and Winger 1995; Bickel et al. 2000; Hursh and Silberberg 2008). Foundational contributions by Hursh, Winger, Bickel, and others have established behavioral economics approaches as a fundamental tool for the quantitative assessment of self-administration behavior. The approach has been widely applied to evaluate motivation across different

✉ Carrie R. Ferrario
ferrario@umich.edu

1 Center for the Study of Complex Systems, University of Michigan, Ann Arbor, MI, USA

2 Department of Pharmacology, University of Michigan, Ann Arbor, MI, USA

reinforcers and testing conditions, to study effects of pharmacological or behavioral interventions on demand, and to address abuse liability and neurobehavioral underpinnings of substance abuse.

The behavioral economics approach as currently applied focuses on the study of the so-called *demand curve*, which represents the consumption of food or drug as a function of the price paid or work performed to obtain it. Such curves are computed by fitting an appropriate mathematical form to raw self-administration data, and aspects of the fitted curve provide summary statistics describing features such as the maximum effort exerted to obtain a given reinforcer and the price at which that maximum effort occurs. Most previous work has made use of an exponential form for the demand curve proposed by Hursh and Silberberg (2008), but, as we discuss, this form and the accompanying methodology have some shortcomings that create challenges for comparisons across studies and for reproducibility of results.

In this paper, we present a new approach for fitting demand data that addresses these shortcomings while providing easy calculation of key metrics of preferred consumption, price sensitivity, and motivation. We begin with a brief introduction to the relevant concepts of economics followed by a description of the current methodology. We then describe our proposed approach, outlining its main features and providing a simple recipe for data analysis that gives straightforward quantitative answers to questions relevant to the study of demand, motivation, and addiction. We couch our discussion primarily in the language of animal self-administration experiments, but, in principle, the same analysis approaches could be used for human consumption data without modification.

Economics of consumption

Consider a typical self-administration experiment in which an animal is given the opportunity to perform work, such as lever presses, in return for a desirable good. In most cases, the good is either food or drug; for the purposes of illustration, let us say it is a dose of drug. The "price" of the drug can be varied by the experimenter, usually using one of two methods: either they can vary the number of lever presses required to receive a fixed dose or they can vary the dose received for a fixed number of lever presses. Either way, one can define the price P of the drug as the number of lever presses per unit of drug received, measured for instance in milligrams. Thus:

$$\text{Price}, P = \frac{\text{Number of lever presses}}{\text{Milligrams of drug received}}. \quad (1)$$

(See Table 1 for a summary of the variables used in this paper.)

In the most straightforward version of the experiment, the experimenter allows the animal to "buy" repeated doses of drug at a set price and records the total amount of drug Q consumed during a session of fixed length (anywhere from a few minutes to hours, depending on the drug, the question at hand, and so forth). The procedure can then be repeated for a range of different prices to measure consumption Q as a function of price P.

The data produced by experiments of this kind have a characteristic form. First, there is normally a clear preferred amount of drug that a particular animal will consume in the allowed time when price is not an issue. Even if the price is reduced practically to zero so that drug is essentially free, the animal will not consume an unlimited amount but will stop when it reaches its preferred level of consumption. Traditionally, this level is denoted Q_0.

Next, if we now raise the price slightly, so that drug is not free but still very cheap, the animal will still be willing to do the modest work required of it and will take its fill of drug, meaning its consumption will still be Q_0. But if we raise the price enough, the effort will start to become a factor and the animal will consume less drug. And if the price is very large—if the animal has to do a million lever presses, say, to receive a dose of drug—then consumption must be zero, since it is physically impossible to perform this many lever presses in the allowed amount of time.

Thus, when plotted against price, we expect consumption to look something like Fig. 1a. The data points in this figure show actual consumption against price for a rat self-administering cocaine.[1] Observe how the points are roughly flat in the left part of the plot but fall off beyond a certain point, denoted approximately by the vertical dashed line, as the price rises. If one were to continue the measurements far enough to the right of the plot, they would eventually reach zero when the price becomes so high that the rat receives no drug at all. Note that the graph is plotted on logarithmic scales, a standard practice that allows us to capture the typically wide range of values of both price and consumption.

The demand curve

Data of the kind shown in Fig. 1a can be summarized by fitting a *demand curve* to it. An example is shown as the solid curve in Fig.1a; it represents the expected consumption level Q as a function of price P. The demand curve has become a standard tool in the analysis of consumption data, capturing in a single graph the willingness of an animal to work for a range of outcomes. The utility of this approach has been demonstrated repeatedly and demand curve analyses have provided essential information

---

[1] Previously unpublished data kindly provided by A. B. Kawa, L. Longyear, and T. E. Robinson.

Table 1 Summary of the variables used in the theory described here, along with alternative notations for the same quantities used by other authors

Variable	Meaning	Other notations
P	Price, work performed per unit of good received	C
Q	Consumption	
R	Revenue, total work performed at a given price point	O
Q_0	Preferred consumption level when price is negligible	
R_{max}	Maximum work performed to obtain goods at any price	O_{max}
P_{max}	Price at which maximum work is performed	
$\widetilde{P}_{max}$	Normalized P_{max}, equal to $Q_0 P_{max}$	P_{max}, nP_{max}
α	Fitted parameter in exponential model, inversely proportional to P_{max}	
$\tilde{\alpha}$	Normalized α, equal to α/Q_0, also called "essential value," inversely proportional to $\widetilde{P}_{max}$	α

about many aspects of consumption and drug-taking behavior in both humans and animal models (see, for example, Hursh and Silberberg (2008); Bentzley et al. (2013); Aston and Cassidy (2019)). As we will see, the demand curve allows us to describe motivation for drugs or food in a quantitative manner, placing numbers on concepts that may otherwise be accessible only via more qualitative approaches, and allowing comparisons across different reinforcers (see Hursh and Winger (1995) and Hursh et al. (2005) for reviews). To do this, we borrow some concepts from economics, starting with the so-called elasticity of demand.

Elasticity of demand

Suppose that we know the demand curve for a particular experiment—the solid line running through the data points in Fig. 1a. (We will see shortly how to extract such curves from data.) In general, the demand curve falls off as price increases, since we expect an animal to consume less of a good as the work required to obtain it increases (the so-called *law of demand*). The *price elasticity of demand E* measures exactly how the demand curve falls off with increasing price. For instance, if we double the amount of work a rat has to do to receive a dose of drug, will the rat consume the same amount of drug overall? Half as much? A quarter?

The elasticity is the ratio between the fraction the price goes up by and the fraction the consumption goes down by. For example, if price goes up by 10% and as a result consumption falls by −20%, then the elasticity is $E = -20/10 = -2$. Note that elasticity is normally a negative number, as here.

More generally, suppose that the price P increases by an amount dP. Then the *fraction* that price increases by
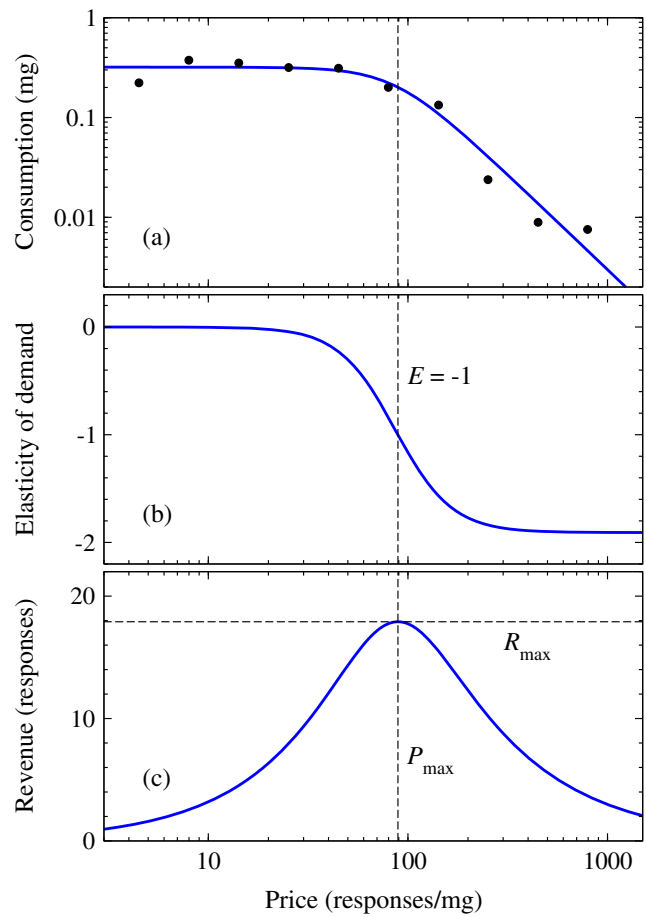

Fig. 1 **a** Data points: total consumption of self-administered cocaine by a single male rat as a function of price P measured in lever presses per milligram. Solid curve: the demand curve reconstructed by fitting to Eq. (21). **b** The elasticity corresponding to the demand curve in **a**. Note that the value of the elasticity is always negative and that the point where elasticity equals −1 coincides with the point at which the revenue (plotted in **c**) reaches its maximum. **c** The "revenue" corresponding to the demand curve in **a**, i.e., the total work performed at each price point. The maximum revenue R_{max} falls at price P_{max} which coincides with the point at which elasticity is −1

is dP/P. If at the same time consumption Q goes down by dQ then the fractional decrease in consumption is dQ/Q. The elasticity is the ratio of these two fractions (fractional decrease in consumption over fractional increase in price), which is

$$E = \frac{dQ/Q}{dP/P} = \frac{P}{Q}\frac{dQ}{dP} \qquad (2)$$

The quantity dQ/dP is the derivative of consumption with respect to price.[2] The derivative is the slope of a graph of Q against P. Because the graph is downward sloping in this case, the slope is negative, and hence again E will be a negative number.

---

[2] Technically, the derivative is the limiting value of dQ/dP for infinitesimal dP and dQ.

Alternatively, we can note that

$$\frac{d\log Q}{d\log P} = \frac{dQ/Q}{dP/P} = E, \tag{3}$$

which is the same elasticity again, meaning that E is also equal to the slope of a graph of $\log Q$ against $\log P$. In other words, if we plot the demand curve on log scales (as in Fig. 1a), then the elasticity is the slope.

Figure 1 b shows the elasticity for the demand curve of Fig. 1a. Note that, in general, elasticity is not a single number: it varies with price. In the left part of Fig. 1a, for instance, the demand curve is flat and hence the elasticity in Fig. 1b is close to zero. But, in the right part, the demand curve slopes downward quite steeply, meaning the elasticity is large and negative.

The elasticity is widely used in (human) economics as a measure of the price sensitivity of goods, including psychoactive substances (Ayers and Collinge 2004; Perloff 2016). A large (negative) elasticity of demand indicates a good that is highly price sensitive: small increases in price will substantially decrease demand. A small elasticity indicates a relatively price-insensitive good. For instance, the elasticity of soft drinks has been measured to be about -3.8 at prevailing prices (Ayers and Collinge 2004), indicating substantial price sensitivity—if the price of a soft drink is increased, people will simply stop drinking it. On the other hand, the elasticity of cigarettes is estimated to be much smaller, around -0.4 (Becker et al. 1994), indicating significantly lower price sensitivity—people will continue to smoke even if the price of cigarettes goes up.

These results suggest that elasticity could be used as a measure of the reinforcing or motivating effects of drugs or food, not only in humans but also in animal experiments. In practice, however, it is rarely used in this way in the animal literature. Instead, elasticity has primarily been of interest because of its role in estimating the maximum work that animals perform, as we discuss in the next section. (There are claims in the literature of using elasticity to quantify motivation, but in most such cases no value of the elasticity is actually reported. Instead, researchers typically report derived measures such as the price at which maximum work is performed, as we discuss in the "Parameters for the exponential demand curve" section.)

Measuring motivation

One of the primary uses of demand curves in animals is for quantifying responding for food or drug. How motivated are animals to take a drug? Can we define a single number to quantify motivation? How does motivation change over time? How does it compare between different reinforcers? One approach is to look at the total amount of work an animal is willing to perform over the course of an experimental session. In the jargon of economics, this total amount of work is called the revenue, denoted R, although in the present context, it may be more useful to think of the R as standing for "responses," since it is simply equal to the number of lever presses or other work the animal performs.[3]

We have seen that the price P is defined as the amount of work that must be performed to receive one unit of the desired good, such as the number of lever presses per milligram of drug (see Eq. (1)). If we know the number of lever presses per milligram and we also know the total number of milligrams consumed Q, then the total number of lever presses—the revenue, R—must be the product of the two:

$$R = PQ. \tag{4}$$

Figure 1 c shows a plot of the revenue corresponding to the demand curve in Fig. 1a. Observe how the revenue starts at a low value on the left of the plot in the regime where the animal is required to perform only a little work to receive drug, rises to a maximum around the turning point in the demand curve, then falls off again as the price becomes too high and the animal abandons trying to obtain drug. These observations suggest two possible measures of the motivation potential of food or drugs: (1) the maximum amount R_{max} of work—the maximum revenue—the animal is willing to commit to obtaining goods at any price,[4] or (2) the price P_{max} at which this maximum occurs. In Fig. 1c, R_{max} corresponds to the height of the peak in revenue (the horizontal dashed line) and P_{max} corresponds to the price at which that peak occurs (the vertical dashed line). Perhaps more intuitively, P_{max} corresponds roughly to the turning point in the demand curve, Fig. 1a, at which the curve falls off from its initial plateau. Thus, P_{max} measures the maximum price the animal will tolerate before it gives up and reduces its consumption.

To calculate P_{max}, we maximize Eq. (4) with respect to P by differentiating thus:

$$\frac{dR}{dP} = \frac{d(PQ)}{dP} = Q + P\frac{dQ}{dP} = Q\left[1 + \frac{P}{Q}\frac{dQ}{dP}\right] = Q(1+E) \tag{5}$$

where E is the elasticity as before and we have used Eq. (2). Setting Eq. (5) to zero, we then find that the maximum revenue is achieved when the elasticity E is equal to -1, and thus, we can find P_{max} by finding the price at which this occurs. We give examples of this calculation in the "Maximum work performed for the exponential demand curve" and "Elasticity and measures of motivation" sections. Once we have determined P_{max}, we can substitute the result back into Eq. (4) to find the corresponding value of the revenue, R_{max}.

---

[3] The notation R for revenue is a standard one in economics, but, in the animal literature, one sometimes sees revenue denoted O for "output." See also Table 1.

[4] Sometimes also denoted O_{max}

Ease of calculating P_{max} and R_{max} is one of the reasons why the demand curve and the elasticity are useful concepts. In principle, one could imagine measuring P_{max} and R_{max} directly from the data by asking what the maximum number of responses is at any price and at what price that maximum falls. Ad hoc methods for doing this have been proposed, for instance, by Hursh and Silberberg (2008) (see also Oleson et al. (2011)). The results returned by these methods, however, are limited to the specific values of price and responding measured in the experiment and so give only a general indication, and they are moreover prone to measurement fluctuations and hence can be unreliable. Calculations based on demand curves are more robust and repeatable (Bentzley et al. 2013).

Both R_{max} and P_{max} are reasonable measures of motivation, but they are not equivalent. If two animals display the same P_{max} for a given reinforcer but the first has higher R_{max}, it implies that the first is willing to do more work than the second for the same amount of consumption. Conversely, if they display the same R_{max} but the first has higher P_{max} then the first is willing to do the same amount of work for less consumption. Certainly, these two measures could be correlated, but they are not the same thing.

R_{max} has found use in human studies, where strong associations have been observed between its value and, for instance, post-intervention alcohol consumption (MacKillop and Murphy 2007; MacKillop et al. 2009), although it should be noted that there were also strong relationships in these studies among P_{max}, R_{max}, and other demand metrics (see "Example calculations" and Fig. 6 for further discussion). Hursh and Winger (1995) have argued that R_{max} is essentially independent of the potency or magnitude of a reinforcer, and for this reason, support its use for comparisons between different reinforcers. The value of P_{max}, by contrast, varies with potency and thus can be difficult to compare across reinforcers, although one can by suitable normalization create a potency-independent version of P_{max} that remedies this issue, as discussed in the "Normalized price and comparisons between different reinforcers" section. At the same time, R_{max} can be difficult to measure in some situations, particularly when using progressive ratio schedules or when drugs have sedative or stimulant effects on motor performance, whereas P_{max} is relatively easy to measure (for additional discussion, see Richardson and Roberts 1996). Thus, both P_{max} and R_{max} can be useful and both are widely employed in the field, although arguably P_{max} is more common. Which to adopt in any individual situation will depend on the details of the experiment.

Normalized price and comparisons between different reinforcers

In addition to their use in basic data analysis, demand curves are used as a way to compare behavior across experiments on different drugs or other reinforcers. Can we tell, for instance, whether animals have greater motivation for food or drugs? Or for one drug over another?

One way to perform such comparisons is to use the measure R_{max} defined in the "Measuring motivation" section above, which is the maximum work animals are willing to perform to obtain a good, at any price. Hursh and Winger (1995) argue that R_{max} is well suited to comparisons between different reinforcers and describe it as "a sensitive tool for direct comparison and quantitative ordering of demand, both within and across the drug classes (stimulant, sedative, and opioid)."

Conversely, the measure P_{max} (also defined in "Measuring motivation"), which is the price at which animals exert their greatest effort to obtain goods, is not well suited to answering such questions because it is not clear how one should compare prices for different goods. Is a price of 10 lever presses per milligram of cocaine higher or lower than 10 lever presses per milligram of amphetamine? The answer depends on the potency of the drugs in question: 1 mg might have a strong effect for one drug but only a weak effect for another. To make a meaningful comparison, we need to normalize the price by a suitable factor that represents the typical magnitude of drug intake in milligrams (or other suitable units). Fortunately, we have exactly such a factor to hand, namely the preferred consumption level Q_0.

If we divide dose by Q_0, we get a number that is independent of potency: every dose is specified as a fraction of the preferred consumption level for the same good. Thus, a suitable normalized price, as first proposed by Hursh and Winger (1995), is given by replacing the dose in Eq. (1) by dose divided by Q_0:

$$\text{Normalized price}, \widetilde{P}$$
$$= \frac{\text{Number of lever presses}}{(\text{Milligrams of drug received})/Q_0} \qquad (6)$$
$$= Q_0 \frac{\text{Number of lever presses}}{\text{Milligrams of drug received}} = Q_0 P.$$

Combining this approach with our measure, P_{max}, we can then write a potency-independent measure of motivation thus:

$$\widetilde{P}_{max} = Q_0 P_{max} \qquad (7)$$

This measure appears to work well in practice for comparing demand across different reinforcers.

Equations (6) and (7) are not in precisely the form used elsewhere. In much of the literature, for instance, the price is denoted C instead of P and the normalized price $\widetilde{P}$ is denoted P. This, however, is merely a matter of notation. A more substantive difference is that many experimenters avoid the use of P_{max} as a measure of motivation in favor of another parameter commonly denoted α. We discuss α in detail below ("Maximum work performed for the exponential demand curve"), where we show that it is in fact essentially equivalent to P_{max} but has some disadvantages that P_{max} does not share.

Fitting the demand curve

Data such as those shown in Fig. 1a already give us a rough outline of the demand curve. But they also inevitably display statistical fluctuations and moreover give us the consumption at only a small set of discrete price points. We can reduce the effects of fluctuations and interpolate between price points by fitting a suitable curve through the data, recovering the entire demand curve, as shown by the blue line in Fig. 1a. The fitting procedure itself is straightforward—there are a range of software packages that will do the job. A crucial question, however, is what mathematical form the fitted curve should take. We need to specify a form that is flexible enough to fit the data we see in experiments on a range of different consumables, different animals, and different procedures, while at the same time following the common-sense requirements that the curve be flat at first, then drop off, and go to zero as price becomes large.

The exponential demand curve

The most common mathematical form used for demand curves in self-administration experiments of the kind considered here is the "exponential" form advanced by Hursh and Silberberg (2008):

$$Q = Q_o e^{k(e^{-\alpha p}-1)} \tag{8}$$

This equation relates the consumption Q to the price P using three parameters. The first is Q_0, which we have already discussed—it is the preferred level of consumption when the price is low enough to have no limiting effect on intake. The other two parameters are k and α, which we look at more closely in the following section.

Because demand curves are normally plotted on logarithmic scales, one often sees Eq. (8) expressed in terms of the logarithm of Q. Taking the log of both sides of the equation we find that

$$\log Q = \log Q_0 + k\left(e^{-\alpha P}-1\right), \tag{9}$$

where "log" denotes the natural logarithm (base e).[5] The two forms, Eqs. (8) and (9), are entirely equivalent and contain the same information.

Parameters for the exponential demand curve

Though it has been widely employed, the demand curve defined in Eq. (8) has some shortcomings. Specific issues include difficulty estimating or interpreting the parameters k

and α, difficulty estimating the elasticity, and an unrealistic nonzero value of consumption at large prices.

The parameter k and the limiting value of Q A disadvantage of Eq. (8) is that it fails to meet one of our fundamental criteria for a demand curve, that it go to zero as price becomes large. As we have said, it is axiomatic that the curve should go to zero: if the price of drug is a million lever presses per dose then the animal is necessarily going to consume no drug. As P goes to infinity, however, the form in Eq. (8) tends to the limiting value $Q_0 e^{-k}$, which is always nonzero.

This conflict causes a number of problems. First, it is undesirable to fit a curve to data when we know the curve to have a different shape from the data. Just as one should not fit a straight line if one knows the data to follow a rounded form, so one should not fit data that must go to zero with a form that does not. This is, however, a somewhat theoretical objection. A more practical issue arises when we attempt to estimate the parameter k. Since k controls the value of consumption when price becomes large, one should be able to determine k by measuring this value. This, however, is not possible in the present case, since as we have said there is no such value in practice: consumption always goes to zero.

There are ways around this difficulty. One could for instance estimate k using some other type of fitting procedure. Hursh and Silberberg (2008) take a different route, avoiding the problem by not fitting k to the data at all. Instead, they choose its value themselves, writing that "the value of k is generally set to a common constant across comparisons because it merely specifies the range of the data." In related work Bentzley et al. (2013) write that the value of k is "chosen based on the maximum observed range of consumption." Specifically, they compute the range spanned by the observed values of log Q for all sessions and set k equal to the largest such range. Gilroy et al. (2019) suggest a slight variant of this procedure, calculating the same maximum range but then adding 0.5, to guard against the possibility that k "does not reflect the full range of observed consumption values."[6]

In practice, however, approaches such as these are somewhat unsatisfactory because they determine k using an ad hoc recipe rather than by fitting to the data. Such recipes can result in different experimenters using different values of k—and hence reaching different conclusions— even when fitting to the same set of data points. Consider Fig. 2, which shows the same cocaine self-administration data that appears in Fig. 1a. The four curves in the figure show the best fit of the exponential demand curve to these data for four different values of k. The blue (solid) curve

---

[5] Some authors have employed common logarithms (base 10) instead, in which case some of the formulas and parameter values are modified, but the scientific outcome is unchanged

[6] The value of k must also always be greater than the mathematical constant $e = 2.718$ for P_{max} to be well defined (Gilroy et al. 2019). If k is smaller than this then there is no point on the demand curve where elasticity is -1 and hence, there is no P_{max}—see "Measuring motivation".
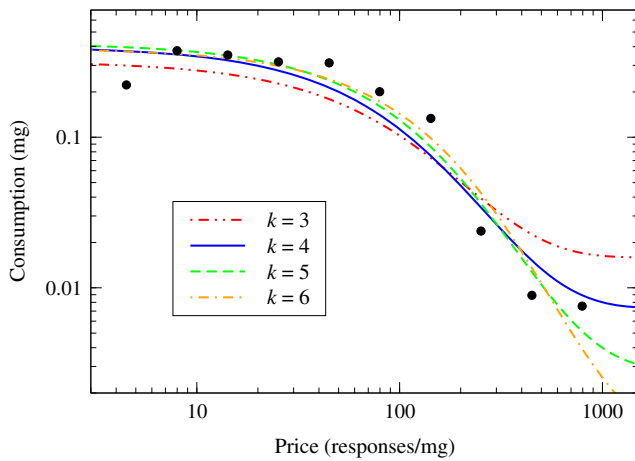
Fig. 2 Best fits of the cocaine self-administration data from Fig. 1a to the exponential demand curve of Eq. (9) for various choices of the parameter k as indicated

shows the fit when k is chosen according to the prescription of Bentzley et al. that k should be set to the largest range spanned by $\log Q$ over all sessions. Since we are looking at only a single session in this case, k is simply equal to the range of the data, which gives a value of about 4. (The exact value is 3.91, but we use 4 for simplicity.)

On the other hand, if we were examining these data as one session out of many, it is likely that at least one other session would have a larger range of $\log Q$, meaning that we would have to use a larger value of k. The green and orange (dashed and dot-dashed) curves in Fig. 2 show fits with $k = 5$ and 6. For comparison, we also show one fit with a smaller value of $k = 3$ (red, dot-dot-dashed).

The values of the parameters of the fit for each curve are shown in Table 2. As we can see, the values span quite a wide range. The value of Q_0 varies from 0.319 to 0.420, for example, an increase of 32%, and $P_{\max}$ shows a similar increase of 38%. Also shown in the table is the quantity we call $\tilde{\alpha}$, which is described in "Maximum work performed for the exponential demand curve" and which is widely used as a measure of motivation in the literature. The fitted values of this quantity vary over a broad range from a low of 0.0047 to a high of 0.0149, an increase of 217%. The quantity $\widetilde{P}_{\max} = Q_0 P_{\max}$, which we recommended in the "Normalized price and comparisons between different reinforcers" section as a measure of motivation, fares better, varying from 37.5 to 43.6, an increase of just 16%, but even this variation is large enough to inject significant uncertainty into the results, given that it is provoked solely by making different choices for the parameter k.

The parameter α The parameter α is also somewhat problematic, although for different reasons. In principle, this parameter is unexceptionable: it plays the role of an exponential constant and can be estimated in straightforward fashion by fitting to

data such as those in Fig. 1a. The problem is that in much of the literature α is said to be a measure of elasticity and it has been extensively treated as such (e.g., Koffarnus et al. (2011); Bentzley et al. (2013); Lacy et al. (2019)). As discussed in "Elasticity of demand", elasticity is a well-defined and useful concept in economics, but α is not a measure of elasticity. Rather, α is a measure of the price at which the animal performs maximum work, equivalent to the quantity $P_{\max}$ discussed in "Measuring motivation" above. This point is made clearly in the literature, but, at the same time, misinterpretation of α is also common. For example, Bentzley et al. (2013) state correctly that "α is an inherently normalized parameter and equivalent to [normalized $P_{\max}$], as these variables are inversely proportional", but also that "α is a measure of demand curve elasticity." To some extent, this may be a matter of semantics—the scientific conclusions are largely unaffected—but, in the interests of clarity, it is good to be precise about the role of the variables.

A further issue with α is that its value is strongly influenced by the choice of k. In the following sections, we examine the elasticity for the exponential demand curve in detail and explain the role played by α.

Elasticity for the exponential demand curve

Another issue with the exponential demand curve arises when we attempt to estimate the corresponding elasticity of demand. The elasticity for Eq. (8) can be calculated by taking the derivative of the logarithmic form in Eq. (9), which, as shown by Bentzley et al. (2013), gives

$$E = -k\alpha P e^{-\alpha P}. \qquad (10)$$

Unfortunately, this value depends fundamentally on the parameter k, whose value, as we have said, may be chosen differently by different experimenters depending on the precise recipe they are using. To see an example of why this matters, let us calculate the maximum value of the elasticity $E_{\max}$ for given values of the parameters Q_0, α, and k. To do this, we differentiate Eq. (10) with respect to P and set the result to zero giving:

$$-k\alpha e^{-\alpha P} + k\alpha^2 P e^{-\alpha P} = 0. \qquad (11)$$

Canceling a number of factors and rearranging we find that $\alpha P = 1$, and, substituting back into Eq. (10), we find the maximum value of the elasticity to be

$$E_{\max} = -\frac{k}{e}, \qquad (12)$$

where $e = 2.718\ldots$ is the base of the natural logarithm.

Table 2 Values of the parameters Q_0 and α for the four fitted curves in Fig. 2, each for a different value of k as indicated, along with the calculated values of the quantities $\tilde{\alpha}$, $P_{\max}$, $\widetilde{P}_{\max}$, and $R_{\max}$

k	Q_0	α	$\tilde{\alpha}$	$P_{\max}$	$\widetilde{P}_{\max}$	$R_{\max}$
3	0.319	0.000476	0.01491	130.2	41.5	10.4
4	0.400	0.00379	0.00948	94.3	37.7	11.3
5	0.420	0.00267	0.00636	96.9	40.7	13.0
6	0.390	0.00183	0.00469	111.8	43.6	14.4

Thus, the maximal value of the elasticity depends only on k (and the mathematical constant e). Since k is determined using an ad hoc recipe that can give different answers for the same data depending on experimental context, this means that the value of $E_{\max}$ is largely arbitrary.

The net result is that, in most cases, values of the elasticity determined from data fits to Eq. (8) (or Eq. (9)) are not informative. This may be one reason why elasticity has not found wide use in the analysis of self-administration data.

Maximum work performed for the exponential demand curve

We have seen that the price $P_{\max}$ at which maximum work is performed falls at the point where the elasticity E is equal to -1 ("Measuring motivation" and Eq. (5)). For the exponential demand curve, the elasticity is given by Eq. (10) and hence, maximum work occurs when

$$-k\alpha P e^{-\alpha P} = -1, \tag{13}$$

which can be solved to give

$$P_{\max} = -\frac{W(-1/k)}{\alpha}, \tag{14}$$

where W is the so-called Lambert W-function (Olver et al. 2010; Gilroy et al. 2019).[7] Once we have the value of $P_{\max}$, it is straightforward to compute the corresponding value of $R_{\max}$ by substituting into Eqs. (4) and (8), which gives

$$R_{\max} = P_{\max} Q_0 e^{k(e^{-\alpha P_{\max}}-1)}$$
$$= -\frac{W(-1/k)}{\alpha} Q_0 e^{k[e^{W(-1/k)}-1]}. \tag{15}$$

In current approaches to data analysis, the standard procedure is to hold the value of k constant over different sessions and different animals, meaning that $W(-1/k)$ is also constant, so that Eq. (14) implies that $P_{\max}$ is inversely proportional to α and hence, as discussed in the "Parameters for the exponential

demand curve" section, the two quantities measure the same thing, a point that has been made clearly by, for example, Bentzley et al. (2013). The detailed relationship between α and $P_{\max}$ however still depends on the value of k since Eq. (14) can be rearranged to read

$$\alpha = -\frac{W(-1/k)}{P_{\max}}. \tag{16}$$

This means that even when working with identical data, from animals with the same $P_{\max}$, different experimenters will arrive at different results for α if they use different values for k. Under the circumstances, therefore, we do not recommend using the value of α as a measure of motivation, given that $P_{\max}$ itself contains the same information but is not directly dependent on k.

Some writers have argued in favor of using α because for some formulations of the demand curve it automatically incorporates a normalization by a factor of Q_0 of the kind discussed in "Normalized price and comparisons between different reinforcers". One can rewrite the demand curve of Eq. (9) thus:

$$\log Q = \log Q_0 + k\left(e^{-(\alpha/Q_0)Q_0 P}-1\right)$$
$$= \log Q_0 + k\left(e^{-\tilde{\alpha}Q_0 P}-1\right), \tag{17}$$

where

$$\tilde{\alpha} = \frac{\alpha}{Q_0}. \tag{18}$$

Equation (17) is exactly equivalent to the original form of Eq. (9) but expresses the consumption as a function of the normalized price $Q_0 P$. This is the form in which the exponential demand curve is mostly commonly written,[8] and it involves a change in the definition of α according to Eq. (18). Combining Eq. (18) with Eq. (16), we find that

$$\tilde{\alpha} = -\frac{W(-1/k)}{Q_0 P_{\max}} = -\frac{W(-1/k)}{P_{\max}}, \tag{19}$$

where we have used the definition of $\widetilde{P}_{\max}$ from Eq. (6).

In other words, $\tilde{\alpha}$ is inversely proportional to the normalized measure $\widetilde{P}_{\max}$ and hence measures the same thing. Like $\widetilde{P}_{\max}$, it is also independent of potency. Because of this, $\tilde{\alpha}$ has been used widely as a measure of motivation and is in fact probably the most commonly used such measure. It is sometimes referred to as the "essential value," and often denoted simply α, though we prefer the notation $\tilde{\alpha}$ to avoid confusion with the original α parameter of Eq. (8). While $\tilde{\alpha}$ is an

---

[7] Some authors have used common (base 10) logarithms instead of natural (base e) logarithms, in which case the solution takes a slightly different form—see Gilroy et al. (2019).

[8] The price in this context is often written as C rather than P. This, however, is merely a notational choice.

appropriate normalized measure of motivation, however, it still suffers from the same shortcoming as the unnormalized α, that it depends on the choice of k, via Eq. (19). We saw an example of this issue in "Parameters for the exponential demand curve" (see Table 2). A further reason to avoid the use of $\tilde{\alpha}$ is that it is specific to the exponential demand curve form. The value of $\widetilde{P}_{max}$ can be calculated for any demand curve, but $\tilde{\alpha}$, by its nature as a parameter of the exponential form, can only be calculated if one uses that form.

An alternative form for the demand curve

In the "Fitting the demand curve" section, we examined the use of fitted demand curves as a way of quantifying the variation of consumption with price and discussed the widely used exponential form, Eq. (8), which has played an important role in the field but has some disadvantages—in particular that it does not go to zero as price becomes large and that it depends crucially on the parameter k which is not determined by a fit to the data. Here, we propose an alternative form for the demand curve which behaves in many ways like the exponential form but eliminates these shortcomings. The form we propose is also mathematically simpler, making solution of the resulting equations more straightforward. In the accompanying materials, we provide a software program that calculates demand curves and parameters such as P_{max} and R_{max} from experimental data using our proposed form, as well as giving visualizations of the demand and revenue curves and best-fit parameter values.[9] Instructions for using the program are given in the Supplementary Information.

Form of the proposed demand curve and parameters

The functional form we propose for the demand curve is

$$Q = Q_0 \left[1 + \left(\frac{P}{P_0} \right)^b \right]^{-a/b}. \tag{20}$$

The parameters, depicted in Fig. 3, are as follows:

Q_0: the preferred consumption level, as previously, i.e., the height of the plateau in the curve.

P_0: the price at the point where the curve starts to fall off, where the work is no longer worth the outcome.

a: the slope of the right-hand part of the curve where it falls off.

b: a parameter controlling the width of the "knee" or transition region between the left and right parts of the curve.

Figure 4 shows the effect of varying each of these parameters on the shape of the demand curve.

This form for the demand curve is flexible enough to fit a variety of different data types (see also Figs. 5 and 7). It has three basic regions, one flat, one curved, and one downward sloping, and hence it can fit data with any of these forms, or a combination of all three.[10]

Since it is common to plot the demand curve on logarithmic scales, one can also rewrite Eq. (20) in terms of the log of Q thus:

$$\log Q = \log Q_0 - \frac{a}{b} \log \left[1 + \left(\frac{P}{P_0} \right)^b \right]. \tag{21}$$

The two forms, Eqs. (20) and (21), are entirely equivalent and contain the same information. One can employ either natural (base e) or common (base 10) logarithms—the results are identical either way.

Elasticity and measures of motivation

We can repeat the analyses of "Elasticity for the exponential demand curve" and "Maximum work performed for the exponential demand curve" for this new demand curve. The elasticity, E, can be computed from the logarithmic derivative, Eq. (3), which gives

$$E = -\frac{a}{1 + (P_0/P)^b}, \tag{22}$$

and the maximal value of E, the equivalent of Eq. (12), occurs when P goes to infinity, giving simply

$$E_{max} = -a. \tag{23}$$

Thus the elasticity in this case bears a very simple relationship to the values of the parameters.

The value of P_{max} is given, as previously, by the point at which the elasticity equals -1, i.e., by

$$-\frac{a}{1 + (P_0/P)^b} = -1. \tag{24}$$

Solving for P we find that

$$P_{max} = \frac{P_0}{(a-1)^{1/b}}. \tag{25}$$

This is a more convenient form than Eq. (14), which involves the special function W and requires a complex iterative procedure to calculate P_{max} (Hursh and Silberberg 2008). Equation (25) by contrast can be evaluated using only a simple calculator or spreadsheet.

---

[9] See http://umich.edu/~mejn/demandcurve for details.

[10] Note that, since Eq. (20) has four parameters, it should not be fit to data with less than four price points, although such data are rare in practical situations.
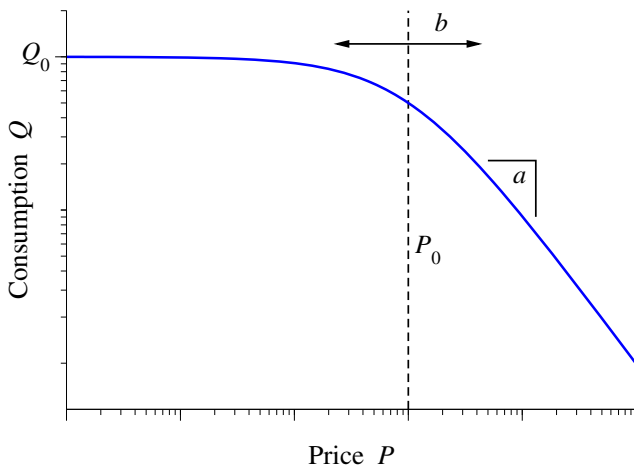
Fig. 3 Example of the demand curve form of Eq. (20), along with an indication of the role played by each of the parameters Q_0, P_0, a, and b

As discussed in "Normalized price and comparisons between different reinforcers," it is useful when comparing values for different goods to normalize the value of P_{max} to give a potency- or magnitude-independent measure of motivation. The relevant measure in the present case is

$$\widetilde{P}_{max} = Q_0 P_{max} = \frac{P_0 Q_0}{(a-1)^{1/b}}. \qquad (26)$$

Given the value of P_{max}, we can compute the corresponding value of R_{max} by substituting into Eqs. (4) and (20) to get

$$R_{max} = P_0 Q_0 \left[\frac{(a-1)^{a-1}}{a^a}\right]^{1/b}. \qquad (27)$$

Combining Eqs. (26) and (27), we can also derive a direct relation between $\widetilde{P}_{max}$ and R_{max} thus:
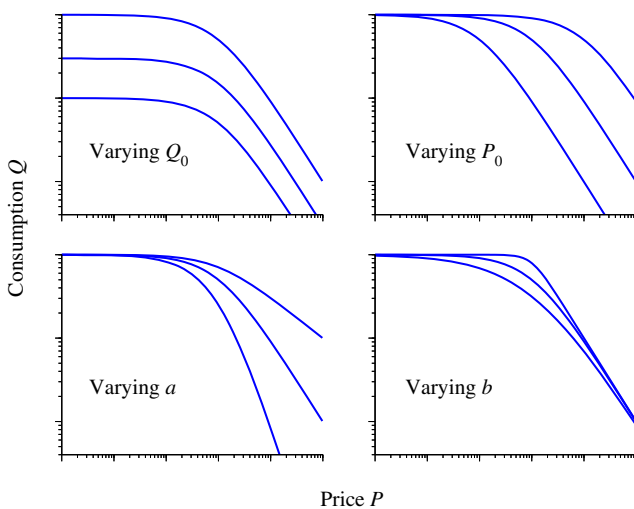
Fig. 4 The effect of varying each of the four parameters in Eq. (20)
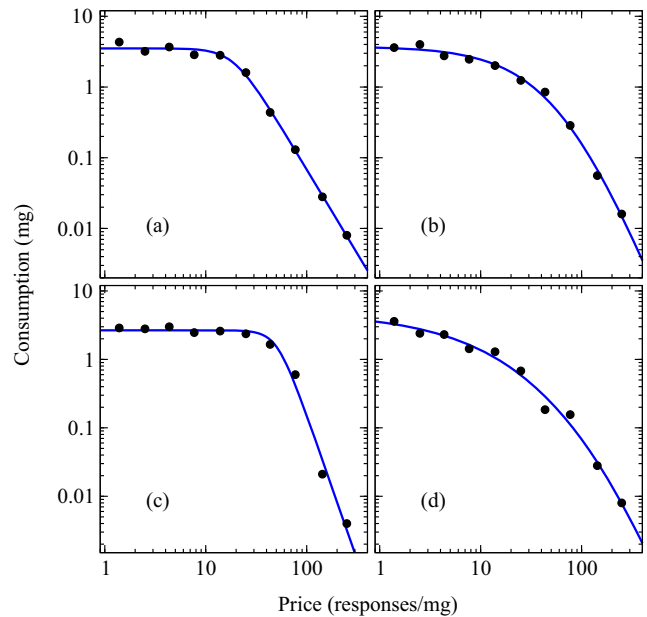
Fig. 5 Cocaine consumption as a function of price in four individual male rats. The blue lines show the best fit of the proposed demand curve form, Eq. (21)

$$R_{max} = \widetilde{P}_{max} \left(\frac{a-1}{a}\right)^{a/b}. \qquad (28)$$

The quantity in brackets is never greater than 1, meaning that we always have $R_{max} \leq \widetilde{P}_{max}$, with the equality occurring in the limit where b becomes large.[11] In practice, b can quite often take values as large as 100 or more, in which case we expect R_{max} and $\widetilde{P}_{max}$ to be essentially equal. These observations shed some light on the question raised in "Measuring motivation" of the extent to which R_{max} and $\widetilde{P}_{max}$ measure the same thing: the answer appears to be that in some cases (but not all) they do and that overall we expect them to be correlated. We test this conclusion against data in "Example calculations."

We do not recommend using the values of the fitted parameters P_0, a, and b themselves as measures of motivation or behavior. When using the exponential form of Eq. (8) (or Eq. (9)), it is common practice, as described in "Maximum work performed for the exponential demand curve," to use the fitted value of the parameter $\tilde{\alpha}$ (also denoted α elsewhere) as an indicator of reinforcement or motivation, but we do not recommend using the parameters of Eq. (20) in this way since they do not have a clear behavioral interpretation (with the exception of Q_0, which is certainly informative). Instead, we recommend using the derived quantities R_{max} and $\widetilde{P}_{max}$.

---

[11] The same result also holds for the exponential demand curve. Equation (15) can be rewritten as $R_{max} = \widetilde{P}_{max} e^{k(e^{-\alpha P_{max}} - 1)}$, but $e^{k(e^{-\alpha P_{max}} - 1)} \leq 1$ always, so this implies that $R_{max} \leq \widetilde{P}_{max}$.

Table 3 Values of Q_0 for the four fitted curves in Fig. 5 along with the calculated values of the quantities P_{max}, $\widetilde{P}_{max}$, and R_{max}

	Q_0	P_{max}	$\widetilde{P}_{max}$	R_{max}
a	3.52	17.3	60.8	41.6
b	3.74	27.0	100.8	33.6
c	2.65	40.3	107.1	85.0
d	4.84	19.0	91.7	15.0

Example calculations

Figure 5 shows example fits[12] of the demand curve of Eq. (21) to data from four different male rats self-administering cocaine under conditions in which each lever press results in an infusion of drug (FR1) and price is varied within session by systematically reducing the dose of drug from 0.72 to 0.004 mg/kg.[13] As the figure shows, there are a range of different qualitative forms in the data, some with a clear plateau followed by a drop-off, in the classic shape of Fig. 1a, others with a gentler curved form with less of a clear "knee." Nonetheless, the proposed form for the demand curve fits all of the data sets well.

After performing the fits, one can use the fitted parameter values to compute measures such as P_{max}, $\widetilde{P}_{max}$, and R_{max}. The results are shown in Table 3.

Note that there is wide variation in the values of both $\widetilde{P}_{max}$ and R_{max} and that large values of one do not always correspond to large values of the other. As discussed in "Elasticity and measures of motivation," we expect R_{max} and $\widetilde{P}_{max}$ to be correlated, and they must satisfy the constraint $R_{max} \leq \widetilde{P}_{max}$, but they are separate measures and can differ substantially.

To shed more light on the relationship between R_{max} and $\widetilde{P}_{max}$, we show in Fig. 6 a plot of the values of R_{max} against those of $\widetilde{P}_{max}$ for cocaine self-administration by nine different male rats over ten sessions each, with one point for each session for a total of 86 points (with four sessions omitted for reasons given below). The plot shows substantial correlation between the two measures ($R^2 = 0.716$) and the fact that $R_{max} \leq \widetilde{P}_{max}$ is clear in that all points lie on or below the diagonal dashed line at which $R_{max} = \widetilde{P}_{max}$. These findings are consistent with previous work examining relationships between R_{max} and $\widetilde{P}_{max}$ (MacKillop and Murphy 2007; MacKillop et al. 2009).

The results shown in Fig. 6 omit data from four sessions, one for which the demand curve had no point with elasticity -1 (and hence no P_{max}) and three that returned values of $\widetilde{P}_{max}$ that were extreme outliers, one having a value of almost 500. If these three outliers are included in the fit then most of the correlation disappears ($R^2 = 0.155$). It is worth asking therefore what causes these outliers. Recall that $\widetilde{P}_{max} = Q_0 P_{max}$ (Eq. (25)), so that large $\widetilde{P}_{max}$ values can be generated either by large P_{max} or by large Q_0. We see both behaviors in the present case: one of our three outliers is caused by a large value of Q_0, which appears to be due to the fact that the measured price range failed to span the "knee" in the demand curve, but the other two outliers are caused by large values of P_{max}. This suggests that $\widetilde{P}_{max}$ may be a (slightly) less reliable measure of motivation than R_{max}, which shows no extreme outliers, at least in the data we have examined. This issue would be an appropriate one for further investigation

Finally, in Fig. 7, we show fits of our proposed demand curve form to data from a range of published studies on consumption of drugs of different drug classes by rats, monkeys, and, in one case, humans. The new form fits this diverse selection of data well, even when the data do not follow the traditional demand curve form as price increases, as in Fig. 7d for example. It is important to note that in such cases the fit will still return values for P_{max}, R_{max}, and other quantities but that they may not be meaningful because the price range
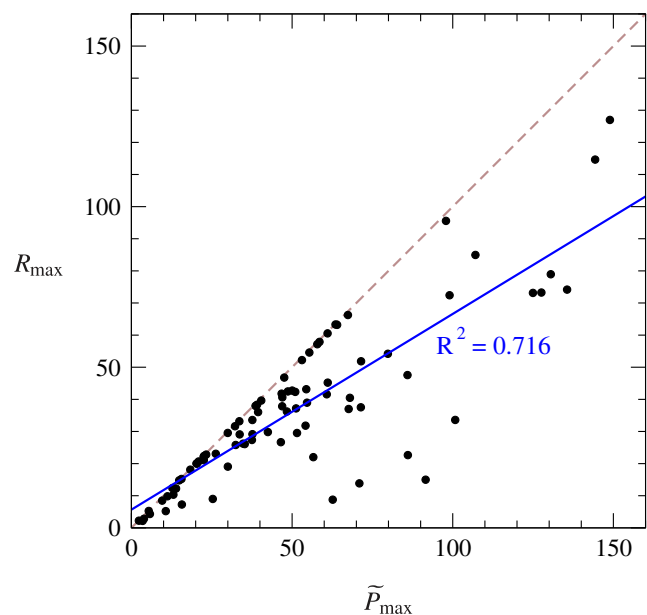

Fig. 6 Plot of the values of R_{max} against $\widetilde{P}_{max}$ for demand curves for cocaine consumption by nine male rats over 86 separate sessions. The values of the two quantities are substantially correlated ($R^2 = 0.716$) and clearly obey the rule $R_{max} \leq \widetilde{P}_{max}$ derived at the end of "Elasticity and measures of motivation." (The dashed line shows the point at which $R_{max} = \widetilde{P}_{max}$, and all points lie on or below this line.) Three extreme outliers on the horizontal axis have been omitted from the plot and from the fit, as discussed in the text
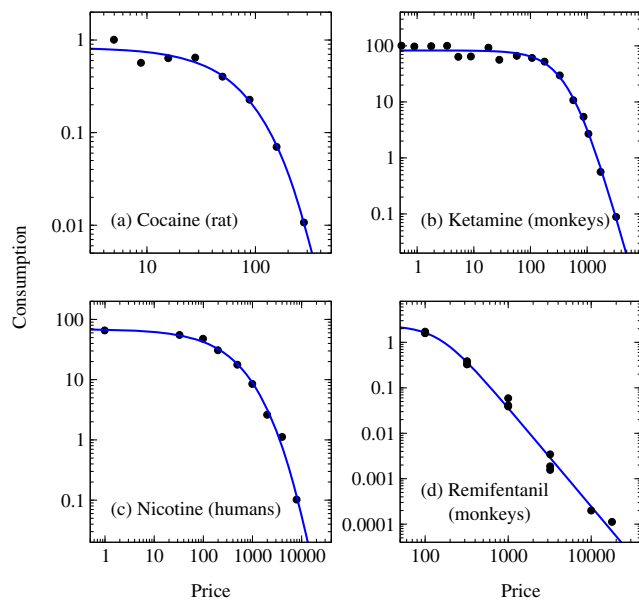
---

[12] We perform our fits using a least-squares best fit of Eq. (21) to the logarithm of consumption, discarding data points with zero consumption, since one cannot take the log of zero. It is also possible to fit consumption data directly to Eq. (20) and it is a straightforward modification of the approach to do so, but in practice we find this gives visually poorer fits. See also Koffarnus et al. (2011).

[13] Previously unpublished data kindly provided by C. Carr and T. E. Robinson.

Fig. 7 Fits of the demand curve of Eq. (21) to data on self-administration of four different drugs. **a** Cocaine in a male rat from Bentzley et al. (2013) (Fig. 3). **b** Ketamine in male monkeys from Winger et al. (2002) (Fig. 3, upper panel). **c** Nicotine in humans from Giordano et al. (2001) (Fig. 1). **d** Remifentanil in male monkeys from Winger et al. (2006) (Table 1)

examined does not cover the plateau and/or knee of the demand curve and hence does not provide a good estimate of Q_0 or P_{max}. For this reason, it is important to inspect all fits visually, not just for goodness of fit to the data, but also to verify that the behavior of interest is actually captured by the range of prices probed in the experiment.

Conclusions

The pioneering work of those who brought microeconomics concepts to the study of food and drug self-administration provides a rich foundation on which to build. In this paper, we have discussed in detail the use of demand curves and elasticity to quantify the motivating effects of food and drugs in behavioral experiments. We have highlighted a number of issues with current methodology in this area and proposed a new mathematical framework for the analysis of consumption data that remedies these issues. This framework incorporates a proposed mathematical form for the demand curve that is flexible enough to fit data for many different reinforcers and experimental protocols and provides established metrics of demand and the formulas needed to calculate them from fits to data. The equations require only basic algebra for their implementation, although we also provide software implementing our entire analysis pipeline, including data fits, plotting, and the calculation of summary statistics (see Supplementary Information). Although our primary focus has been on the analysis of data from animal experiments, the proposed approach is also suitable for human data. The

use of individual metrics and their interpretation will certainly vary from one context to another and will always be at the discretion of the experimenter, but the general framework has wide applicability across many experimental settings.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

References

Aston ER, Cassidy RN (2019) Behavioral economics demand assessments in the addictions. Curr Opin Psychol 30:42–47

Ayers RM, Collinge RA (2004) Microeconomics. Prentice Hall, Upper Saddle River

Becker GS, Grossman M, Murphy KM (1994) An empirical analysis of cigarette addiction. Am Econ Rev 84:396–418

Bentzley BS, Fender KM, Aston-Jones G (2013) The behavioral economics of drug self-administration: a review and new analytical approach for within-session procedures. Psychopharmacology 226: 113–125

Bickel W, Marsch L, Carroll M (2000) Deconstructing relative reinforcing efficacy and situating the measures of pharmacological reinforcement with behavioral economics: a theoretical proposal. Psychopharmacology 153:44–56

Cartwright E (2018) Behavioral economics, 3rd edn. Routledge, New York

Galuska C, Banna K, Willse L, Yahyavi-Firouz-Abadi N, See R (2011) A comparison of economic demand and conditioned-cued reinstatement of methamphetamine-seeking or food-seeking in rats. Psychopharmacology 22(4):312–323

Gilroy SP, Kaplan BA, Reed DD, Hantula DA, Hursh SR (2019) An exact solution for unit elasticity in the exponential model of operant

demand. Experimental and Clinical Psychopharmacology 27:588–597

Giordano LA, Bickel WK, Shahan TA, Badger GJ (2001) Behavioral economics of human drug self-administration: progressive ratio versus random sequences of response requirements. Behav Pharmacol 12:343–347

Hursh SR (1993) Behavioral economics of drug self-adminstration: an introduction. Drug Alcohol Depend 33:165–172

Hursh SR, Silberberg A (2008) Economic demand and essential value. Psychol Rev 115:186–198

Hursh SR, Winger G (1995) Normalized demand for drugs and other reinforcers. J Exp Anal Behav 64:373–384

Hursh SR, Galuska CM, Winger G, Woods JH (2005) The economics of drug abuse. Mol Interv 5:20–28

Kahneman D (2011) Thinking fast and slow. Farrar, Straus and Giroux, New York

Kawa A, Bentzley B, Robinson T (2016) Less is more: prolonged intermittent access cocaine self-administration produces incentive-sensitization and addiction-like behavior. Psychopharmacology 233:19–20

Koffarnus MN, Hall A, Winger G (2011) Individual differences in rhesus monkeys' demand for drugs of abuse. Addict Biol 17:887–896

Lacy RT, Austin BP, Strickland JC (2019) The influence of sex and estrous cyclicity on cocaine and remifentanil demand in rats. Addict Biol 25:e12716

MacKillop J, Murphy JG (2007) A behavioral economic measure of demand for alcohol predicts brief intervention outcomes. Drug Alcohol Depend 89:227–233

MacKillop J, Murphy JG, Tidey JW, Kahler CW, Ray LA, Bickel WK (2009) Latent structure of facets of alcohol reinforcement from a behavioral economic demand curve. Psychopharmacology 203:33–40

Oleson EB, Richardson JM, Roberts DC (2011) A novel IV cocaine self-administration procedure in rats: differential effects of dopamine, serotonin, and GABA drug pre-treatments on cocaine consumption and maximal price paid. Psychopharmacology 214:567–577

Olver FWJ, Lozier DM, Boisvert RF, Clark CW (2010) NIST handbook of mathematical functions. Cambridge University Press, Cambridge

Pantazis C, James M, Bentzley B, Aston-Jones G (2019) The number of lateral hypothalamus orexin/hypocretin neurons contributes to individual differences in cocaine demand. Addiction Biology Jul 11.

Perloff JM (2016) Microeconomics: theory and applications with calculus, 4th edn. Pearson, New York

Richardson N, Roberts D (1996) Progressive ratio schedules in drug self-administration studies in rats: a method to evaluate reinforcing efficacy. J Neurosci 66:1–11

Winger G, Hursh SR, Casey KL, Woods JH (2002) Relative reinforcing strength of three N-methyl-D-aspartate antagonists with different onsets of action. J Pharmacol Exp Ther 301:690–697

Winger G, Galuska C, Hursh S, Woods J (2006) Relative reinforcing effects of cocaine, remifentanil, and their combination in rhesus monkeys. J Pharmacol Exp Ther 1318:223–229