
Taxing Goods and Services

Taxes on goods and services—excise duties, stamp duties, and, most importantly, VAT—raise nearly 30% of tax revenue in the UK and are important parts of all modern tax systems. VAT especially has become an increasingly important part of the tax system in the UK and elsewhere. Over 150 countries worldwide, including every OECD country apart from the US, employ VAT systems.

Of course, a tax on consumption levied at a uniform rate on goods and services is economically equivalent to a flat-rate expenditure tax, as described in Chapter 2 and which we look at in detail in Chapters 13 and 14.

We organize our discussion of indirect taxes into four chapters. In the next chapter, we look at the design of VAT, explaining why VAT is, in general, a well-designed consumption tax. In Chapter 8, we give particular consideration to the taxation of financial services. Then, in Chapter 9, we look at a particular reform to VAT in the UK. In this chapter, we consider the choice of *base* and *structure* for indirect taxes.

The first question to consider concerns what should be taxed. We will argue in favour of taxing goods and services when they are purchased to consume, but not when they are purchased to use as inputs in the production of other goods and services. We will also argue against the use of transactions taxes, such as stamp duties, on the sale or purchase of assets.

The main issue we address in considering the structure of indirect taxes is whether the same tax rate should be levied on all goods and services. The current VAT system in the UK uses several different rates. We look at whether the efficiency arguments favour a uniform or a differentiated VAT system, and then consider issues of equity. There are some quite subtle and

complex arguments here, but, in practice, the case for differentiated rates looks weak and that for a broader, more uniform, structure looks strong. Imposing a uniform rate avoids distorting the choices people make between different goods. And the case for imposing lower VAT rates on goods and services consumed disproportionately by poorer people is weakened in a modern tax system where we have income taxes and benefit payments available to achieve distributional goals more effectively.

Tax rates are also differentiated by imposing excise duties in addition to VAT on a small number of goods, such as tobacco products, alcohol, and motor fuel. A strong case for this can be made when the consumption of particular goods and services generates externalities, i.e. when consumption has positive or negative spillovers to other people's welfare. But we should recognize that a large part of the argument for high alcohol tax and, especially, high tobacco tax is behavioural, based on the negative effects of their consumption on the consumer herself—the 'internality' effect.

6.1. WHAT TO TAX

The starting point for our analysis is that it is the consumption of goods and services which is the appropriate tax base. This fits closely with our arguments in Chapter 2. In general, consumption and expenditure by consumers can be considered the same thing—though in the case of very long-lived goods such as housing, consumption occurs over such long periods that we might prefer to levy a tax as consumption occurs rather than when the house is first built and sold.

VAT is structured such that it is paid in the end by consumers. Businesses in general do not pay VAT when they buy inputs into their production processes. VAT is also designed to be a tax on consumption, not a tax on transactions. We will see in the next chapter how VAT is designed such that these statements are true. We show here why, in general, it is appropriate to have a tax neither on inputs into the production process nor on transactions.

6.1.1. The Production Efficiency Argument against Taxing Inputs

Production efficiency occurs when the economy cannot increase the output of one good without having to produce less of another. In an idealized market economy without any taxes, the prices of inputs act as signals that guide individual firms to make choices that ensure production efficiency is achieved.

Taxes can be levied on produced inputs¹ as well as on final consumption goods. However, taxes on produced inputs would distort the input choices of firms and result in a loss of production efficiency. In a famous paper, Diamond and Mirrlees² show that the tax system should be designed to ensure production efficiency is attained. This implies that produced inputs should not be taxed, so that all taxes should fall on final consumption goods.

The reasoning is straightforward. What matters for the welfare level of the economy is the amount of consumption enjoyed by each consumer. Consumption choices are determined by the prices of final goods relative to the wage rate. Any set of prices for final goods that can be obtained using a combination of input taxes and final goods taxes can be achieved by final goods taxes alone. Taxes on input prices are therefore superfluous in respect of their effects on consumption choices. They can, however, affect production decisions in a way that creates production inefficiency. If the tax system results in production inefficiency, then welfare can be easily increased: more of a good that someone enjoys can be produced and consumed without using any additional labour. A tax system that places the economy in a position where there are such unexploited gains cannot be efficient.

In fact, when there are externalities or other market failures, the strict conditions needed for this result may often not be met. The case for taxing environmentally damaging inputs into production remains. But this apart, the requirement for production efficiency is powerful and a key reason for the use of VAT in preference to taxes that burden intermediate transactions.

¹ A produced input is an input into the production process that is itself the output of an earlier production process. Such inputs can equally be called *intermediate inputs*.

² The argument is demonstrated in Diamond and Mirrlees (1971) for an economy with no market failures.

The UK tax system (in common with many others) does indeed tend to eschew taxation of inputs.³ But there is a set of taxes levied not on final consumption but on transactions. These are stamp duties. We come to these briefly now.

6.1.2. Should We Tax Transactions?

Taxes on share dealing, house purchase, and land sales are an important part of the fiscal landscape in the UK. These ‘stamp duties’ have a long history and their continued use reflects the ease with which such taxes can be levied, given the need for people to register their ownership of these items. However, they are unattractive from an economic point of view.

This lack of attractiveness stems from the fact that any tax on transactions will reduce expected welfare by discouraging mutually beneficial trades. Welfare is maximized when assets are owned by the people who place the highest value on them. Taxing transactions will affect who owns an asset, and so can disrupt the efficient pattern of ownership.

The value of a good or service is determined by the flow of benefits that are derived from owning it. So a consumption tax can be levied either on the purchase price of the good or service when it is first sold or on the flow of benefits over time. A transactions tax does not do this and it always seems preferable to tax the benefits directly. For example, if the benefit of owning the asset comes in the form of income, as in the case of share dividends, then it can be taxed as income. Alternatively, if the benefit comes in the form of a flow of consumption services, it can be taxed along the lines that we discuss for housing in Chapter 16. In no case do we find the arguments for a transactions tax compelling.

Of course, if a good is sold only once before final consumption, then a transactions tax (on the value of the transaction) and a VAT have the same economic effect: they raise the price of the good by the rate of tax. The difference between the two taxes arises when the good is sold more than once before consumption. A durable good, such as a house, can be sold

³ With one major exception—the taxation of business property through the business rate (see Chapter 16).

many times over its lifetime. A financial asset can have an unlimited potential lifetime and be traded many, many times—hence potentially being liable for a transactions tax many, many times. A transactions tax will reduce the price of assets that are traded relatively frequently, it will reduce the number of mutually beneficial trades, and it will fall arbitrarily heavily on those who, for whatever reason, engage in more transactions.

Stamp duty on house transactions, for example, taxes according to the number of times a house changes hands over its lifetime. Houses vary considerably in the number of times they are traded, but there is no good economic argument for taxing more-frequently-traded housing more. Worse still, a tax on transactions reduces the incentive to trade in housing and leads to less efficient usage of the housing stock. A tax on the consumption value of housing would make sense (as we see in Chapter 16) but a stamp duty on transactions does not.

An argument for a transactions tax has been put forward, however, in cases where transactions may not always be efficiency enhancing. In particular, this case has been made for a tax in some financial markets, where it is argued that trade is excessive and in some situations destabilizing: the transaction itself induces a negative externality. Then a transactions tax can be seen as a way to reduce excessive speculative activity, and consequent price volatility. This was the justification for the original ‘Tobin tax’ on financial transactions,⁴ which involved the taxation of transactions on foreign exchange markets at a very low rate, but which can be applied to any financial transaction.

While superficially appealing, there are reasons to doubt the premise on which this argument is based. It has been argued that speculation could only be profitable if it reduced volatility, while empirical research has not found any clear link between speculation and volatility.⁵ There is some evidence, though, that greater use of financial transactions taxes would affect incentives to trade. The prices of more-frequently-traded shares appear to be responsive to announcements of changes in stamp duty in the UK.⁶ So, for

⁴ Tobin, 1978.

⁵ A recent contribution to this literature is Radalj and McAleer (2005).

⁶ Bond, Hawkins, and Klemm, 2004.

better or worse, transactions taxes are likely to have an impact, even when levied on purely financial assets.

The Tobin tax has gained renewed support in the wake of the recent financial crisis,⁷ though not, it seems, because of any convincing evidence that its existence would have reined in the growth of banks' balance sheets or dealt with the kind of asset market inflation that spawned the crisis. It is now seen more as a means to obtain additional tax revenue from the financial sector. But it is important to be clear on where the incidence will lie. There is no particular reason to believe that the owners of financial sector companies would bear the burden of a tax on foreign exchange transactions. It might well be passed on to consumers in the form of higher import prices. More general financial sector transactions taxes would likely be passed on to savers in the form of lower returns.

This is not to deny that we might want to rethink the taxation of the financial services sector, in particular because it is currently undertaxed as a result of not being subject to VAT. Moreover, the degree of undertaxation grows as the standard rate is raised. We consider this in Chapter 8. But, in general, there is a weak economic case for taxing transactions rather than the income from, or consumption of, the asset or good that is changing hands. While current anger at the financial sector may be justifiable, it is difficult to make a compelling case for using a transactions tax to deal with what are essentially regulatory issues. It is preferable to start by removing the favourable treatment of financial services in VAT.

Ultimately, the existing transactions taxes on housing and share dealing have little compelling economic logic behind them.

6.2. EFFICIENCY, EQUITY, AND UNIFORMITY

So, taxes should only be levied on goods and services used for consumption. But should the same tax rate be applied to all, or are there good reasons for applying different rates to different goods and services?

⁷ In fact, even before the current episode, Stiglitz (1998) was making this argument.

There are some goods—alcohol, tobacco, petrol—where a case for a high rate of tax can readily be based on the harmful effects of their consumption on others. Indeed, in the case of tobacco especially, the long-term harmful effects on the consumer herself can also be the basis for a high tax rate. We discuss these issues briefly in Section 6.3. But first we consider the arguments for and against uniformity in situations where there are no such spillover effects.

In general, of course, there is an initial presumption in favour of uniformity to avoid distortion of consumption decisions. It is costly in welfare terms if the tax system results in me buying a different set of goods from the ones that I would have bought had all goods been treated the same. Uniformity also avoids the complexities and political lobbying that are inevitable concomitants of a differentiated regime. It can reduce both the administrative and compliance burdens of the tax system. It has practical advantages for businesses, which currently need to work out the VAT categories into which their products and purchases fall. They face much more complicated calculations and paperwork if more than one category applies. Uniformity avoids the sometimes farcical process of deciding exactly which goods should be taxed at which rates. It would allow us to dispense with the need for court cases to establish whether Jaffa Cakes are cakes or biscuits.⁸

Uniform taxation also simplifies the politics of decision-making and makes tax policy less vulnerable to lobbying pressure and short-term political considerations. A system with differentiated rates invites interest groups to lobby for lower rates for their own products, or at least rates as low as those on other preferentially treated commodities.⁹ Arguably, this has been important in past decisions to create concessions in the VAT base, such as the reduced rate on domestic energy.

⁸ Chocolate-covered biscuits are subject to VAT, while chocolate-covered cakes are not. McVitie's produced a giant Jaffa Cake for the court to illustrate that their product was really a miniature cake, not a biscuit, arguing that cakes generally go hard when stale while biscuits go soft. The VAT Tribunal ultimately upheld the makers' claim. *United Biscuits (UK) Ltd (VTD 6344)*.

⁹ This argument is made very strongly by Buchanan in Buchanan and Musgrave (1999).

These arguments in favour of a simple uniform system seem persuasive. But there are important arguments against.

First consider just the efficiency issues. Suppose the government has to raise a certain amount of revenue and cares only about the total deadweight loss caused by the tax system but not its distribution across people. As we saw in Chapter 2, a lump-sum tax, with everyone paying a fixed amount, would be efficient. But it is generally impractical. Now suppose that purchases of commodities can be taxed. In a world in which indirect taxes are passed straight on to consumers in the form of higher prices, the efficiency loss resulting from a commodity tax will depend on the extent to which the rise in price reduces people's demand for the good or service in question.

The deadweight loss or 'welfare cost' of a tax is greatest where it has the largest impact on people's purchasing behaviour. It follows that the deadweight loss is smallest when higher rates of indirect tax are imposed on goods for which demand is relatively inelastic (in other words, where the higher price will do relatively little to deter people from buying them) than on those for which there is relatively elastic demand. This *inverse elasticity rule*¹⁰ suggests that since goods generally differ in their price elasticity,¹¹ optimal tax rates would be differentiated across commodities.¹² In practice, accurately implementing the recommendations of the inverse elasticity rule would require knowledge about the responsiveness of demand to price for individual goods. It would also require decisions on the level of disaggregation at which to distinguish between different goods and services, and on how often to vary the tax rates as market conditions change and alter the optimal rates. The more closely the authorities want to adhere to the

¹⁰ For a formal demonstration of this result, see Atkinson and Stiglitz (1980) or Myles (1995). The inverse elasticity rule is valid when there are no substitutability or complementarity relationships between commodities. When such relationships are admitted, the efficient tax system is described by the *Ramsey rule*. This rule is more general but has the same overall consequences as the inverse elasticity rule.

¹¹ Blundell, Pashardes, and Weber, 1993.

¹² Because there is a budget constraint relating labour income to total expenditure, we can generally relate conditions on demand elasticities to work and leisure choices. Uniformity will be efficient when all goods are 'equally substitutable' for leisure. Otherwise, goods that are more complementary with leisure time will bear higher taxes.

most efficient set of tax rates over time, the greater the costs both of administering the system and of complying with it. But, in general, this rule would suggest a case for higher tax rates on goods such as basic foodstuffs, which are ‘necessities’ and for which price elasticities tend to be relatively low, than on other goods.

This is the main efficiency argument for having differential tax rates.¹³ In fact, much of the deviation from uniformity we see in the UK VAT system and elsewhere is driven by considerations of equity—goods such as food and domestic fuel, which form larger parts of the budgets of poorer people, are taxed at zero or reduced rates as a means of promoting ‘fairness’. Such concerns for fairness lie behind most opposition to the notion of a uniform VAT.

It is worth distinguishing three egalitarian arguments for differentiated VAT. The main one relies on the fact that poorer people spend a larger portion of their income on certain goods. A slightly different argument suggests that there are some goods that should face lower tax rates because they are in some sense ‘essentials for life’. A third and separate argument is that, separate from income, spending patterns might themselves reveal something particular about a person’s needs. The second and third of these arguments we consider separately in the next subsection. Here we focus on the first.

When indirect taxation is considered in isolation, and when there are concerns for equity, there looks to be a strong case for differentiating tax rates to help low-income households by imposing lower taxes on goods that they consume disproportionately. But indirect taxes should *not* be considered in isolation from the rest of the tax and welfare system. Where the government is able to levy a progressive income tax and pay welfare benefits that vary according to people’s needs and characteristics, this will generally prove a much more effective means of meeting its equity objectives.

However, there are some circumstances in which governments cannot levy progressive income taxes at all effectively. The income tax (and benefit)

¹³ There are others, including the suggestion that lower taxes can offset the effects of market power where firms are able to charge above the efficient price (Boadway and Pestieau, 2003), but we do not pursue those here.

system may be constrained, for administrative reasons, to be relatively simple. This might typify a developing country that has relatively few direct tax and benefit instruments available. In this case, a government might care about distribution but can only use indirect taxation to achieve its redistributive goals. It will then want to tax at a lower rate (and possibly even subsidize) the necessities, which make up a larger share of the expenditure of the poor. Since necessities, such as basic food items, also tend to be price inelastic, the government faces a trade-off between efficiency and concerns for redistribution. This can result in keeping tax rates on price-inelastic goods such as food relatively low even though, for pure efficiency reasons, one would want to tax these goods at higher rates.

But the UK and other developed economies do have access to sophisticated direct tax and benefit systems. As we demonstrate in Chapter 9, it is possible to introduce a uniform VAT in the UK whilst changing the direct tax and benefit system to produce an outcome with similar distributional (and work incentive) features to those that are achieved with extensive zero-rating.

However, there does remain a more subtle argument for indirect tax differentiation which relies on the fact that even a sophisticated direct tax and benefit system cannot achieve a 'first-best' outcome. Underlying abilities and needs are not observed with perfect accuracy. A high-ability person may put in little work effort and earn the same amount as a low-ability individual who works very hard. If both have the same earnings, both will pay the same tax even though the high-ability person works less. This information constraint limits the redistribution that can be achieved via an income tax.

In this case, there may be an argument for distorting choices across commodities if the preferences for some commodities are related to work and leisure choices or to earning capacity.¹⁴ To develop this argument, note that the government would like to redistribute from the more able to the less able, but the extent to which this can be done is limited since a high-wage individual can always work less and benefit from the redistribution. These are the incentive compatibility constraints that enter any modern discussion of tax design. The higher the tax rate on higher earned incomes, the more the high-wage individual will choose to work less. Out of this comes the simple

¹⁴ Technically referred to as 'non-separabilities with work'; see Atkinson and Stiglitz (1976).

rule—tax goods that are ‘complementary’ to leisure as this will discourage the high-wage individual from taking too much leisure. Of course, since there is a deadweight loss to such tax distortion, there is a limit to the extent to which we want to tax such complementary goods. But some degree of higher taxation of goods that are ‘complementary’ to leisure activities and, conversely, lower taxation of goods that are ‘complementary’ to work will be warranted.

The argument here is quite straightforward. Goods that take more time to consume, such as restaurant meals or theatre performances, require more leisure time. Hence, taxing them more highly is a form of leisure tax and will encourage those who consume them to work longer hours and take less leisure. In general, by taxing goods and services that are associated with leisure more heavily—and goods and services associated with work more lightly—we can partially offset the disincentives to work that a redistributive tax system inevitably creates. A similar effect is possible by imposing lower tax rates on goods that are associated with longer hours of work (ready meals, for example).

Assessing the practical significance of this argument is not easy, although there are some instances where it might be compelling. For example, taxing childcare services (during working hours) less heavily than other goods and services would help offset the disincentive to work created by other parts of the tax system. The same might apply to some kinds of public transport (particularly peak-time travel). Subsidy of goods that increase the time available for work, or its effectiveness, such as (some) medicines, can be justified. Perhaps less compelling, goods and services that are most useful during leisure time (such as fishing rods, suncream, and cooking ingredients) are candidates for higher tax rates, while reduced tax rates on takeaway and ready meals, dishwashers, and repair services might encourage people to do more paid work instead of doing these activities themselves.

6.2.1. Egalitarianism and Horizontal Equity

A conceptually different argument for differentiation rests on what James Tobin has described as *specific egalitarianism*¹⁵—the idea that there are specific domains in which we seek to limit inequality as an end in itself. For example, differentiation of tax rates could be justified to avoid taxing ‘life’s essentials’. Poorer households spend more of their money on cigarettes than rich households, but one rarely hears calls for these items to be given preferential treatment: concerns about ‘fuel poverty’ have greater resonance than concerns about ‘cigarette poverty’. What distinguishes commodities such as food and domestic fuel (along with the likes of water and sewerage services and burial and cremation) is not that they take up a larger share of poorer households’ budgets but that they are essentials of life. Unlike cigarettes, people must unavoidably buy a certain amount of these goods.

Many goods where such egalitarian sentiments prevail, such as education and health care, are provided by the state. However, many are allocated through the operation of the market. To distinguish specific egalitarianism from the argument for generalized redistribution, we must believe that people will still choose to buy ‘too little’ of these goods even if they have the money to do so. The goal is to encourage people to buy these goods in particular rather than giving them enough money in general. In one sense, this is a more coherent argument for differentiation than is the general equity argument, since we could not achieve this outcome more efficiently using other tools that we currently have. On the other hand, it is an argument that sits uncomfortably with a belief that people are generally able to make the right decisions for themselves.

The third egalitarian argument for a differentiated VAT is that some spending patterns might convey extra information about consumers’ abilities or needs and hence might be useful ‘tags’ for the tax system to achieve specific distributional ends. A good example in the current UK tax system is that vehicles for people with disabilities are exempt from VAT on the cost of adaptation and exempt from payment of the annual vehicle excise duty. As with broader distributional objectives, the case for differentiating indirect tax rates for this reason depends on there not being more efficient alternatives.

¹⁵ The concept of specific egalitarianism is discussed in Tobin (1970).

This might be so if other distributional tools have undesirable disincentive effects or if it is difficult to target certain groups using them.

There may be one or two exceptions of this sort, but, in general, when other more direct instruments exist, using differentiation in the indirect tax system to achieve distributional objectives is likely to be costly and inefficient.

We should round off this discussion by reminding ourselves that there is, in fact, another type of equity for which uniformity is important. This is what we described in Chapter 2 as *horizontal equity*, or the desire to tax similar people in similar ways. Individuals with the same income or wealth may differ in the way that they like to spend their money. When the rate of tax is differentiated across goods, some individuals are rewarded and others penalized in a way that can appear rather arbitrary. At present in the UK, we subsidize those who spend large amounts of money on designer clothes for their children but tax those who spend similar amounts on, perhaps rather educational, toys. Those with a taste for music are taxed; those with a taste for magazines are not. We clearly want to avoid differentiating rates between people who differ only in *inconsequential* characteristics (such as a taste for Jaffa Cakes rather than chocolate-covered biscuits).¹⁶

At the extremes, non-uniform taxation can even appear discriminatory when differences in tastes and needs mirror characteristics such as age and gender, where society has explicit anti-discrimination policies.

6.2.2. Balancing the Arguments

For efficiency reasons, it is always a good thing for consumer prices to line up with marginal costs. If markets are competitive, then non-uniformity

¹⁶ This consequence of this approach is demonstrated in Atkinson and Stiglitz (1976). Assume there are two consumers. One consumer has a strong taste for vanilla ice cream relative to chocolate ice cream but does not care greatly about the choice between red and white wine. The other consumer has a strong taste for red wine but does not mind about ice cream. It is then efficient (applying the inverse elasticity rule) to tax vanilla ice cream and red wine at a higher rate than chocolate ice cream and white wine. The taxes are differentiated, but the two consumers have the same level of welfare.

drives prices and costs apart in a way that is always undesirable in itself and can be justified only by some countervailing consideration.

Where the government is looking to raise revenue, and price elasticities of demand for different goods differ, then there is an efficiency case for imposing high taxes on those goods with lower price elasticities. This will generally mean higher taxes on necessities. On the other hand, if the government is concerned about equity, and does not have an effective direct tax and benefit system to achieve redistribution, then there is a case for imposing lower taxes on those goods that take a greater part of the budget of the poor. However, where there are effective progressive direct taxes and benefits available, these will do a better job at redistributing. In this case, differential commodity taxation will enable more redistribution only where commodity purchases reveal something about effort. Lower taxation on goods that are complementary to work may allow more effective redistribution.

In general, though, it is hard to implement most of these distinctions. With the likely exception of childcare, the case for the greater complexity associated with differentiated tax rates is not proven. The equity and revenue-raising goals of any differentiated indirect tax plus an income tax can be achieved more efficiently by replacing these taxes with a uniform indirect tax and a more progressive income tax. The progressivity of the income tax achieves redistribution; the lack of differentiation in indirect taxes avoids distortions in choice of commodities; both help raise revenue.¹⁷ This is not surprising. Income taxes and benefits can be closely designed to achieve distributional outcomes. Indirect taxes are generally very blunt instruments. While we do not doubt the case for redistributing to low-income households, differential commodity taxes are an inefficient way to achieve that end, for the reasons that we have outlined.

Differentiation of the tax rates on commodities redistributes on the basis of how much people spend *on particular items*. But if we wish to redistribute from those most able to pay to those least able to pay, we can do so more accurately based on the level of people's *total* expenditure—or their earnings,

¹⁷ This statement is justified in Laroque (2005a) and Kaplow (2008).

income, or wealth.¹⁸ Unless someone's spending on particular items can tell us something about his ability or his needs that his overall spending and income cannot, differential indirect tax rates cannot help us achieve the redistribution we want. So, the general case for differentiating commodity tax rates as a way of achieving redistribution is, in the presence of other tools such as an income tax, a weak one.

In Chapter 9, we will explore the implications of extending the VAT base while looking for ways of compensating some of the households that lose. While we cannot compensate each and every household, we will find that, for the most part, the argument that there are better ways of dealing with distributional concerns is largely borne out in practice. This result is somewhat complicated, but not undermined, by the potential for a progressive income tax to create work disincentives.

In sum, the efficiency arguments for differential tax rates are important but, in our view, can be very hard to operationalize in practical terms. The one exception to this is that there is probably a strong case for exempting childcare costs from VAT because, in many cases, spending on childcare is so closely related to the choice over how many hours to work. The equity arguments might also be powerful in the absence of other parts of the tax and benefit system that can redistribute more effectively. But those other parts do exist, and they can redistribute more effectively without creating the distortions that differential VAT rates bring with them. In addition, considerations of horizontal equity—treating similar people similarly—argue in favour of uniformity.

6.3. SPILLOVERS AND BEHAVIOURAL CHANGE

So far, we have avoided discussion of the most significant divergences from uniformity. In the UK, and in most other countries, alcohol, tobacco, and automotive fuels (petrol and diesel) are subject to taxes in addition to the standard VAT. These are usually levied as excise duties—a tax per unit

¹⁸ Chapter 13 discusses which—or what combination—of these is the most appropriate tax base.

bought, rather than a tax proportional to value (though there is also an additional *ad valorem* element to the taxation of tobacco products in the UK). These duties continue to account for a substantial element of UK exchequer revenues. In 2010–11, fuel duties were expected to raise over £27 billion, tobacco duties over £9 billion, and various alcohol duties £9.5 billion.¹⁹ There are also a series of smaller environmentally related taxes, insurance premium tax, and additional taxes on betting and gaming.

There are convincing arguments for these sorts of differentiated tax rates where the consumption of a particular good or service creates spillover costs or benefits for individuals other than the consumer. An obvious example is the imposition of higher tax rates on products, such as petrol, the consumption of which harms the environment. We devote separate chapters to considering environmental taxes in some detail (Chapters 10–12).

The basic principles for taxing spillovers are straightforward. In general, the additional or reduced taxation on any good should reflect the damage or benefit that the consumption of a little more of it creates. This ensures that the private decision about how much of the good to consume takes into account the impact on others. In practice, there are limits to this. Estimating the damage or benefit—and the appropriate tax differential—is often difficult. Indeed, the external effects of consumption may vary dramatically according to where, when, and by whom the consumption is taking place. Most moderate drinkers, for example, impose no costs on the rest of society. But the costs imposed by a minority, through associated accidents and crime, may be very high.

Importantly, even though taxation can be an effective tool with which to influence the amount of socially costly or beneficial consumption in the economy, it may not always be the *best* tool. If the harm is large enough, an outright ban would likely be the best policy: for example, alcohol is taxed, while drugs perceived as more damaging are prohibited. Similarly, compulsion may sometimes be the best way to achieve beneficial spillovers: for example, it seems preferable to require vehicle owners to take out motor insurance rather than merely to impose a lower tax rate on it.

¹⁹ HM Treasury, 2010b, table C11.

But it would be a mistake to think that taxes on tobacco and alcohol in particular are entirely, or even largely, about reducing negative effects on others. They are actually in some part justified by the rather more contentious argument that the consumption of alcohol and, especially, tobacco harms the consumers themselves. Rather than creating an externality—an impact on other people—the consumption of these goods may create what some authors have called an *internality*²⁰ and others refer to as ‘an externality to one’s future self’²¹—an impact on the consumer. The idea is that taxes can encourage people to avoid acting against their own self-interest. They might act against their own self-interest, and might require action to prevent them doing so, either because they do not understand the costs of their consumption decision, or because the goods are addictive (and, in particular, addictive in ways that are not fully anticipated), or because the decision is in some sense not rational.

In all these cases, the government can act paternalistically and use taxes to change prices and encourage people to change their behaviour. The general principle is similar to the case of spillovers: tax rates should include a component to reflect self-harm, with a larger tax where self-harm is deemed to be larger. Such arguments are certainly influential in the taxation of alcohol and tobacco, although (as with spillovers) it is hard to know what level of taxation would be proportionate to the costs to the individuals concerned. The costs of consuming a particular product may also differ across different groups. Children and young adults, for example, might be thought particularly susceptible to the consequences of drinking and smoking. As it happens, there is also evidence that such groups may be the most responsive to taxes.²²

Underlying many of these arguments is the observation that individuals make inconsistent choices at different times during their life. This may be especially true at the point at which smoking begins. The majority of smokers start smoking when they are young. They are unlikely to make the decision to do so with a clear view about the future. This is well illustrated by a survey that followed a group of school seniors in the US who smoked a

²⁰ Gruber, 2003a.

²¹ Viscusi, 1995.

²² See e.g. Chaloupka and Wechsler (1997).

pack or more a day. Of those who expected still to be smoking in five years' time, 72% were still smoking. Of those expecting *not* to be smokers five years hence, 74% were still smoking.²³

If consumers cannot control their short-term desire to smoke, even though they would expect it to be of long-term benefit to control it, then in more reflective moments a smoker might actually prefer higher taxation as a counterweight to his or her short-term lack of control. This is a plausible view of human psychology (and indeed supported by experimental evidence²⁴), but the implications for policy are not straightforward. The tax rate required to offset people's lack of self-control would vary widely. Taxation could also be seen as penalizing fully-rational individuals who choose to smoke and drink when they are entirely cognizant of the (current and future) benefits and costs.

Similar sorts of arguments have also been made for taxing fatty foods. This is much more complex than taxing alcohol and tobacco, in part because, of course, moderate consumption of most foods is beneficial, but perhaps even more importantly because of the extreme difficulty in defining a category of fattening or harmful foods different from all others. We already have a situation in the UK in which certain unhealthy foods are standard rated for VAT while others are zero rated. Ice creams, biscuits, and potato crisps fall into the first category; cakes, tortilla chips, and chocolate cookies bought from a bakery fall into the second. This illustrates both the use of taxation to deter consumption of unhealthy foods and the difficulty of making a clear assignment into healthy and unhealthy categories.

6.4. CONCLUSIONS

The main conclusion from the analysis in this chapter is that there is a strong case for a move to a broader-based and more uniform system of indirect taxation. There are a few clear-cut situations where there should be deviations from uniformity—taxes on environmental harms, and taxes on

²³ Gruber, 2003b.

²⁴ Gruber and Koszegi, 2001.

goods such as alcohol and tobacco that can have damaging effects on the consumer and on other people, are the obvious examples. But the case for the widespread differentiation in indirect tax rates that we see in the UK at present is not strong. In particular, if we are concerned about equity, then it is much better to use the direct tax and benefit system to achieve the distributional outcomes that we favour than it is to use differential indirect tax rates. We will look in detail at how this might be done in Chapter 9.

There are reasons other than equity for favouring differential tax rates, including a desire to tax more lightly the consumption of those goods associated with work. This is likely to provide a strong case for a low (perhaps zero) VAT rate on childcare. One could make a case for some other goods and services in this category, but, in the absence of strong evidence to the contrary, our view is that the advantages in terms of simplicity of a single rate are likely to outweigh any possible advantage from differentiating tax rates for this or other reasons of efficiency.

Prior to these conclusions is the conclusion that economic efficiency is best served by taxing consumption goods and not by taxing either produced inputs or transactions *per se*. We tackle some of the issues associated with this principle in Chapter 16, where we consider stamp duty (a transactions tax) and business rates (a tax on a produced input). In this context, the question of what we might do with financial services is dealt with briefly in Chapter 8.

Finally, indirect taxes need to be seen in an international context. Significant changes to the economic environment in recent years have brought international tax issues to the fore. For EU member states, for example, the most important of these was the completion of the single European market in January 1993, which directly affected the operation of the VAT system. In an international context, there is a fundamental question over *where* taxation should take place. In practice, indirect taxes are almost exclusively levied on a destination basis at present—that is, in the country where purchases take place rather than in the country where production occurs (as would happen if taxes were levied on an origin basis). While this is likely to continue to be the case, it does create a number of inefficiencies associated with cross-border shopping and problems of administration, to which we return in the next chapter.