1 3. Lesson 2: Social opportunity costs

In Parts 1 and 2, we saw how prices in competitive markets reflect the opportunity costs faced by producers and consumers. For many writers on economics, including Hazlitt, this is the beginning and end of the story. The conclusion they draw is that government action that takes society away from the market allocation can only be for the worse.

In reality, however, markets don't work in the idealised fashion assumed in simple tracts like Economics in One Lesson. As a result, many opportunity costs arising in the process of production and consumption aren't reflected, or aren't fully reflected, in market prices.

To begin with, there is nothing special about the particular market equilibrium we observe at any given time. There is an infinite range of possible allocations of property rights, each corresponding to different social choices, and each associated with a different competitive equilibrium.

Second, the actual outcome in a market economy differs greatly from the ideal competitive equilibrium. Markets for vital services like health and education work poorly or don't exist at all. Social and economic problems including unemployment, pollution and monopoly are further examples where markets don't work in the way that Hazlitt assumes. This large class of problems is collectively known as 'market failure'. Although market failures are many and varied, all involve the failure of market prices to reflect opportunity goods.

One type of market failure, the cycle of boom and bust that gives rise to mass unemployment, is so severe and so pervasive that it has become the subject of a special branch of economics, called

macroeconomics. The name, which refers to the study of the economy at an aggregate level, is distinguished from microeconomics, the study of individual prices and markets and the way they interact in equilibrium.

The evidence from macroeconomics is that, for the economy as a whole, resources are not always allocated on the basis of opportunity cost. Rather, there are long periods of recession and depression where productive resources sit idle, so that their opportunity cost, in effect, is zero.

The inability of markets to resolve questions of distribution, and the various forms of market failure form the basis of Lesson Two

Lesson 2: Market prices don't reflect all the opportunity costs we face as a society.

In Part 3, we will look at Lesson 2 in detail. We first examine how the logic of opportunity cost applies to the distribution of income and wealth. Next we will look at a variety of forms of market failure, drawing on the classic work of Francis Bator (1958). Finally, we will consider how to interpret the classic macroeconomic problems of recession, unemployment and inflation in terms of opportunity cost.

1.1 Property rights, and income distribution

The competitive equilibrium we talked about in Lesson 1 is not the unique product of spontaneous social processes. Rather it depends on the allocation of property rights on which trade is based. Before we can trade in markets, we must determine who owns what. This

determination is subject to the logic of opportunity cost, but can't be reduced to market transactions.

Presented with this problem in the abstract, most people would prefer an egalitarian initial allocation, leading to outcomes where everyone is better off than they were before entering into trade, and no one is much better off than anyone else. In reality, though, there is no starting point at which we get to make a once-for-all choice. People enter the world with endowments of all kinds that are determined, in greater or lesser measure, by those of their parents.

1.1.1 What Lesson 2 tells us about property rights and income distribution

In any market economy, the outcome of interactions between individuals, families, businesses and governments depends on the initial allocation of property rights and resources that determines the starting point for trade and employment. Those property rights include not only ownership of houses, factories and so on, but the set of rights and obligations created by taxation and welfare systems, and the legal framework within which economic activity takes place.

The range of possible initial allocations and institutions is vast, and so is the range of possible market outcomes they can generate. In fact, according to economic theory, any final outcome that is consistent with the technological possibilities available to society, and that takes full advantage of the possibilities for trade, can arise as the market outcome, given the right initial allocation.

What this means is that the choice of any particular starting point, and the resulting market outcome, entails an opportunity cost, namely, forgoing all the alternative possibilities. Increasing the allocation of rights to one person or group will, in general, reduce what is available for everyone else, and this will be reflected in the market outcome.

1.1.2 The starting point

If we are going to consider changes in the distribution of income and wealth, what should we take as out starting point? There are various possibilities, many of which are of theoretical interest, but not of much practical use.

Hazlitt doesn't spell out the starting point for his analysis. However, his analysis is based on the implicit claim (spelt out in more detail by Bastiat) that there is a natural distribution of private property rights, that exists prior to any government activity such as taxation and the payment of welfare benefits.

This is nonsense. It is impossible to disentangle some subset of property rights and entitlements from the social and economic framework in which they are created and enforced.

The ordinary meaning of "property" refers to a specific kind of control over resources, most completely realised in freehold ownership of land. In the idealised model which forms the basis of much thinking about property

Most of the time, we take the existing allocation of property rights for granted. This is, however, an example of exactly the fallacy pointed out by Bastiat, that of focusing on what is seen and ignoring the unseen

alternatives. All property rights began with a decision by governments to create and enforce someone's right to use a particular good, asset or idea, and to regulate the way in which that right might, or might not, be transferred to others.

In some of the cases discussed in Section 2, such as those of telecommunications spectrum and fishing quotas, the rights were created relatively recently, and the process by which they were created is well documented. In somewhat older cases, such as that of the 19th century innovations which created limited liability corporations, the history has been forgotten by all but a few specialists. Going even further back, property rights in land and in ordinary goods (chattels, in legal parlance) are mostly taken for granted, even though they are all derived, in the final analysis, from a state-created legal framework.

Propertarians like Hazlitt want to pare back government to the minimum necessary to protect the property rights of which they approve. These include rights over land and houses, private sector financial assets and personal possessions.

There are two main difficulties with this.

First, propertarians disagree among themselves as to which government functions should be retained, and which property rights should be maintained. For example, some support core government functions like police and fire services while others want these to provided, on a market basis, to those willing to pay for them. Similarly, some propertarians, support the idea that the creators of ideas should have unlimited 'intellectual property' in those ideas, while others believe that 'information ought to be free'.

Moreover, while propertarians almost invariably oppose ‘welfare’ benefits paid out of tax revenue, such as social security, there is no clear dividing line between these benefits and contractually obligatory payments such as pensions for public and private workers.

The fine distinctions between Austrians, minarchists, objectivists, and anarcho-capitalists are too complex and tedious to be detailed here. The point is that any attempt to define, on the basis of logical first principles, a ‘natural’ set of property rights, independent of government, runs rapidly into quicksand.

The second problem is that any attempt to strip all rights and entitlements back to a minimal set corresponding to a naive notion of ‘private property’ would not produce anything like the existing distribution of private property rights. Some kinds of private property would become much more valuable, and others much less so. An example can be seen in the mass privatisations that followed the end of Communism in Russia and other countries in the former Soviet bloc., These processes greatly enriched a handful of oligarchs and greatly impoverished everyone else, leading to the loss of the little .

It is impossible to describe a proposed starting point based on such a radical change with any accuracy. So, we can’t really say what the opportunity cost of shifting property rights from one person to another might be in such a situation.

It makes sense, therefore, to start thinking about the initial allocation with reference to our actual position rather than to some or other theoretical ideal.

In most modern societies, governments collect a substantial proportion of national income in taxation revenue. Some of this revenue is spent

on the provision of public services, and some on ‘transfer payments’ such as social security, unemployment and disability insurance, and assistance to poor families.

The starting point therefore includes both the existing set of property rights of workers, the employment position of worker and the rights and obligations of members of the community to receive government services and benefits and to pay the taxes necessary to finance those services and benefits.

1.1.3 The opportunity costs of redistribution

There are many policy changes that will improve the starting position for some members of the community. Examples include

(A) Reducing marginal rates of income tax above some income level, which will benefit those with taxable incomes above that level.

(B) Increasing the duration of intellectual property rights such as copyrights and patents, which will benefit the owners of those rights

(C) Increasing the number of publicly funded places in colleges and universities, which will benefit the young people who are enable to attend

(D) Increasing social security payments and unemployment insurance, which will benefit those who are unable to work because of age or inability to find a job

(E) Increasing the minimum wage

Over the past 40 years, we have seen substantial changes of types (A) and (B) in the United States and elsewhere around the world. The top

marginal rate of income tax has been reduced from ... to ... . The maximum term of copyright protection has been extended from ... Other measures, such as the use of ISDS provisions in trade agreements, have created a variety of new and expanded property rights for corporations.

By contrast, there have been few changes of types (C), (D) and (E). On the contrary, public funding of universities has been reduced, eligibility for social security has been tightened and the real value of the minimum wage has been reduced.

This outcome reflects the logic of opportunity cost. To finance increased expenditure on some goal or to reduce the taxes paid by one group, the government must find offsetting cuts in expenditure or increased taxes elsewhere, or else accept a larger deficit, incurring a debt that will have to be serviced in the future. The least unattractive of these options, as evidenced by the choices of policymakers, will constitute the opportunity cost of providing the benefit.

Creating new property rights or extending old ones provides the owner with control over resources, including ideas, that were previously accessible to all. Users other than the owner will either be excluded from the resource or will have to negotiate terms with the owner; the associated costs represent the opportunity cost.

1.1.4 Opportunity cost of redistribution: Example

Any change in the allocation of rights and obligations will create benefits for some people and costs for others. Consider a simple example: a reduction of 0.1 per cent in the top marginal rate of

income tax, currently 39.6 per cent, providing roughly $1 billion in additional net income to those with pre-tax incomes over $400 000.

The opportunity cost of such a policy is the offsetting measure needed to finance it. Possibilities include a reduction in public expenditure, an increase in other forms of taxation or the issuance of debt that must be repaid in the future. For concreteness, let's suppose that the tax cut is finance by a reduction in unemployment insurance payments.

How large a reduction is needed? Both collecting taxes and operating unemployment insurance schemes involve administrative costs. Collecting taxes is costly, as is administering unemployment insurance.

The opportunity cost of this policy is not limited to the $1 billion transferred from high income earners to the unemployment.The opportunity cost of a transfer payment includes the value of resources spent in administration, as well as the amount transferred.

But, as advocates of the free market will be quick to point out, that's not all.

Reducing tax rates on high-income earners will lead to changes in the opportunity costs they face. In particular, the opportunity cost of taking additional leisure time, namely the additional expenditure that could be enjoyed with a higher post-tax income, increases as tax rates fall.

This change in opportunity costs, often referred to as an 'incentive effect' means that high-income earners will tend to allocate more time to work, and less to leisure, when tax rates are reduced. Some of the resulting extra income will flow back to the government in the form of tax revenue, partly offsetting the initial cost of the tax cut.

More importantly perhaps, the lower are tax rates, the less effort high income earners, and their lawyers and accountants, may be expected to put into schemes to avoid or reduce tax liabilities.

From the viewpoint of someone paying a tax rate of 40 per cent, and not concerned with the ethics of tax avoidance, a scheme that turns a dollar of taxable income into 70 tax-free cents is well worth while. The benefit of 70 cents exceeds the opportunity cost of 60 cents of disposable income. So, we can expect lower marginal tax rates to be associated with some reduction in the resources devoted to tax avoidance.

On the other side of the transfer, it it is often argued more generous unemployment benefits reduce the opportunity cost of remaining unemployed, namely the income foregone, and therefore make the unemployed less keen to seek work. The evidence on this point is mixed in the US context, but there is probably at least some effect.

Taking all these points into account, the opportunity cost of a $1 billion reduction in the tax paid by top income earners will be a reduction of less than $1 billion in the net benefits that can be paid to the unemployed.

For those concerned solely with ‘efficiency’ or maximizing the market value of GDP, that’s sufficient to resolve the issue. Cutting taxes on the rich, and impoverishing the poor even further, will generally increase GDP.

But GDP is an arbitrary aggregate, which tells us nothing about the social opportunity costs and benefits of different allocations of rights and obligations. To assess the desirability of a redistribution of rights,

such as a reduction in marginal tax rates we need to answer two questions

First, what is the opportunity cost? In this case, how much do we have to reduce net payments through unemployment insurance in order to finance a cut in marginal tax rates.

Second, how should we weigh the benefits to some (in this case, the wealthy) against the opportunity costs borne by others (in this case, the unemployed).

We will address these questions in the next two sections.

1.1.5 TANSTAAFL and the Laffer hypothesis

Scratch someone with a TANSTAAFL bumper sticker and you're pretty sure to find a believer in the so-called so-called 'Laffer curve'. The idea associated with that phrase represents the ultimate 'free lunch': the claim that by cutting income tax rates for high income earners, it is possible to generate more tax revenue, which can then be used to make everyone better off.

Everyone knows the story of how Laffer drew a graph on a napkin, illustrating the point that tax rates of 100 per cent would result in a cessation of economic activity and therefore yield zero revenue. Since a tax rate of zero will also yield zero revenue, there must exist some rate of taxation that yields a maximum level of revenue. Increases in tax beyond that point will harm economic activity so much that they reduce revenue.

Wanniski christened this graph the 'Laffer curve', but as Laffer himself was happy to concede, there was nothing original about it. It can be traced back to the 14th century Arabic writer Ibn Khaldun. Laffer credited his own version to the nemesis of supply-side economics, John

Maynard Keynes. And while few economists had made much of the point, that was mainly because it seemed too obvious to bother spelling out.

What was novel in Laffer's presentation was what might be called the 'Laffer hypothesis', namely that the United States in the early 1980s was on the descending part of the curve, where higher tax rates produced less revenue.

Unfortunately, as the old saying has it, Laffer's analysis contained a mixture of correctness and originality. The Laffer curve was correct but unoriginal. The Laffer hypothesis was original but incorrect.

For the Laffer hypothesis to be supported, tax cuts would have to increase revenue more rapidly than would be expected as a result of inflation and normal income growth. In fact, as Richard Kogan of the Center on Budget and Policy Priorities reported, income tax receipts grew noticeably more slowly than usual in the 1980s, after the large cuts in individual and corporate income tax rates in 1981.

To the extent that there was an economic response to the Reagan tax cuts, and to those of George W. Bush twenty years later, it seems largely to have been a Keynesian demand-side response, to be expected when governments provide households with additional net income in the context of a depressed economy (See section ...).

There have been attempts to make the fantasy economics of the Laffer hypothesis more respectable, using an idea called 'dynamic scoring'. Studies using this idea have not supported the extreme claims made by Laffer, but they have suggested that a large proportion of any cut in taxes, particularly taxes on capital income, will be recouped in the form of additional revenue.

Dealing with this issue in detail is beyond the scope of this book. However, my earlier book Zombie Economics explains some of the problems with the dynamic scoring approach, which have led to its being abandoned by serious economists.

Moving on from the fantasy world of the Laffer hypothesis, a large number of economic studies have attempted to estimate the relationship between tax rates, economic activity and revenue. The most common finding is that the incentive effects of a dollar in tax cuts generate around twenty cents of additional economic activity. Given a top marginal tax rate of 40 per cents, around eight cents will flow back to the government in the form of tax revenue.

The incentive effects of transfer payments like unemployment insurance are less well understood, but it seems reasonable to use a similar estimate here: that a reduction in unemployment insurance would yield some additional job search and work effort, resulting in around 20 cents of additional economic activity for each dollar in reduced assistance.

In addition, we might estimate 10 cents in the dollar for the administrative costs of the tax and welfare systems, including the resource costs of tax avoidance.

Putting all of these effects together, a plausible estimate is that increasing the incomes of the wealthy by one dollar, through lower tax rates, implies an opportunity cost of 50 cents, in reduced transfers to the poor and unemployed.

Does such a change make society as a whole better or worse off. Answering this question inevitably involves a value judgement. But that doesn't mean economics has nothing to say about the question. We

can use opportunity cost reasoning to clarify our thinking about issues of income distribution.

1.1.6 Weighing opportunity costs and benefits

Changes in the regulation of labor and capital markets and in taxation and expenditure policy have greatly enhanced the income and wealth of the best-off members of society (the so-called 1 per cent), and have yielded more modest, but still substantial, improvements in the position of those in the top 20 per cent of the income distribution (broadly speaking, professionals and business owners and managers).

On the other hand, incomes for the rest of the community have grown much more slowly than might have been expected based on the experience of the decades from 1945 to 1975. The substantial technological advances of recent decades have had little impact on the (inflation-adjusted) income of the median US household. For many below the median, incomes have actually fallen (real wages, welfare reform).

In the absence of the tax cuts of the 1980s, the associated cuts in public expenditure and financial and industrial relations policies that benefitted business, the incomes of the wealthy would not have increased as much as they have done. Those on median and lower incomes would have done substantially better. But how should we compare those gains and losses?

Economists and philosophers have been looking at this question for a long time and in many different ways. The answers most consistent with opportunity cost reasoning can be described by the following ‘thought’ experiment, developed explicitly by John Harsanyi and John

Rawls in the mid-20th century, but implicit in the reasoning of earlier writers like Jeremy Bentham, John Stuart Mill and Friedrich von Wieser.

First consider yourself in the position of both the high income beneficiary and the low income loser from such a change. Next, imagine that you are setting rules for a society, of which you will be a member, without knowing which of these positions you might be in. One way to think of this is to imagine life as a lottery in which your life chances are determined by the ticket you draw.

Now consider a choice between increasing the income of the better off and the worse off person. Presumably, if the dollar increase were the same in both cases, you would prefer to receive it in the case where you are poor rather than in the case when you are rich.

The reasons for this preference are obvious enough. For a very poor person, an additional hundred dollars could mean the difference between eating and not eating. For someone slightly better off, it may mean the difference between paying the rent and being evicted. For a middle class family, it might allow an unexpected luxury purchase. For someone on a million dollars a year, it would barely be noticed.

Economists typically present this point in terms of the concept of marginal utility, a technical term for the benefits that are gained from additional income or consumption. As argued above, the marginal utility of additional income decreases as income rises. It follows that a policy that increases the income of the rich and decreases that of the poor by an equal amount will reduce the utility of the poor more than it increases the utility of the rich.

Few mainstream economists would reject this analysis outright . However, many prefer to duck the issue, relying on a distinction

between ‘positive’ economics, concerned with factual predictions of the outcomes of particular economic policies and ‘normative’ economics, concerned with ‘value judgements’ like the one discussed above. The debate over the justifiability or otherwise of this distinction has been going on for decades and is unlikely to be resolved any time soon.

More importantly, constructs derived from economics are often used, implicitly or explicitly, in ways that imply that an additional dollar of income should be regarded as equally valuable, no matter to whom it accrues.

The most important of these constructs is GDP, the aggregate value of all production in the economy.GDP per person is the ordinary average (or arithmetic mean) income of the community. GDP per person treats additive changes in income equally no matter who receives them.

Used correctly, as a measure of economic activity, GDP can be a useful guide to the short-term management of the economy. In the short run, weak GDP growth is commonly an indicator of a recession, suggesting the need for expansionary monetary and fiscal policies.

Unfortunately, measures of GDP and GDP per person are commonly misused, as an indicator of living standards and economic welfare more generally. There are many reasons why this is inappropriate, but the failure to take account of the distribution of income is most important.

It is easy enough to see that, if the opportunity cost of a given increase the income of a better-off person is an equal increase in the income of a worse-off person, then the change is for the worse.

What about the case when we the choice is between a given increase for the worse off person and a larger increase for the better off

person? How big does the opportunity cost have to be before it outweighs the benefit. This question, raising once again the thought experiment mentioned above, can be answered in many different ways.

One answer, which seems close to the views typically elicited when people are asked questions of this kind, is to treat equal proportional increases in income as being equally desirable. That is, an increase of $1000 in the income of a person on $10 000 a year is seen as yielding a benefit comparable to that of an increase of $10 000 in the income of a person earning $100 000 a year. Conversely, if the opportunity cost of the $10 000 benefit to the high income earner is a loss to the low income earner of more than $1000, the cost exceeds the benefit.

It's surprisingly easy to turn this way of looking things into a measure of living standards over time. If, instead, we want a measure that treats proportional changes equally, all that is needed is to replace arithmetic mean measures such as GDP per person with the geometric mean we all learned about in high school (and most of us promptly forgot).

The geometric mean has the property that, if all incomes increase by the same proportion, so does the geometric mean. So, it's a better measure of the growth rate of incomes across the community than the usual arithmetic mean. It can also be justified mathematically, in terms of the theory of expected utility. For those interested, the details are spelt out in an optional section.

The geometric mean is equal to the arithmetic mean when incomes are distributed exactly unequally. But the more unequal is the income distribution, the greater the gap between the arithmetic and geometric

means. For this reason, the ratio of the arithmetic to the geometric mean is often used as a measure of income inequality.

We can look at the changes in these measures using data from the US Census Office, and some simple computations (details available on request). From 1967 to 2013, arithmetic mean income per household (in 2013 dollars) rose from $66 500 to $104 000, an increase of 56 per cent. But the geometric mean rose by only 34 per cent, from 50 000 to 67,500. The ratio between the two rose from 1.32 to 1.54, indicating a substantial increase in inequality.

The idea that equal proportional increases are equally valuable, and therefore that the geometric mean is a good measure of economic welfare or wellbeing is not the only answer to the question posed above. Another, leading to a strong version of egalitarianism, is always to prefer the increase to the worse off person . In this case, welfare is measured by the minimum income.

There's no way of reaching a final resolution on questions like this. But it's worth observing that a policy aimed at maximising the geometric mean of income would be substantially more egalitarian than anything that has ever been seen in a market economy.

For example, calculations by Peter Diamond and Emmanuel Saez, using a method equivalent to the geometric mean approach, suggest that the top marginal tax rate, after taking account of disincentive effects should be between 70 and 80 per cent.

These rates are far above those found in any country today. And while the top marginal rate was at or above this level in the 1950s, generous exemptions and other loopholes meant that the effective rate was much lower.

It’s not surprising that political outcomes are less egalitarian than an opportunity cost estimate would suggest. The thought experiment leading to the geometric mean gives everyone equal weight, as in an ideal democracy. In practice, however, the well off have more weight in democratic systems than do the poor; and of course the disparity is even greater in undemocratic and partly democratic systems. So, while there are good arguments for more strongly egalitarian approaches, policies aimed at maximizing geometric mean income will inevitably be found well to the left of centre in any feasible political system.