*Erasmus Journal for Philosophy and Economics,*
*Volume 11, Issue 2,*
*Autumn 2018, pp. 1-33.*
*https://doi.org/10.23941/ejpe.v11i2.326*

# On the Very Idea of a Just Wage

JOSEPH HEATH
*University of Toronto*

**Abstract:** The way that wages are determined in a market economy produces results that strike most people as morally counterintuitive, if not positively unjust. I argue that there is an important and easily defensible principle underlying the system—it is designed to channel labour to its best employment, the way that it does any other resource. But many consider this defence too minimal, and so strive to find a thicker, more robust moral principle that can be used to defend the market, using concepts like 'contribution', 'effort', 'laziness', 'skill' or 'talent'—all of which combine to provide a concept of 'desert', or 'fairness' in compensation. The objective of this paper is to caution against such overreach. I begin by articulating what I take to be the central principle underlying the determination of wages. I go on to discuss three different ways that both critics and defenders of the market have sought to go further than this, by introducing thicker moral concepts to the discussion, and why each of these initiatives fails. My central contention will be that markets are structurally unable to deliver 'just' wages, according to any everyday-moral understanding of what justice requires in cooperative interactions.

**Keywords:** wages, distributive justice, marginal productivity

**JEL Classification:** D630, E240, E250

## I. INTRODUCTION

In 2013, Harvard economist N. Gregory Mankiw was widely excoriated for his attempt to defend the incomes of the top one per cent in America. Led by Robert Solow, who lamented the way that Mankiw's "cheerful blandness" drew attention away from the "occasional unstated premises, dubious assumptions, and omitted facts" in his argument, critics were quick to pick apart almost every aspect of the article (2014, 243). ("The 1 percent needs better defenders", declared *The Economist* magazine.) And yet there is one, highly problematic presupposition that

AUTHOR'S NOTE: My thanks to Daniel Hausman, Andrew Lister, as well as participants at the 2015 Ethics of Economic Institutions Conference at Utrecht University, for helpful comments.

not only went unquestioned, but was even accepted by many of Mankiw's critics. Mankiw at one point observes that, "[i]n the standard competitive labour market, a person's earnings equal the value of his or her marginal productivity" (2013, 30). He then goes on to treat this conception of marginal productivity as equivalent to that individual's "contribution to society" (30). On this basis, he assumes that if "the Left" has some concerns about the distribution of income, it must because of the "various reasons that real life might deviate from this classical benchmark" (30).

It is no surprise that if one treats the market as a system of natural justice, whose essential tendency is to ensure that the principle 'to each according to his or her contribution' is respected, then this will generate an enormous presumption in favour of the pattern of wealth distribution that it generates. Indeed, one could see in the article Mankiw struggling even to understand what sort of concern could be animating 'the Left'—after all, why would you not want to reward people based on their contribution? And yet, the suggestion that 'marginal productivity' corresponds to some intuitive or morally compelling basis for the distribution of reward is one that was intensely debated in the early 20th century, and is widely regarded as having been refuted. More generally, the idea that marginal productivity is equivalent to contribution is just one example of an unfortunate tendency many people have of taking concepts that are drawn from everyday morality and the informal social sphere (or what Jürgen Habermas refers to as the 'lifeworld' [1987]), tailored to mediating face-to-face interactions among individuals, then 'reading them in' to patterns that arise in a market economy.

Over the course of his article, Mankiw actually articulates three rather different principles that he takes to govern the reward of labour in the market. In the introduction, he claims that "because people earn the value of their marginal product, everyone has the appropriate incentive to provide the efficient amount of effort" (21). This suggests a consequentialist perspective, according to which wages are largely about providing the correct incentives, with an eye toward the more general goal of promoting economic efficiency. Later on, however, prior to articulating the rather different view that reward reflects contribution, Mankiw suggests that higher reward is associated with superior "talent", and that the relative lack of intergenerational mobility in the United States is due to the heritability of major dimensions of talent, including

IQ (25). The issue of 'talent', along with its supposed rewards, has also played a major role in recent philosophical discussions of market inequality.[1] And yet this issue is a rather marginal one in modern labour economics, one that arises primarily in discussions of the ability of 'superstars' to command economic rents.[2] It is not central to any discussion of everyday wage differentials. Indeed, there is a large empirical literature on inter-industry wage differentials, all of which suggests that various aspects of 'ability', including IQ, play no role in explaining the prevailing patterns.[3] So again, it is very far from obvious that a principle derived from small-scale cooperative interaction, like the idea that 'talent' should be related to greater reward, can be read into the operations of a market economy.

Indeed, even a cursory examination of the empirical literature on wages is sobering, since there remains so much that we do not know or understand. The one thing that can be said with certainty, however, is that the way wages are set in a market economy strikes most people as morally counterintuitive, if not positively unjust. This is why Friedrich Hayek was so strenuous in his insistence, not that markets are just, but rather that markets be treated as *exempt* from such forms of moral assessment. "The manner in which the benefits and burdens are apportioned by the market mechanism would in many instances have to be regarded as very unjust if it were the result of a deliberate allocation to particular people", he wrote (1976, 64). The only adequate defence, in his view, is to insist upon the impersonality of the market mechanism, along with the unplanned and unforeseen character of its results.[4]

Although I think Hayek's view is unduly pessimistic about the possibility of justifying overall market outcomes, the core observation is correct. People have a variety of everyday-moral concepts that arise in the context of managing cooperative labour in small-scale, face-to-face interactions. These include concepts like 'contribution', 'effort',

[1] Most importantly in G. A. Cohen's *Rescuing Justice and Equality* (2008).
[2] See Rosen (1981), and, more generally, Frank and Cook (1996).
[3] For a survey of this literature, see Thaler (1989). On IQ, see Blackburn and Neumark (1992).
[4] In truth, Hayek takes several different positions on this question, not all of which are consistent. After the passage cited, for instance, he goes on to say that market institutions came to be accepted "because it was found that they improve for all or most the prospects of having their needs satisfied" (64). I am emphasizing the idea that market outcomes are exempt from normative assessment simply because Hayek spends so much time and energy developing the 'spontaneous order' concept that one assumes he also saw it as having significance for our thinking about markets. For more general discussion, see Lister (2013).

'laziness', 'skill', or 'talent'—all of which combine to provide a concept of 'desert', or of 'fairness' in compensation.[5] Yet when one looks at the broad patterns of compensation in a market economy—not the one per cent, but more prosaic examples, like how much the custodial staff earn, compared to the lawyers they clean up after; or how much teachers make, compared to public relations consultants; or how much garment workers make in Los Angeles, compared to their counterparts in Bangladesh—it is not difficult to show that the central organizing principles of the labour market are such that the outcomes will essentially be orthogonal to these moral concerns.

Unlike Hayek, I think that there is an important and easily defensible principle underlying the market determination of wages—the system is designed to channel labour to its best employment, the way that it does any other resource. The problem is that many people consider this defence too minimal, and so strive to find a thicker, more robust moral principle that they can use to defend the market. This leads them to the overreach that Hayek cautioned against. What Mankiw's argument reveals is that there is still a great deal of confusion surrounding these normative questions. My objective in this paper is therefore something of a tidying-up operation. I will start by articulating, in the narrowest way possible, what I take to be the actual principle underlying the determination of wages in a market economy, and why that principle should be regarded as providing a general presumption in favour of those outcomes. I will then go on to discuss three different ways that both critics and defenders of the market have sought to go further than this, by introducing thicker moral concepts to the discussion, and why each of these initiatives fails. My central contention will be that markets are structurally unable to deliver 'just' wages, according to any everyday-moral understanding of what justice requires in cooperative interactions—and so we should stop trying to either defend or criticize them in those terms.

[5] In an influential discussion, Joel Feinberg described this concept of "personal desert" as a "natural" moral notion, "one that is not logically tied to institutions, rules and practices", and contrasted it with what a person might be "eligible" for, according to the institutional rules (1970, 56). Since I am skeptical about the existence of a pre-institutional morality, I am opting here to describe the relevant notions as 'everyday-moral' rather than 'natural'.

## II. ON PRICES

If we adopt the traditional definition of an economy as a system that allocates scarce resources among their competing uses, then it is easy to see that any complex economy that hopes to achieve this allocation non-arbitrarily will need to have a system of prices. A price, in this sense, can be understood as simply a quantitative 'score' assigned to a particular use of a resource.[6] The need to assign such a score is driven by the need to compare one use against another, in order to determine which is best (according to whatever conception one may of have of what counts as 'best').

When thinking about the role of prices, a useful comparison may be drawn to the way that a chess-playing computer program works. In order to decide what move to make, the program goes through and systematically examines each of its options. Every available move is the top node in a decision tree, composed of possible moves and more-or-less probable countermoves. The program investigates each tree to a certain depth, then assigns a score to each outcome along every branch of the solution tree. So, for example, the capture of a piece will be worth a certain number of points, depending on its rank; the loss of a piece will result in a commensurate loss of points. Then, based on its estimation of how likely each countermove is, the program will assign a net score to each available move. It will investigate millions of permutations, then look at its moves and choose the one with the highest net score.

The problem that must be solved, when it comes to the production and distribution of goods in the world, is not all that different (Berliner 1999, 159-164). Imagine that some miners strike a particularly rich vein, and so are able to extract an extra hundred tonnes of iron ore. The question for society then becomes: 'What shall we do with it?' There are thousands of different applications. Should it be used to make well pumps? Frying pans? Radiators? Or should it be refined into steel, then made into kitchen knives? Car doors? Roof tiles? The important point is that, no matter how one thinks such questions should be answered—or through which process—some quantitative basis of comparison between different uses will be required. Whatever general objective one thinks the economy should be aimed at satisfying, each different use will satisfy it to a greater or lesser degree.

[6] For discussion, see Berliner (1999, 70-82).

In this respect, the comparison to chess is slightly inapt, in that human players typically manage to play without the need to engage in explicitly quantitative assessment. Similarly, in a very small-scale economy, it may be possible to allocate resources and goods to their best employment without quantification. Thus the stipulation above, that we are concerned with the situation in a 'complex' economy, is important. The key point lies in the recognition that the optimization problem involved in determining the best use of resources is subject to a combinatorial explosion, as new goods are introduced into the economy, because goods are used to produce other goods, and so each decision made about the level of production of one good has implications for the production of multiple other goods. As planners in the former communist nations found, producing a plan that is even *consistent* is an enormous challenge, without getting into the question of optimality (Nove 1991, 86). The more general point is that the entire process cannot even get started until some set of prices has been introduced.

Once we accept the need for prices, the question becomes what basis we should use to determine them. This comes down to the question of what objective we would like the economy to serve. There are a variety of possible answers to this question, but the one that has come to prevail in our society is that the economy should aim at maximizing 'the satisfaction of human wants'. Extrapolating from this generates the familiar idea that the price of goods should reflect their *relative scarcity*.[7] Such prices are produced by balancing two considerations. First, how much people *want* a particular good, measured in terms of what they are willing to give up to obtain it, and, second, the opportunity cost of producing that good, measured in terms of what *other* wants could be satisfied through the production of some other good. These two constraints are better known as 'demand' and 'supply',

[7] Kornai (1992, 149-153) provides an interesting account of the consequences of deviating from this principle in the centrally planned economies of the former Soviet bloc. The major problem with their system of administratively determined prices, he argues, was that planners were given more than one principle to apply, which in turn generated contradictions. "One function of price in market coordination is to convey information in a concise form on the relative scarcity of resources and products. No such information is conveyed by the prices here described. In fact, they impart almost no useful information at all, as it is almost all lost in the conflict between the disparate pricing principles" (152). This observation is highly relevant to the present discussion, since, as we shall see, the demand for 'just' wages in many cases amounts to a demand that the determination of prices, in the case of wages, be done in accordance with more than one principle.

respectively. This generates, again, the familiar idea that the satisfaction of wants is maximized when the amount of 'want' satisfied by a particular use of resources is identical to the amount of 'want-satisfaction' foregone with other potential uses of those resources (since that makes it impossible to increase want-satisfaction by shifting resources out of one employment into another). This is to say that want-satisfaction is maximized when prices are set to the point at which supply is equal to demand.

It is important to recognize that, while we do not have much choice but to use some system of prices, the decision to have specific prices be determined by relative scarcity is very much a choice, based on a normative judgement about what should constitute the overarching objective of the economy. Most obviously, the existing arrangement is one that defers to individuals when it comes to determining what is to count as a 'want'.[8] The presence of a normative judgement here is sometimes obscured by the fact that the specific way this commitment is institutionalized in our society—namely, through a competitive market—operates in a decentralized fashion, without any central locus of planning or calculation. As a result, it may appear that 'scarcity prices' arise spontaneously, and therefore that they are part of the natural order. In this regard, the 'socialist calculation' debate of the early 20th century was quite illuminating, in that it showed how an entirely planned and obviously artificial order might still choose to use the principle of scarcity pricing as a basis for allocating resources and goods (see, for instance, Lange and Taylor 1938; Lerner 1944).

In any case, within an economy such as our own, in which all prices are scarcity prices, it is not difficult to explain why the wages earned by workers tend to be what they are. Wages are prices—in this case, the price of labour—and they are determined by more-or-less the same forces of supply and demand that determine every other price in the economy. Of course, they are also subject to various distortions and rigidities, including minimum wage legislation, cross-subsidization across employee groups within firms, as well as various forms of market power due to unionization or employer monopsony. Thus when I talk about 'market wages', what I am referring to is the general tendency of markets to push wages toward the level at which the supply of labour is

[8] The principal alternative would be an economy organized around some set of perfectionist commitments, where economic production was aimed at satisfying an 'objective' set of values.

equal to the demand for labour in a competitive market, and the price that this implies. The important point is that markets have no special way of rewarding labour. In principle, it gets treated like any other commodity.

If one were to ask what *justifies* any particular wage level, the answer would be straightforward. It would be the same as the answer given to anyone who inquires about the price of any other good in the economy. The scarcity price is the one that channels resources to their best employment, in terms of the satisfaction of human wants. Markets, of course, institutionalize this only imperfectly, but this is nevertheless the *principle* underlying the determination of wages. Deviation from the market wage will tend to generate misallocation of labour, so that workers will spend their time producing goods that, relatively speaking, people do not want so much, when they could have been spending their time producing goods that people want much more. This is a circuitous way of describing a situation in which some portion of their effort is *wasted*.

This answer is, as far as I am concerned, satisfactory, in the sense that it provides a plausible justification for the basic principle of wage-determination under capitalism, while at the same time explaining why deviation from that principle is likely to have undesirable consequences. It does, however, suffer from some deficiencies at the rhetorical level. The most important one is that it justifies a given wage rate by appealing to efficiency effects that are only manifest at the level of the economy as a whole. Furthermore, those effects are only felt in the medium term, in general and on the whole, and when most other prices in the economy are determined in the same way. Finally, because of the decentralized nature of production and price-determination under capitalism, it is often not possible to trace out explicitly the precise negative impacts caused by deviation from the market wage. Thus the justification of the wage appeals to very abstract properties of 'the system', and has practically nothing to say about the specific transaction that is being undertaken between employer and employee. And yet the transactional level is the one at which most people deploy whatever ideas they may have about fairness, morality, or justice in interpersonal relations.

Now of course there is a view that seeks to justify wages in purely transactional terms. This is the voluntaristic theory advanced by Robert Nozick (1974, 150-152), and defended in more ad hoc ways by many other libertarians. According to this view, given a just initial allocation of goods, whatever transactions people subsequently agree to are just, simply because there is no more to justice than what people voluntarily agree to. The problem with this argument is not that it fails to justify the rate of wages under capitalism, but rather that it justifies too much, including too many different wage rates. Indeed, it comes close to saying that 'whatever is, is good'. For example, it fails to provide any basis for preferring the wage rate determined in a competitive market over one in which some party has significant market power. Indeed, while Nozick had much to say about the importance of exchange, he had nothing to say about the importance of competition—which is arguably the more important institutional feature of capitalism. And yet, the inability to find anything wrong with monopoly pricing is a fairly major deficiency in any normative reconstruction of capitalism.

When we turn to the more conventional moral intuitions that people have about the market, what we find is that they often appeal to thicker, more robust principles, which they seek to apply at a transactional level. For example, many people believe that workers should be paid a 'fair' wage, or that compensation should be based on what individuals 'deserve'. It is here that most of the problems begin. There is a strong temptation to take categories drawn from everyday morality, used to organize small-scale individual interactions, and try to map them onto the movement of wages in a market economy, in order to declare the system 'just' or 'unjust'.[9] This is, I will argue, a mistake. Because specific market prices are not normatively patterned—but are rather the outcome of a complex system of incentives—it is simply inappropriate to evaluate them by applying everyday moral categories, which are typically oriented toward the evaluation of cooperative interactions. This is the intuition that was expressed, in an unsatisfactory way, by Hayek with his notion of the economy as a 'spontaneous order'. The point was put more perspicuously by Jürgen Habermas, using the distinction between 'system' and 'lifeworld'. In certain domains of interaction, what he refers to as 'lifeworld' contexts, outcomes are directly patterned by a system of shared norms or values. In other domains, which he refers to as 'systemic', outcomes are achieved indirectly through coordination of

[9] For some examples of this, one may consider the debate among business ethicists over sweatshop labour. As several commentators pointed out, many of the criticisms made of the wages paid to workers in underdeveloped countries are actually objections to the basic principle of wage determination in *all* markets. See Powell and Zwolinski (2012) and Maitland (1997).

action-incentives, and as a result, are not directly determined by the normative system (Habermas 1987, 154, 171).[10] One of the pitfalls that a critical social theory must avoid, he claims, is to evaluate 'systemic' outcomes using everyday 'lifeworld' categories of analysis. This does not mean that these outcomes are exempt from normative evaluation (the system must still be, as Habermas puts it, "anchored" [1987, 173] in the lifeworld). It simply means that they should not be evaluated naively, using thick concepts drawn from everyday morality, but must instead be evaluated in terms of overall system performance, using more formal or abstract concepts. It is this constraint that various conceptions of 'just' or 'fair' wages typically violate.

## III. THE PRODUCT OF LABOUR

Perhaps the oldest conception of justice in compensation is the idea that the *product* of labour constitutes its 'natural reward'. 'Whatsoever a man soweth, that shall he also reap'—workers should get what they have produced. It is easy to imagine this being something like a natural law, since an isolated individual (for instance, 'alone on a desert island'), would naturally enjoy certain benefits precisely to the extent that she was willing to labour to produce them. Yet it is also an ancient observation that, once two or more individuals begin to work together cooperatively, it becomes increasingly difficult to determine how much each person has contributed, especially if the forms of labour involved are heterogeneous. The problem becomes even more difficult—potentially intractable—when other factors of production are introduced, whose owners make some claim on the product. As capital goods came to play an increasing role in production, particularly with the introduction of industrial machinery, this problem began to be felt more acutely. As a result, everyday notions of what each individual has 'contributed' to the production process begin to fail us.[11]

This observation might easily have led to the conclusion that focusing on the 'product' of labour is simply not a good point of

[10] For discussion, see Heath (2011).

[11] This was the conclusion famously arrived at by Thomas Hodgskin in 1825. "Between the commencement of any joint operation, such as that of making cloth, and the division of its product among the different persons whose combined exertions have produced it, the judgment of men must intervene several times, and the question is, how much of this joint product should go to each of the individuals whose united labours produce it? I know no way of deciding this but by leaving it to be settled by the unfettered judgments of the labourers themselves" ([1825] 1969, 86). He concludes that it should therefore be left up to the "higgling of the market" (86).

departure for thinking about what constitutes a just wage. This might indeed have occurred, had it not been for the rise of 'marginalist' thinking in the late 19th century, along with the claim that wages, in a competitive equilibrium, will be equal to the 'marginal product' of labour. The idea soon arose that the contribution made by each factor of production at the margin might provide a basis for disaggregating their respective contributions, and, thus, serve as a principled basis for determining the entitlements of those who contribute each factor. This was the thesis defended most famously by John Bates Clark, who began his 1899 book *The Distribution of Wealth* by declaring it "the purpose of this work to show that the distribution of the income of society is controlled by a natural law, and that this law, if it worked without friction, would give to every agent of production the amount of wealth which that agent creates" ([1899] 1931, v). The law in question is what he calls the 'law of final productivity' (where 'final' refers to the last, or marginal, unit of any factor added to a production process).

The marginal-productivity theory of wage determination is simple to state, but conceptually can become rather tricky. The central idea is that, holding other factors of production constant, hiring one more worker will increase production by a given amount, which will in turn increase the firm's revenue. The gain is referred to as the marginal revenue product, and it will typically begin to decrease once the number of workers exceeds a certain threshold. Profit-maximizing firms will continue to hire workers so long as the marginal revenue product exceeds the marginal cost associated with hiring an additional worker. When the two are equal, the firm will stop hiring. It follows that, since the cost of hiring workers just is the wage, under equilibrium the wage will be equal to the marginal revenue product (which is to say, the value of the marginal product at prevailing prices) of the last worker hired.

The theory is fairly clear when stated in this way—even clearer when represented on a graph. Any attempt to translate these concepts into everyday terms, however, is fraught with difficulty. Clark, for instance, starts out his book claiming that this system "assigns to everyone what he has specifically produced" (v), which makes it sound as though, under such a system, each worker is to receive his or her *actual product*. This is an important ambiguity, since the phrase 'labour is paid its marginal product', *sounds* like saying, to any given worker, 'when you get hired, the firm will pay you a wage that is equal to the amount that *you* contribute to the firm's output'. It is easy to see, however, that this is not correct. Indeed, as the term 'final productivity' suggests, this system gives to each worker a wage equivalent to the contribution made to revenue by the *last* worker hired. Not only is this not the same as the actual contribution made by any of the infra-marginal workers, it will in the normal run of cases also be less than the *average* contribution made to production by each worker. [12] Since the marginal worker, by hypothesis, makes the least contribution to production, one might reasonably wonder why *everyone* should be paid a sum equal to that individual's contribution.[13]

The answer, roughly, is that everyone is paid that sum because everyone could become the last worker, simply by being fired. This is more intuitive if one thinks of the margin, not in terms of the firm adding workers, but rather in terms of subtracting them. The marginal product can then be defined as the amount that the firm would lose, by removing any one of its workers from the production process (assuming, of course, homogeneity of labour). Arthur Cecil Pigou, for example, offered a useful clarification of the concept in precisely these terms:

> The marginal net product of a factor of production is the difference that would be made to the aggregate product by withdrawing *any* (small) unit of the factor. The marginal unit is thus not any particular unit. Still less is it the worst unit in existence—the most incompetent workman who is employed at all—as some writers have supposed! It is *any* (small) unit out of the aggregate of units, *all exactly alike*, into which we imagine the aggregate to be divided. Though, however, the marginal unit is thus *any* unit, it is not any unit *however placed*. On the contrary it is any unit *conceived as placed at the margin* (1952, 133, emphasis in original).

The result, however, is that the 'marginal product' of labour is a hypothetical construct, one that does not exactly correspond to any of our intuitive ideas about what an individual can be said to have

---

[12] In this context, it is worth observing that worker co-ops would typically hire up until the point at which marginal revenue *per worker* was equal to marginal cost. It is only when the capitalist-owner is introduced that one gets wages set equal to *individual* marginal revenue product. For discussion, see Ward (1958). There are of course efficiency arguments for the latter arrangement.

[13] This argument does still come up in contemporary discussions. See, for instance, Schweickart (1980). Schweickart also makes the reverse argument, that capital is overpaid relative to its contribution, because each investor receives the high rate of return required to extract capital from the most reluctant (that is, marginal) investor. In later work he appears to abandon both arguments—see Schweickart (2002, 30).

contributed. Clark, it should be noted, acknowledges this in his more detailed discussion of wages, where he distinguishes the 'actual productivity' of labour from the 'effective productivity'. If all workers are interchangeable, then if any one worker should desert his position, the employer will immediately rearrange the assignment of workers to positions, so that only the least important job remains undone. A worker's effective productivity can then be defined as "the loss that his employer suffers when the man departs, and when the employer rearranges his force so that the more necessary kinds of work are still done" ([1899] 1931, 105). Because of this, "the effective productivity of any one of them is equal to the absolute productivity of the final or marginal one, whose work can best be dispensed with. We shall find that all wages are naturally gauged by the effective, rather than the absolute, productivity of the men who get them" (105).

This is all perfectly fine, as far as the economics of it are concerned. Yet many people may find that their moral convictions become somewhat attenuated in the passage from 'absolute' to 'effective' product, especially since the latter is defined in terms of a counterfactual.[14] It amounts to saying to the worker: 'you are not being paid an amount that reflects your actual contribution, but rather what you *would* have contributed, had a set of circumstances, which in fact do not obtain, actually obtained'. This does not seem like an explanation likely to stifle all objections, much less steamroll a committed union negotiator. Indeed, it helps to show how the connection between actual product and marginal product is, as Amartya Sen put it, essentially spurious. "It might, of course, be seen as a 'convenient fiction', but that fiction is a whole lot more convenient for some than for others" (1985, 16). Indeed, Alfred Marshall, in his much-admired *Preface* to Langford Price's book, *Industrial Peace* (1887), was much more cautious than Clark, choosing to describe the determination of wages under capitalism as potentially "fair" (x), but as necessarily falling short of being "absolutely just" (xi). He offered various reasons for this, most of them variants on the idea that, in a market economy, wages are simply

[14] It is notable that, in his discussion, Mankiw switches between discussing the 'marginal product' of labour, which evokes the neoclassical theory of wage-determination, and 'productivity' in the Mirrlees model (2013, 26), which is a quite different conception. In the Mirrlees model, more productive workers are modeled as though they were actually supplying a greater quantity of labour per unit of time worked, and are paid an amount that reflects this quantity (Mirrlees 1971, 176). Here workers are being paid in accordance with their *actual* productivity, but this is introduced as a modeling assumption, not as a conclusion of the analysis.

influenced by too many factors that are arbitrary from the moral point of view. The concept of a 'fair' wage was, in Marshall's usage, intended as faint praise.

There is, however, a much more significant problem with the way that the concept of marginal productivity has been understood among those hoping to make use of it in a normative argument. There is a temptation to think of the 'marginal product' as some sort of objective quantity, out there in the world, that *determines* the wage rate. This is, however, not correct—indeed, it is an instance of the fallacy of misplaced concreteness. It would be just as correct to think of the wage rate as determining the marginal product, although this is also potentially misleading. The correct thing to say would be simply that both the wage and the marginal productivity of labour are jointly determined under equilibrium. In other words, neither exists prior to the other, they are fixed simultaneously by the equilibrium, the point at which they are equal. As Daniel Hausman (1992) has observed, the *causal* relations could run either way—while one might increase wages by raising marginal productivity, one might also increase marginal productivity by raising wages.

This is a rather technical way of putting a point that can be given a more intuitive formulation. A key feature of wages is that, while they are just one more price, like any other in the economy, labour is also an input in the production of virtually all other goods. Thus a change in the price of labour has a significant impact on the price of almost everything else. A given firm's marginal revenue product curve is going to be affected by a number of these prices, and so the marginal productivity of labour is going to depend, in part, on the wage rate. More generally, this means that it will not be possible to draw the 'supply' and the 'demand' curves for labour, then look to see where they intersect to discover the market price, since changes in the price will tend to shift those very curves. As Hausman writes, "[i]n general one cannot sensibly consider what demand for labour would be, were the wage larger than it is, prices being what they are, because if the wage were larger, relative prices would not be what they are" (157).

All of this is a rather elaborate way of making the point that 'marginal productivity' does not mean what many people think it means, and certainly does not correspond to any plausible conception of 'how much a worker produces'. Once this is recognized, it goes a long way toward explaining a number of phenomena that casual observers of the market often find quite puzzling. For example, there is the fact that workers in different countries doing more-or-less the same job are often paid vastly different wages. This is often described, correctly, as a consequence of 'higher productivity' in the more developed country. Many people go on, however, to misinterpret this as the claim that specific workers, working for a particular company, earn more than their foreign counterparts, because the former actually produce more than the latter. This is obviously absurd.[15] Workers at an automobile factory in Mexico earn on average less than $4 per hour, while workers in a comparable factory in Canada typically earn $40 per hour (Valdenebro 2014, 25). And yet the Canadian workers are clearly not ten times more productive than the Mexicans, especially since the factories use approximately the same productive technology, not to mention the same work process. Workers in Canada earn more because the *average* productivity of Canadian workers in the economy as a whole is higher. The benefits of increased productivity are diffused across the entire labour force; they are not captured (for long) by any particular batch of workers.

Of course, if one thought that markets had some tendency to reward each worker based upon the amount that he or she actually produces, then one would be inclined to see the disparity in wages between Canadian and Mexican auto workers (or Italian and Korean shoemakers, or French and Chilean wine producers, and so on) as both flagrantly and self-evidently unjust. Not only that, but *maintaining* such an injustice, against the dominant tendency of the market, would seem to require massive global collusion, not to mention significant use of force. This is, in fact, how many left-wing critics of globalization see things, and is what fuels a number of popular conspiracy theories. The more prosaic explanation is simply that market wages do not reflect the 'contribution' that workers are making, in any concrete sense of the term, which is why workers in different countries, who are making what would appear to be exactly the same contribution, may nevertheless earn vastly different wages. If one thinks of wages as scarcity prices, this is entirely unmysterious. Because of the limited mobility of labour (and, for more complicated reasons, capital) across national borders, and because the

[15] The suggestion, in other words, is that the *piece rate* being earned by workers in different countries is the same, which is not true. The suggestion that it must be is usually based upon a failure to understand comparative advantage, and thus the assumption that any manufacturing facility located in a high-wage country must enjoy some absolute advantage over one in a low-wage country.

relative scarcity of different kinds of labour—and indeed, of labour in general—differs from country to country, one can find vastly different wage rates for workers producing essentially the same output.

## IV. TALENT, SKILLS, OR NATURAL ABILITY

Another popular theory of wage-determination, briefly mentioned by Mankiw (2013, 25), is based on the thought that wages are related to the 'talent' of the employee.[16] This specific claim is part of a broader family of views, which identifies employee ability—skills, training, talent, and so on—as a major determinant of wages. The thought is that these employees produce more value, which gives employers both the ability and the incentive to pay them more, in order to motivate them to greater work effort, and thus, to maximize mutual benefit.

At first glance, this theory may seem to be the same as the previous one—that employees are paid based on their productivity, so the more they produce the more they will earn. A moment's reflection, however, is sufficient to establish that they cannot be exactly the same theory, since the central presupposition of neoclassical wage theory is the homogeneity of labour. Indeed, it is precisely because any worker can be replaced by any other that earlier theorists considered it 'fair' to pay them the marginal product, rather than the actual. As soon as one starts to talk about 'talent', however, it is clear that we have abandoned homogeneity as an assumption, and hence moved out of the neoclassical framework. This may not be such a bad thing, since it also represents a significant move in the direction of greater realism—the world we live in is one in which there are often vast differences in ability between 'good' and 'bad' employees (Frank 1985, 59-61). The problem is that the discussion of 'talent' in many cases seems to move, not just outside the framework of neoclassical economics, but outside the framework of competitive labour markets entirely. Thus the discussion gets sidetracked into a debate over the disposition of *economic rents*, while ignoring the more fundamental questions about the way that ordinary wages are determined in a market economy.

[16] This argument seems to be a particularly tempting one for university professors, who may be inclined to regard the various comforts that they enjoy as a reward for being smart. (Without being *ad hominem*, I think it is worth noting that all of the major academic contributors to the argument over the incomes of the top one per cent are members of that group.) Thus Mankiw, for instance, puts considerable emphasis on IQ, despite that fact that, while IQ is a strong predictor of educational achievement, it is not a very strong predictor of future income (Strenze 2007, 415).

Much of the philosophical discussion of the relationship between wages and 'talent' seems to have its origin in some rather opaque remarks made by John Rawls, in *A Theory of Justice* (1999), in which he suggests that "the existing distribution of income and wealth" would be "the cumulative effect of prior distributions of natural assets—that is, natural talents and abilities—as these have been developed or left unrealized, and their use favored or disfavored over time by social circumstances and such chance contingencies as accident and good fortune" (72). Rawls goes on to argue that this distribution cannot provide the baseline for any compelling conception of "equality of opportunity", because "it permits distributive shares to be improperly influenced by these factors so arbitrary from a moral point of view" (72).

Even casual inspection of the key sentence should be enough to persuade anyone that there is a lot going on in this argument. There has in fact been significant disagreement among Rawls's commentators about exactly what he was claiming, and, more generally, how much of a role this argument played in motivating his more general position.[17] The argument acquired prominence mainly because Nozick dedicated 18 pages of *Anarchy, State and Utopia* (1974, 213-231) to a discussion of it. As an alternative way of thinking, Nozick put forward a theory of property rights founded on a principle of self-ownership, or that individuals have a right to control and to exclude others from the use of their own bodies. This suggests that, even if 'natural talents and abilities' are arbitrary from the moral point of view, each individual nevertheless has a right of control over his or her own talents. When combined with a principle that licenses voluntary transactions between individuals, this is sufficient to show that individuals have a right to contract with others for the exercise of those talents, and that whatever terms they agree upon are just *eo ipso*. As a result, if the talented are able to gain more from the sale of their labour than the untalented, there is nothing to be impugned in this arrangement from the standpoint of justice.[18]

[17] Some have claimed that it provides the key rationale for the difference principle. Others have claimed that it is a digression, which plays no role at all in motivating Rawls's central claims. See Sandel (1998, 73-82), Gorr (1983), and Pogge (1989, 73-81).
[18] This may explain why Mankiw makes the puzzling kidney-redistribution argument (2013, 32) against Rawls, who, being a contractualist, believes that principles of justice apply only to the fruits of cooperation, and thus excludes body parts. Mankiw may be thinking that anyone who is willing to countenance redistribution must be denying self-ownership, and so must also have no respect for personal or bodily integrity.

Although not fully explicit in his presentation, this is the normative framework that underlies Nozick's famous 'Wilt Chamberlain' argument (Nozick 1974, 160-164). He presents this argument in order to make a narrow point about how free contracting can disrupt 'patterns', and thus, how an 'end-state' theory of justice winds up being incompatible with liberty. Many of Nozick's readers, however, have seen in this argument a general defence of the type of economic inequality that arises in market economies, which they take to be a consequence of 'the talented' demanding higher wages, coupled with the threat that, unless they are paid more, they will work less. Critics have, therefore, spent a great deal of time and energy arguing that 'the talented' are not entitled to make such demands, simply because their natural abilities are supposedly undeserved, or 'arbitrary from the moral point of view'.[19] Seana Shiffrin, for example, compares the demand made by a talented person for greater compensation to a white person asking to be paid more merely because she is white. Both are instances of individuals "gaining advantage because of a feature that is arbitrary from a moral point of view" (2010, 135). G. A. Cohen compares it to the demands of a kidnapper, who refuses to return one's child unless a ransom is paid (2008, 38-41). Economic inequality, according to this view, is a consequence of individuals with superior natural abilities leveraging those endowments, through something akin to blackmail or extortion, in order to secure additional economic advantages.

If this were an accurate representation of how labour markets function, then the position one took toward natural endowments would wind up being of pivotal significance. And yet the situation that Nozick describes in his Wilt Chamberlain argument is really not a typical one. The desire on the part of spectators in his example is not just to watch a basketball game, but to watch *Wilt Chamberlain* play basketball. This is not a competitive labour market. On the contrary, Wilt Chamberlain is a monopolist in the market for Wilt Chamberlain services. He exercises market power, which is to say: he is, through his supply decisions, able to raise the price of those services. To the extent that he exercises this power, then some fraction of what he earns constitutes an economic rent—a payment that goes beyond what is needed to maintain the factor of production in that employment.

Mankiw, of course, was offering a defence of the top one per cent income earners, among whom one might expect to find some singular

[19] For discussion, see Lister (2017).

talents. Mankiw picks Steve Jobs, the former CEO of Apple Computer, as his preferred example. Yet this conversation, about whether it is acceptable for certain individuals to command a large economic rent, is quite distinct from the general debate over the way that markets determine wages, and whether the economic inequalities that result are acceptable. Here, the core principle that determines the level of compensation that employees receive—in the real world, in which labour markets are segmented, different individual possess different skills, and employees vary in their level of productivity—is not talent, but rather the *scarcity* of the relevant skills.

The central problem with the 'blackmail' model of the wage-determination is that, in a moderately competitive market, the talented are in no position to dictate terms to employers in the way that Wilt Chamberlain is with his fans. There is, of course, the fact that without demand for a particular service, no amount of 'talent' in the world can give it economic value. More importantly, however, if too many people possess a certain talent, then its economic value will also be quite low. As Frank Knight memorably observed, "[t]he value of a productive service varies from zero to indefinite magnitude, according to its scarcity. The most vital ministrations become valueless if offered in superabundance, and the most trivial performance becomes exceedingly valuable if sufficiently unique and rare, as when a human monstrosity satisfies an economic demand by letting people look at him" (1923, 599). Thus the economist's classic answer to the question why diamonds are expensive yet water is free applies with equal force to the determination of wages.

Theorists like Shiffrin and Cohen are of course aware that what counts as a 'talent' varies with demand, and thus over time. They do not appear to realize, however, that even given a certain level of demand, talent *as such* commands no particular economic return—it all depends on how many other people possess it. One kidnapper may be able to hold out for a ransom, but consider what happens if there are two kidnappers, who have a falling out after the crime is committed. They each have a key to the room where the child is being kept, and so each initiates independent negotiations with the parents over the amount of the ransom. Since neither has any use for the child as such, standard economic theory suggests that the parents should be able to negotiate the ransom down to nothing, or perhaps some small sum, sufficient to cover the cost of transporting the child back home.

Of course, to the extent that a particular talent is rare, that might serve as a source of scarcity, which will in turn tend to raise the wages of those who exercise it. But the higher wage is not a reward for superior talent; it is simply a consequence of the relative scarcity.[20] Indeed, thinking that talent results in higher wages is a clear example of mistaking correlation for causation, with scarcity being the confounding factor. To see this, consider the case of occupations that require a great deal of talent, skill or training, but where wages are quite low, because too many people possess the relevant qualifications. Symphony musicians, for instance, often regard their relatively low wages (frequently in the same range as police officers or firefighters) as a terrible injustice. They are, after all, supremely talented musicians, most of whom have spent their entire early lives competing in, and winning, talent competitions. The problem is that many parents have independent reasons for wanting their children to have the relevant training: the popularity of piano and violin lessons has very little to do with the market for piano players and violinists. As a result, too many people have the relevant training, which in turn drives down wages.

If one were looking for an intuitive way of thinking about the issue, it would be that wages are not determined by 'what you bring to the table', but rather by 'how easily you can be replaced'. If what you bring to the table makes it such that you are very difficult to replace, then it may result in higher wages. But if, all of a sudden, the market is flooded with new arrivals, able to do what you do just as well, then your market wage will tend to decline, even if the job that you do has not changed at all. This is also why wages tend to decline across the entire economy during periods of high unemployment, and to rise when unemployment is low.

Thus 'talent', at least in the everyday sense of the term, is not doing any direct work in determining wage levels. This is something that Cohen eventually came to acknowledge, when he designated as 'the talented' all those who "*are so positioned that, happily for them, they do command a high salary and they can vary their productivity according to exactly how high it is*" (2008, 120, emphasis in original). Now this is probably not quite what he meant to say, since all workers can vary their productivity in the non-technical sense of 'the amount that they

[20] As Robert Frank puts it, in the labour market "a person's fate will often depend much less on his ability in any absolute sense than on how his ability compares with the abilities of others" (1985, 175).

produce' simply by choosing to work more or less, depending on how high their wages are. What Cohen undoubtedly meant to pick out was someone like Chamberlain, who is able to increase his own rate of pay by threatening to work less. But this is just what it means to exercise market power. Thus when Cohen talks about 'the talented', what he really means is 'workers who have market power'. And yet the reason that we encourage competition in markets is to try to eliminate market power, so that prices will gravitate toward market-clearing levels. So whatever concerns there may be about the talented earning high salaries could be addressed simply by making the relevant labour markets more competitive.

As a result, when someone like Mankiw argues that the rich are merely being paid in accordance with their talents, it is overkill to respond, as Solow and others did, 'yes, but they have done nothing to *deserve* those talents!'. Many people regard their own talents and abilities as core features of their personal identity, and have spent considerable time—sometimes decades—cultivating them. In many cases the exercise of these talents constitutes their major life project. Being told that their talent, or perhaps the underlying aptitude, is arbitrary and unearned seems to undermine any basis of valuation. At very least, it is to make an extremely controversial claim. A much less controversial approach is simply to deny that wages are a reward for talent. For every story of how talent has been richly rewarded by the market, one can find a story of how markets have failed to reward some talent, or of how an untalented person has earned some rich reward. Thus the entire question of natural ability or talent is simply orthogonal to the debate over whether the particular wage rates determined by competitive markets are justifiable.

## VI. FAIRNESS IN WAGES

The discussion so far has been focused largely on moral interpretations of labour market dynamics that seek to apply some concept of 'desert', or the common-sense idea that whoever works harder, or contributes more, should be entitled to a larger share of the product.[21] There is, however, another influential approach, which focuses more on the transactional nature of the employment contract and asks whether the wages being paid are 'fair'. Naturally, if one defines fairness as simply 'each person getting what she deserves'—the position that Richard

[21] For a careful critique of the relevant conceptions of desert, see Olsaretti (2004).

Arneson refers to, barbarously, as 'desertitarianism'—then the two positions are not distinct. Yet the appeal to fairness is normally intended to frame the moral issue in a slightly different way. The exchange of labour for wages is an interaction that generates benefits for both parties—that is why it is undertaken. And yet, as with all such cooperative interactions, there is a question about how the benefits of cooperation are being divided up between the parties involved. The common-sense idea is that if one party derives disproportionately greater benefit from the transaction, then the terms of the transaction are 'unfair'. Since the most important determinant of how the benefits of cooperation are allocated is the price, this way of thinking is what underlies much of the traditional thinking about 'just prices'.

Both ancient and medieval discussions of just price regarded the transaction as something of a contest between buyer and seller, where the seller had an interest in obtaining the highest price possible, while the buyer would want the lowest price, and the question therefore became which price was justifiable from the moral point of view. In Book V of the *Nicomachean Ethics*, Aristotle suggested, and many were happy to follow him in thinking, that the price should be fixed at a level that generated some kind of equality, or at least a rough balance, in the benefits going to both sides. The kind of situation that struck them as a paradigmatic instance of *injustice* was when one side was, for some reason or another, much more in need of the transaction than the other, and, therefore, more likely to accept disadvantageous terms. For medieval theorists, writing at a time when an increase in the price of food would often provoke riots, this was not casual conjecture. The grain merchant does not encounter the hungry peasant on equal terms; the former is in a much better position to hold out for a better price than the latter.

It was this sort of asymmetry that concerned Thomas Aquinas, in his influential discussion of the 'just price'. He wrote that "[i]f one man derive a great advantage by becoming possessed of the other man's property, and the seller be not at a loss through being without that thing, the latter ought not to raise the price, because the advantage accruing to the buyer, is not due to the seller, but to a circumstance affecting the buyer" (1920, 1514). It is not too difficult to reconstruct Aquinas's basic intuition within a modern framework—such a reconstruction will also help to show what is wrong with it. Consider the well-known example, introduced by Bruno Frey, of a merchant raising the price of snow shovels in the aftermath of a winter storm. Frey found in his initial study that 83 per cent of respondents in Switzerland considered this price increase 'unfair' (Frey and Pommerehne 1993; Walsh and Lynch 2002). The intuition here is essentially the same as Aquinas's—because of the snow, the buyer's need for a shovel has increased, while the seller remains essentially unaffected. The value of the shovel has increased. Yet, by increasing the price, the seller essentially appropriates that increase in value, and thereby violates the prohibition on reaping where he has not sown.

The situation can be illustrated quite perspicuously using modern supply-demand diagrams. Figure 1 shows the benefit that both buyer and seller derive from a transaction at market-clearing prices. The triangle below the demand curve *D* and the price line *p* constitutes the buyer's surplus, or the welfare gain accruing to the buyer.[22] The triangle above the supply curve *S* and below the price line at *p* constitutes the seller's surplus.[23] Together these two triangles represent the *gain from trade*, or the *cooperative surplus* that is achieved through the exchange. Since the actual quantity of goods in the world remains unchanged, only the distribution of goods is affected, the cooperative surplus naturally takes the form of an *increase in welfare*.

[22] Intuitively, it is the difference between what an individual *would* have been willing to pay for the first, second, third (and so on) unit of the good, and what he *actually* has to pay, when the entire quantity *q* is purchased at price *p*.
[23] Again, it is the difference between what an individual *would* have been willing to sell the first, second, third (and so on) unit of the good for, and what she *actually* receives, when the entire quantity *q* is sold at price *p*.

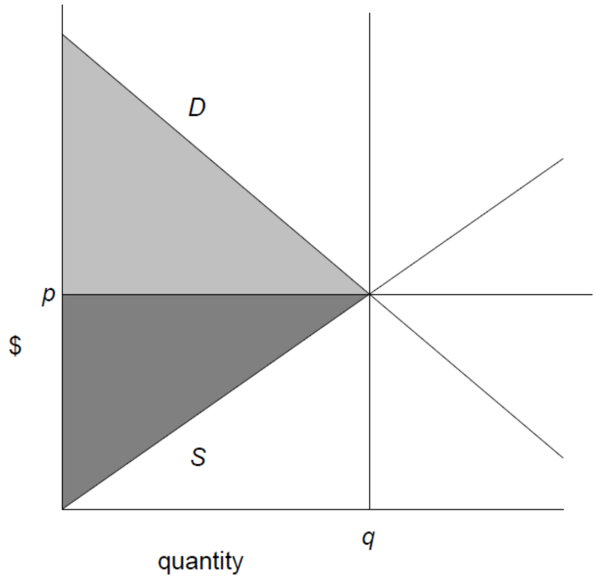


**Figure 1. Buyer's and seller's surplus**

Now, suppose that a snowstorm increases the buyer's need for a shovel. This will manifest itself in the form of a greater willingness to pay, which can be represented graphically as an upward shift in the demand curve, as in Figure 2 with the shift from *D* to *D'*.

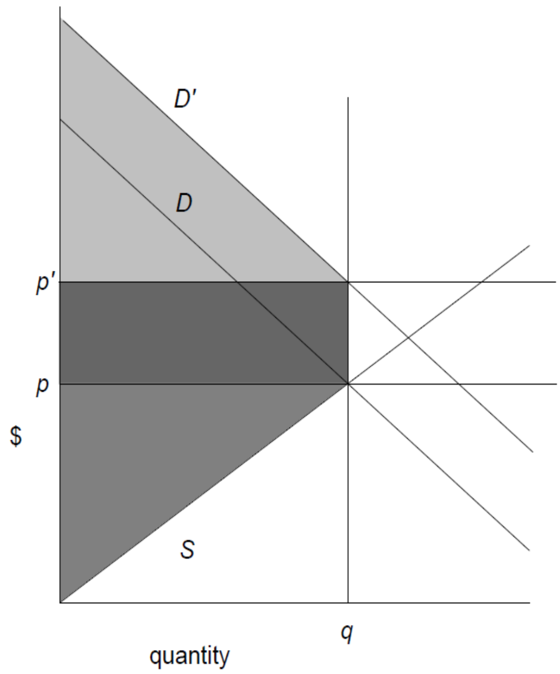


**Figure 2. Welfare effects of price increase**

Suppose now, sensing this increase in willingness to pay on the part of the buyer, the seller increases the sale price of the shovel to *p'*. What is noteworthy is that, although the size of the cooperative surplus has now increased, because the usefulness of the shovel has increased, the entire increase is appropriated by the seller.[24] This appears to be a clear-cut case of the seller benefiting from someone else's misfortune—now, not only does the buyer have to shovel the snow, but she has to pay more for a snow shovel as well!

There are certainly many objections that can be made to the suggestion that this is unfair. I do, however, think it is worth pausing for a moment to observe that the argument as it stands it not *unintelligible*, as some have pretended.[25] Whenever there is a cooperative interaction between two people, it seems like a perfectly reasonable question to ask how the benefits are being divided up between them. And whenever the benefits of changed circumstances are going entirely to one person, it also seems like a reasonable question to ask why. If the answer is 'because something bad happened to the other person', then it also seems reasonable to question whether that allocation of the surplus is *fair*, or whether it doesn't just represent one person taking advantage of his improved bargaining power in an unprincipled fashion. At the very least, the situation looks suspicious.

The most important objection to the Thomistic argument (at least under this reconstruction) is that it is myopic, since it ignores the effects that the price increase will have on the quantity of goods transacted. Indeed, it is undoubtedly the most normatively significant observation of early modern economists to have pointed out that, if the price increases, this will have the salutary effect of motivating individuals to bring more goods to market, at precisely the time that they are needed. So even if sellers do opportunistically increase their price to take advantage of the increase in demand, increased competition from new sellers bringing additional goods to market will tend to push the price down, until eventually it settles somewhere between the original price and the opportunistic one. This is illustrated in Figure 3, with the new price settling at *p''* and an increase in the quantity transacted to *q''*.

[24] The rectangle bordered by the two price lines, *p* and *p'*, the line at quantity *q* and the vertical axis, represents the increased cooperative surplus. At the higher price, *p'*, this entire utility gain goes to the seller, leaving the buyer's surplus unchanged.

[25] For discussion, see Moriarty (2012, 69).

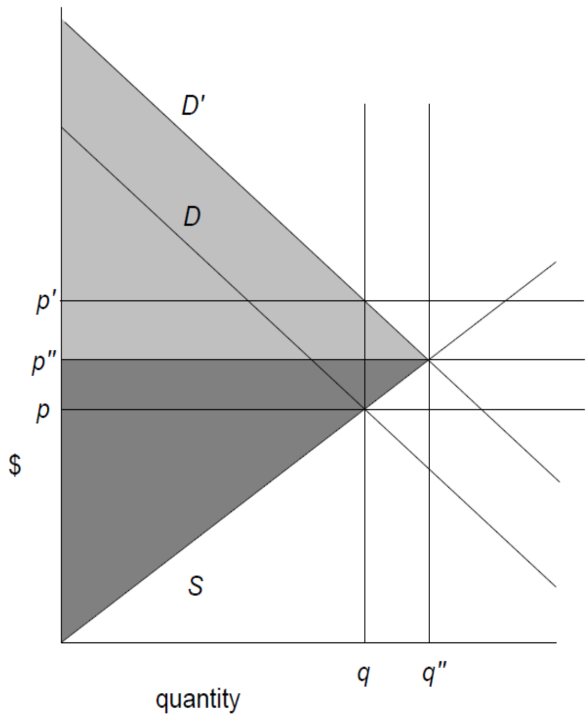


**Figure 3. Incentive effects of price increase**

The important point about this process is that, once it has run its course, it does result in both buyer and seller receiving some fraction of the cooperative surplus. Furthermore, both parties are making a clear contribution to the expansion of the cooperative surplus, since the sellers are not just raising their prices, they are also bringing new goods to market. At the same time, it is important to recognize that the eventual division that will be settled upon is not based upon any principle of distributive justice. It is, on the contrary, essentially arbitrary from the standpoint of distribution, since it will be determined entirely by the slope of the supply and the demand curves.

There was a time when it was popular among economists to claim that this argument showed that there was no such thing as a just price, or that determining the correct price level is simply not a moral question. Prices will be determined 'naturally', in accordance with the laws of supply and demand. This is, however, not the right way to think about it. What the argument shows is that, although the concept of a 'just' or 'fair' price is coherent, there is also an indissoluble tension between efficiency and equality in this domain. There is only one price level that is consistent with maximizing efficiency in the allocations of resources, and that is the market-clearing price. Unfortunately, this price level will result in a division of the cooperative surplus that is pretty much arbitrary from the standpoint of distributive justice. Furthermore, if we choose a price level based upon considerations of distributive justice (for instance, we adopt a policy of cheap bread, to ensure that the poor can eat), then the consequence will be a reduction in efficiency (for instance, the loss of mutual benefit represented by 'Harberger's triangle'—the area right of *q*, between *D'* and *S* in Figure 2). The decision not to forego this potential increase in welfare amounts to assigning efficiency priority over equality. It might best be described as 'ignoring the distributive effects of prices, in order to derive maximum benefit from their incentive effects'. This clearly involves a moral (or at least normative) judgement.

The important point, for the purposes of the present discussion, is that this policy is one that recommends itself, not just for snow shovels, but for all goods and services, across the entire economy, including wages. The average worker is in a situation very similar to the person in need of a snow shovel, in that she finds herself in rather strenuous need of a job, and has much less ability to 'hold out' than the average employer.[26] So, for example, we might consider it unfair that workers are forced to accept lower wages during a recession, or because they happen to live in a poor country. Yet, structurally, this is just the same as the intuition that the price of snow shovels should not be raised after a storm, with supply and demand reversed. The question, then, is whether we want wages to adjust over time, so that labour is directed to its best employment; or whether we are willing to sacrifice these efficiency gains, in order to achieve an outcome that is more attractive from the standpoint of distributive justice.

Now there are very few countries in the world in which there is not *some* interference in labour markets motivated by concerns over distributive justice. Minimum wage legislation provides the least controversial example. The point is that these constitute *interferences* in the labour market. Left to its own devices, there is no reason to think that the labour market will tend to produce wages that are 'fair' or 'just'. And to the extent that we do allow market forces free reign in this

[26] As Alan Manning has observed, it is common in our society for workers to celebrate when they find a job—for instance, taking their family and friends out for dinner and drinks—while it is far less common for employers to react in the same way after filling a vacancy (2003, 4).

domain, it is not because we consider the outcomes to be satisfactory from the standpoint of distributive justice, it is that we regard them as desirable from the standpoint of efficiency.

## VII. CONCLUSION

The preceding discussion has been conducted at a very simple level, both in terms of the moral intuitions being discussed, as well as the economic ideas about the way that supply and demand will affect wages. Obviously there are far more complex and nuanced models of the labour market available, not to mention more sophisticated and precise conceptions of 'justice' or 'desert'. There is no way to rule out the possibility that some exotic model of wage determination, combined with a more fully-developed conception of distributive justice, could wind up showing that the labour market operates as a system of natural justice. The discussion in this paper has been aimed merely at showing that the widespread habit of taking common-sense ideas (or 'lifeworld' norms), derived largely from the ethics of interpersonal interaction, and projecting them onto the operations of the market system, can easily be shown to be problematic. Ultimately, this is because the market is a system of decentralized decision-making, where individuals are given free play to follow incentives, subject only to rather loose legal and moral constraints. Thus it would be extremely surprising to discover that this system wound up conforming to principles that bear much similarity to the ones that govern face-to-face, deliberative decision-making.

Indeed, the best evidence that the two do not line up comes from the fact that, when individuals are given the power to decide wages collectively, in face-to-face deliberations, the outcomes that result tend to deviate from the market outcome in the direction of greater equality. For example, within large organizations—where there is more discretion about what wage will be paid because of the potential for cross-subsidization among employee groups—wages typically diverge from the market pattern. In particular, they have a pronounced egalitarian bias (see, for instance, Levine 1991; Lazear 1989). There is a large literature on the phenomenon of 'pay compression' in large firms. Collective bargaining is also a well-known source of wage compression (Freeman 1996). Thus whenever individuals are in a situation where the outcome they receive is directly 'patterned' by the prevailing set of norms, compensation tends to diverge in a predictable way from the pattern favoured by the market.

Now one might well wonder, if large firms and collective bargaining units are able to achieve much greater equality in the distribution of reward, whether that does not offer a model for the economy as a whole. Why not opt out of the principle of scarcity pricing, not with respect to ordinary goods, but just with labour and wages? Gender pay equity legislation, for instance, in many cases does not allow for any consideration of market conditions. In Canada, only four criteria may be taken into consideration when determining the 'value' of an employee's work: skill, effort, responsibility, and working conditions.[27] This means that, for example, an employer who is experiencing high turnover of employees in one occupational category, cannot increase the wage in just that one category, but must do so across all categories that have the same 'score'. There is nothing to say that one cannot try to reduce pay inequities through this sort of legislation. At the same time, it is important to recognize that the efficiency-equality trade-off described in section VI cannot be legislated away. As long as firms are still operating in a market economy, the dominant effect of pay equity legislation is to increase the costs associated with forming a large organization. After all, instead of having one firm with two classes of workers, there is always the option of splitting the firm in two (or 'outsourcing' one class of work), at which point pay equity legislation no longer constrains how much employees in the two classes must be paid relative to one another. Thus the level of compensation favoured by labour markets it still going to serve as an outside constraint, limiting the extent to which firms can diverge.

Consider the following, very concrete example of the phenomenon. In 2014, Google was subject to some negative media commentary, when it was discovered that security guards working at the 'Googleplex' were being denied certain perks that were available to other employees. For instance, although they were allowed to help themselves to a free lunch at one of the 25 cafeterias that provide complimentary food to all Google employees, they were being denied access to the 'takeout' boxes, which other employees could use to bring meals home with them.[28] This immediately set off an uproar over the treatment of these 'second-class citizens' of Silicon Valley. Google was able to tamp down some of this

[27] See Government of Canada (2016).
[28] See Sydell (2013). See the spin in Tiku (2013).

criticism, however, by drawing attention to a fact that had been somewhat buried in the initial round of media criticism, which was that the security guards in question were not actually Google employees, but worked for firm named SIS, to which Google had contracted out its security services. Thus, in principle, they were not even entitled to the free on-the-job lunches that they had been receiving.

One can see in this little morality play a number of the tensions that are at work in the way that markets determine wages. First, one can see how, within firms, inequality can produce considerable conflict. If some employees are being provided with free food, there is pressure to provide the same to all employees. There is something about working together, shoulder-by-shoulder, sharing the same physical workspace, that triggers the relevant norm. After all, no one was complaining about the fact that security guards down the road, working for another company, were not being given free takeout. And yet, if the rule is going to be 'all Google employees must receive the same benefits', and if the benefits are sufficiently expensive, then that creates a powerful incentive to contract out services where the market wage is too low to justify the benefit package. Again, this is not an ironclad constraint. Google, for instance, after all the negative publicity, decided to 'in-source' its security services, and to hire 200 of its own guards—who will presumably receive compensation packages that vastly exceed market norms. What it shows is simply that the market pattern of compensation, far from mirroring our intuitive conceptions of justice in compensation, instead serves as a source of constraint, preventing many individuals and firms from paying wages that those directly involved would regard as 'just'.

Again, none of this is to suggest that the market pattern of wage-determination is unjustifiable. In every society there is an enormous mismatch between 'jobs that need to be done' and 'jobs that people would like to do'. Achieving alignment of the two requires convincing millions of people to give up their hopes and dreams, a task that will necessarily involve the application of some very tough incentives, deployed on a vast scale. Part of the attraction of the market as an institution is that it accomplishes this more ruthlessly than any other system, largely because of the 'hard budget constraint' that firms face, which serves as a check on sentimentality. To expect then that the market should be able to achieve this enormous task, while at the same time producing outcomes that we regard as 'just' with regard to specific individuals, is to hold out hope for far too much. The market has one job to do, and it does that job very well. Producing 'just' wages, however, is not that job.

## REFERENCES

Aristotle. 1999. *Nicomachean Ethics.* Edited and translated by T. Irwin. Indianapolis, IN: Hackett.

Aquinas, St. Thomas. 1920. *The "Summa Theologica" of St. Thomas Aquinas.* Literally translated by the Fathers of the English Dominican Province. 2nd and revised edition. London: Burns, Oates & Washburne.

Berliner, Joseph S. 1999. *The Economics of the Good Society: The Variety of Economic Arrangements.* Hoboken, NJ: Wiley-Blackwell.

Blackburn, McKinley, and David Neumark. 1992. "Unobserved Ability, Efficiency Wages, and Interindustry Wage Differentials." *The Quarterly Journal of Economics* 107 (4): 1421-1436.

Clark, John Bates. 1899. *The Distribution of Wealth: A Theory of Wages, Interest and Profits.* New York, NY: The MacMillan Company, 1931.

Cohen, G. A. 2008. *Rescuing Justice and Equality.* Cambridge, MA: Harvard University Press.

The Economist. 2013. "The 1 Percent Needs Better Defenders." Accessed January 30, 2017. <https://www.economist.com/blogs/democracyinamerica/2013/06/inequality>.

Feinberg, Joel. 1970. *Doing and Deserving.* Princeton, NJ: Princeton University Press.

Frank, Robert H. 1985. *Choosing The Right Pond: Human Behaviour and the Quest for Status.* Oxford: Oxford University Press.

Frank, Robert H., and Philip J. Cook. 1996. *The Winner-Take-All Society.* New York, NY: Penguin.

Freeman, Richard B. 1996. "Labor Market Institutions and Earnings Inequality." *New England Economic Review* (May/June): 157-168.

Frey, Bruno S., and Werner W. Pommerehne. 1993. "On the Fairness of Pricing—An Empirical Survey Among the General Population." *Journal of Economic Behavior and Organization* 20 (3): 295-307.

Gorr, Michael. 1983. "Rawls on Natural Inequality." *The Philosophical Quarterly* 33 (130): 1-18.

Government of Canada. 2016. "Pay Equity." Accessed January 30, 2017. <https://www.canada.ca/en/employment-social-development/programs/pay-equity.html>.

Habermas, Jürgen. 1987. *The Theory of Communicative Action, vol. 2: Lifeworld and System: A Critique of Functionalist Reason.* Boston, MA: Beacon Press.

Hausman, Daniel M. 1992. "Supply and Demand Explanations and Their *Ceteris Paribus* Clauses." In his *Essays on Philosophy and Economic Methodology*, 147-162. Cambridge: Cambridge University Press.

Hayek, Friedrich A. 1976. *Law, Legislation and Liberty, vol. 2: The Mirage of Social Justice.* London: Routledge & Kegan Paul.

Heath, Joseph. 2011. "System and Lifeworld." In *Jürgen Habermas: Key Concepts*, edited by Barbara Fultner, 74-90. Durham: Acumen.

Hodgskin, Thomas. 1825. *Labour Defended Against the Claims of Capital*. New York, NY: A. M. Kelly, 1969.
Knight, Frank H. 1923. "The Ethics of Competition." *The Quarterly Journal of Economics* 37 (4): 579-624.
Kornai, János. 1992. *The Socialist System: The Political Economy of Communism.* Princeton, NJ: Princeton University Press.
Lange, Oskar, and Fred Taylor. 1938. *On the Economic Theory of Socialism.* Minneapolis, MN: University of Minnesota Press.
Lazear, Edward P. 1989. "Pay Equality and Industrial Politics." *Journal of Political Economy* 97 (3): 561-580.
Lerner, Abba P. 1944. *The Economics of Control.* New York, NY: MacMillan.
Levine, David I. 1991. "Cohesiveness, Productivity and Wage Dispersion." *Journal of Economic Behavior & Organization* 15 (2): 237-255.
Lister, Andrew. 2013. "The "Mirage" of Social Justice: Hayek Against (And For) Rawls." *Critical Review* 25 (3-4): 409-444.
Lister, Andrew. 2017. "Markets, Desert, and Reciprocity." *Politics, Philosophy, and Economics* 16 (1): 47-69.
Maitland, Ian. 1997. "The Great Non-debate Over International Sweatshops." *British Academy of Management Annual Conference Proceedings*: 240-265.
Manning, Alan. 2003. *Monopsony In Motion: Imperfect Competition in Labor Markets.* Princeton, NJ: Princeton University Press.
Mankiw, N. Gregory. 2013. "Defending the One Percent." *Journal of Economic Perspectives* 27 (3): 21-34.
Marshall, Alfred. 1887. "Preface." In *Industrial Peace, Its Advantages, Methods, and Difficulties; A Report of an Inquiry Made for the Toynbee Trustess*, by L.L. Price. London: Macmillan.
Mirrlees, J.A. 1971. "An Exploration of the Theory of Optimum Income Taxation." *Review of Economic Studies* 38 (2): 175-208.
Moriarty, Jeffrey. 2012. "Justice in Compensation: A Defense." *Business Ethics: A European Review* 21 (1): 64-76.
Nove, Alec. 1991. *The Economics of Feasible Socialism Revisited.* New York, NY: HarperCollins.
Nozick, Robert. 1974. *Anarchy, State, and Utopia.* New York, NY: Basic Books.
Olsaretti, Serena. 2004. *Liberty, Desert, and the Market.* Cambridge: Cambridge University Press.
Pigou, A. C. 1952. *The Economics of Welfare.* 4th edition. London: MacMillan.
Pogge, Thomas W. 1989. *Realizing Rawls.* Ithaca, NY: Cornell University Press.
Powell, Benjamin, and Matt Zwolinski. 2012. "The Ethical and Economic Case Against Sweatshop Labor: A Critical Assessment." *Journal of Business Ethics* 107 (4): 449-472.
Rawls, John. 1999. *A Theory of Justice*. Revised edition. Cambridge, MA: Harvard University Press.
Rosen, Sherwin. 1981. "The Economics of Superstars." *The American Economic Review* 71 (5): 845-858.
Sandel, Michael. 1998. *Liberalism and the Limits of Justice.* 2nd edition. New York, NY: Cambridge University Press.
Schweickart, David. 1980. *Capitalism or Worker Control?* New York, NY: Praeger.

Schweickart, David. 2002. *After Capitalism*. Oxford: Rowman and Littlefield.
Sen, Amartya. 1985. "The Moral Standing of the Market." *Social Philosophy and Policy* 2 (2): 1-19.
Shiffrin, Seana Valentine. 2010. "Incentives, Motives, and Talents." *Philosophy & Public Affairs* 38 (2): 111-142.
Solow, Robert. 2014. "Correspondence." *Journal of Economic Perspectives* 28 (1): 243-244.
Strenze, Tarmo. 2007. "Intelligence and Socioeconomic Success: A Meta-Analytic Review of Longitudinal Research." *Intelligence* 35 (5): 401-426.
Sydell, Laura. 2013. "Income Inequality in the San Francisco Bay Area: What It's Like to Live On Low Pay in a Land of Plenty." *National Public Radio*, December 18, 2013. Accessed January 30, 2017. <http://www.npr.org/sections/alltechconsidered/2013/12/17/251992536/ security-guards-at-big-tech-companies-struggle-with-low-pay>.
Thaler, Richard H. 1989. "Anomalies: Interindustry Wage Differentials." *Journal of Economic Perspectives* 3 (2): 181-193.
Tiku, Nitasha. 2013. "The Second Class Citizens of the Google Cafeteria." *Valleywag*, December 19, 2013. Accessed January 30, 2017. <http://valleywag.gawker.com/the-second-class-citizens-of-the-google-cafeteria-1486447319>.
Valdenebro, Alex C. 2014. "Explosión de la Industria Automotriz en México: De Sus Encadenamientos Actuals a Su Potencial Transformador." *Análisis* 1: 1-44.
Walsh, Adrian, and Tony Lynch. 2002. "The Very Idea of Justice in Pricing." *Business & Professional Ethics Journal* 21 (3 and 4): 3-25.
Ward, Benjamin. 1958. "The Firm in Illyria: Market Syndicalism." *The American Economic Review* 48 (4): 566-589.

**Joseph Heath** is Professor in the Department of Philosophy and the School of Public Policy and Governance at the University of Toronto. A fellow of the Royal Society of Canada and the Trudeau Foundation, Heath is the author of several books, both popular and academic. His most recent, *Morality, Competition and the Firm* (Oxford, 2014), is a collection of papers on business ethics and the normative foundations of market economies. He is also the author of *Enlightenment 2.0* (HarperCollins, 2014) and, with Andrew Potter, *The Rebel Sell* (HarperCollins, 2004).

Contact e-mail: <joseph.heath@utoronto.ca>
Website: <chass.utoronto.ca/~jheath>