6 TWO KINDS OF TOP PAY

Paul Ormerod[1]

Introduction

The dramatic rise in the remuneration of top earners relative to most of the labour force is well documented. Forty years ago, the typical compensation of a CEO in America was around thirty times greater than that of the average employee. By the mid 1990s, this ratio had risen to a ratio of 100 to 1, and now it is over 300 times as much. In the UK, to give a smaller-scale example, there is much controversy over the way in which the salaries of university vice-chancellors – university presidents in America – has grown enormously, even in the most undistinguished institutions.

These developments have attracted a great deal of adverse comment in the media. Popular resentment is high. Yet the even greater amounts of money made by both entrepreneurs and stars of popular culture appear to be accepted with equanimity.

For example, in 2018 (according to *Forbes*), George Clooney earned $239 million and Dwayne Johnson was the

1 I am very grateful to an anonymous referee for some stimulating and helpful comments.

second highest among male actors at $119 million.[2] The same magazine noted that, in the year to June 2018, the 100 best-paid athletes made $3.8 billion between them.[3] The boxer Floyd Mayweather topped the list with $285 million. Even these sums are of course dwarfed by the wealth of entrepreneurs such as Mark Zuckerberg and the late Steve Jobs.

The public's apparent unconcern about some people's fabulous incomes may seem paradoxical. However, I will argue shortly that economic theory offers a sound justification for these stupendous numbers in popular culture and in entrepreneurship.

By contrast the rewards of corporate board members are harder to justify, though there is an extensive literature in economics which attempts to do precisely that. It might be thought that this is a challenging task. Until relatively recently in the corporate world, to become seriously rich you had to be an entrepreneur and take risks with your own money. What we have seen with the rise in CEO remuneration is that mere employees who are only risking other people's money, not their own, have also been able to become seriously rich.[4]

2 The world's highest-paid actors 2018: George Clooney tops list with $239 million, *Forbes*, 22 August 2018 (https://www.forbes.com/sites/natalierobehmed/2018/08/22/the-worlds-highest-paid-actors-2018-george-clooney-tops-list-with-239-million/#7154e5e17dfd).

3 Full list: the world's highest-paid athletes 2018, *Forbes*, 13 June 2018 (https://www.forbes.com/sites/kurtbadenhausen/2018/06/13/full-list-the-worlds-highest-paid-athletes-2018/#915d89f7d9f9).

4 This applies to a wider range of employees in the banking and financial sectors, though I do not discuss them here.

In this chapter I consider arguments made by some leading economists in defence of corporate top pay levels. I also offer a broad perspective on the empirical evidence. But a distinguishing feature of my analysis is to make a case that the rise in corporate pay can be accounted for by the spread of social norms.

Top pay in popular culture

The top-ranked *Journal of Economic Perspectives* had a symposium of papers in one of its 2013 issues on the 'top one per cent'. The paper by the leading Harvard economist Greg Mankiw was explicitly entitled 'Defending the one per cent'.

Mankiw relies essentially on two arguments, both of which are based upon technology. He cites with approval the work of Goldin and Katz (2008), who argue that technological change usually increases the demand for skilled labour. Unless society is able to educate and train people so that the supply of skilled labour increases at least as much as the demand, the earnings of skilled workers will rise relative to the rest of the labour force. As Mankiw (2013: 23) puts it, 'the story of rising inequality there is not primarily about rent seeking, but rather about supply and demand'.

We might wonder what the impact of the huge growth in business schools in recent decades has been. These are meant to teach executive skills and so the supply of people capable of filling top executive roles might be presumed to have increased substantially.

Mankiw does concede that the Goldin and Katz arguments apply to the broad changes in inequality in general, and are not necessarily focused on the rewards of the top 1 per cent. But Mankiw invokes a further aspect of technology to account for why pay at the very top has grown so spectacularly.

Changes in technology have allowed a small number of highly educated and exceptionally talented individuals to command superstar incomes in a way which was not possible a generation ago. Brynjolfsson and McAfee argue this strongly in their book *Race Against the Machine*. Mankiw quotes from the book: 'aided by digital technologies, entrepreneurs, entertainment stars and financial executives have been able to leverage their talents across global markets and capture reward that would have been unimaginable in earlier times' (Brynjolfsson and McAfee 2011: 44).

A similar argument was used well before the digital revolution really took off, by Sherwin Rosen (1981) in his paper 'The economics of superstars'. This is the article invoked by Kaplan and Rauh (2013), also defending the pay of the top 1 per cent, in the same issue of the *Journal of Economic Perspectives* as the Mankiw article.

An example which Rosen gave very early on in his paper was 'sales of elementary textbooks in economics are concentrated on a group of best sellers, though there exist a large number of very good and highly substitutable alternatives in the market' (Rosen 1981: 845). This shows an amusing prescience, given that Mankiw is the author of what is perhaps the Number One basic textbook of our times.

Rosen's arguments are based on some subtle concepts in economic theory, and the paper itself rapidly becomes rather dense in its use of maths for the general reader. Even when set out in English, they require some effort to follow.

The key is what is described in the jargon of economics as a 'public good'. It is important to realise that the phrase is not being used here in any natural sense. It has a specific, technical meaning in economics.

A classic example of a public good in economics is defence. Individuals cannot effectively be excluded from consuming it, and the use of it by one person does not reduce the availability of it to others. To introduce some more jargon, a public good is 'non-excludable' and 'non-rival in consumption'.

So, for example, once a decision has been made to provide a country with a nuclear deterrent, no citizen can be excluded from the services which it provides. Indeed, in this case even the most ardent pacifist is obliged to 'consume' it. And the fact that one individual benefits from its existence does not mean that anyone else is prevented from benefiting to the same extent. If I buy up all the bananas on a market stall, until the seller restocks the stall, no one else can buy them. But this is not the case with a public good.

In essence, public goods are not consumed individually, but jointly.

Even in the pre-internet era, radio, television and the telephone had greatly increased the level of connectivity in society. A hundred years ago, for example, the only people who could have any direct experience of Manchester United playing soccer live were those present in the stadium during the game. In 1927, the BBC began broadcasting live football

commentary on radio. Although the BBC[5] began transmitting soccer on television in 1938, the games were few and far between, and regular coverage did not start until the early 1960s. It was only around that time that more than 50 per cent of UK households had a TV. Now, of course, Manchester United can be watched by literally billions around the world, using a variety of delivery channels.

Rosen argued that the provision of many cultural services, for example (and using the word 'culture' in a wide sense) involves joint consumption, not unlike a public good. A performer or author must make the same effort almost regardless of whether 100 or 100,000 people watch him or her, or read the book. The costs of production do not really rise with the size of the market.

However, the difference between this technology and a public good is that people can be, and are, excluded from consumption. Unless you pay, you don't get to see the show or read the book. But the joint consumption features of these products and services means that a relatively small number of sellers can in principle service the entire market. And the more talented they are, the fewer still are needed.

So, as Rosen (1981: 847) puts it, 'the possibility for talented persons to command both very large markets and very large incomes is apparent'.

Kaplan and Rauh argue that, in the age of the internet, the arguments made by Rosen have become even more powerful. Professional athletes and artists, for example, can now reach much bigger markets than ever before.

5 Which then held a domestic monopoly on radio and television in the UK.

Technology enables Wall Street investors and executives to acquire information and trade in hitherto unimaginable amounts.

There is considerable validity in the opinions offered by Mankiw and (via Rosen) Kaplan and Rauh. But these arguments are much more relevant to stars of popular culture and to entrepreneurs than they are to the board members of massive companies which have been around for a considerable amount of time.

As far as entrepreneurs are concerned, we can also invoke the argument made by Schumpeter (1934). A successful innovation, which provides a product or service which did not previously exist, enables the entrepreneur to earn monopoly profits until such time as competitors are able either to replicate the innovation or to produce a superior competitor.

The theoretical justifications for the earnings of pop stars, actors, athletes and entrepreneurs appear to be well founded. In practice, people do not seem to begrudge these individuals their rewards, vast though they may be. They are perceived as being based upon individual merit.

Top pay and corporate executives

In the simple models of basic economics textbooks, individuals are rewarded in line with their productivity. The more value they add to the organisation, the more they get paid.

In the jargon of economics, this is known as marginal revenue productivity theory. In competitive conditions, it

is asserted, employees will receive the value of their marginal contribution to a firm's revenue. Straightforward expositions of this proposition are readily available on the internet, so there is no need to expound the argument in detail here.

Marginal contributions may be fairly easy to see on factory production lines, in restaurants or in fruit picking. When we are considering the pay of, say, CEOs, the theory is more complicated. It reflects an important discussion in economics which dates back to the end of the nineteenth century, and is still very much alive today.

In the decades before World War I, two highly accomplished mathematicians who occupied the top chairs in economics at the time, Alfred Marshall at Cambridge and Francis Edgeworth at Oxford, wrangled over the issue.

Edgeworth thought that, in most situations, there was an inherent indeterminacy about the price which emerged – in the current context the 'price' is the salary of the CEO. He wrote: 'It may be said that in pure economics there is only one theorem, but that it is a very difficult one: the theory of bargain'.[6]

Marshall simplified matters dramatically. He assumed there are so many economic agents in a market that no single one of them can influence the price. This enabled him to draw, in his own best-selling textbook, the supply and demand curve diagrams familiar to generations of students.

6 F. Y. Edgeworth, 'On the application of mathematics to political economy', reproduced in Marchionatti (2004: 137).

At the time, Marshall prevailed. His textbook dominated the teaching of economics for decades. In recent decades, Edgeworth has returned with a vengeance. A lot of modern economic theory is about developing Edgeworth's view that economics is basically about bargaining. It makes theory much more difficult, but potentially more realistic.

The implication is that there may be a wide range of possible outcomes to any given bargaining process, rather than there being a unique one dictated by marginal productivity theory.

Nevertheless, the idea that executive contracts somehow represent an optimal outcome retains a strong following in economics.

A recent and very detailed survey of the literature on executive pay was carried out by Edmans et al. (2017). They spend a considerable time discussing what they call the 'shareholder value' view of CEO pay. This proposes that CEO contracts are the outcome of shareholder value-maximising firms competing in an efficient market for managerial talent.

The authors immediately qualify the concept of optimality in this context. For example, the optimal contract theoretically may be highly non-linear and never be observed in reality. More generally, they describe what they term 'bounded rationality' where boards are not aware of some performance measures that 'could theoretically improve the contract if included'.

A reasonable interpretation of the latter phrase is that it describes imperfect information rather than bounded

rationality in the sense which Nobel Laureate Herbert Simon used the phrase. Simon (1955) believed that in most practical situations it is simply not possible to know the optimal strategy to follow or the optimal decision to make. It is not a matter of a simple lack of information. The environment in which agents take decisions is too complex at a point in time and it evolves in too unpredictable ways over time, for the concept of optimality to have any meaning.

Nevertheless, even though optimality may not be a terribly useful concept in this context, it is of course still possible that the main motivation in determining executive contracts is an attempt to increase shareholder value.

Edmans and his colleagues suggest that there are three main hypotheses put forward in the literature to account for the huge increase in executive pay in recent decades.

The first is the 'shareholder value' one already mentioned. The second is the 'rent extraction' view, which argues that contracts are set by executives themselves to maximise their own rents. The final perspective is that pay is shaped by institutional forces, such as regulation, tax and accounting policies.

The conclusion which is drawn is firm, while at the same time being heavily guarded. From an empirical perspective, the authors argue that no single hypothesis can explain all the evidence. The outcomes which are observed are some combination of all three.

They also emphasise further limitations to our ability to draw unequivocal conclusions. For example, much of the formal theory around the issue has been developed for the 'shareholder value' approach. But they point out

that seemingly innocuous differences in assumptions in models can lead to quite different outcomes. Further, this hypothesis is consistent with a wide range of potential empirical outcomes.

Much of the evidence discussed in the paper is, however, of a detailed and technical nature. In the next section, I set the discussion around executive pay in a broader context.

A broad perspective on the empirical evidence

We might usefully begin with a thought experiment. Imagine someone working for, say, $10 an hour. We wave a wand, and are in a world where everything else remains unchanged but the same worker is now paid just $5 an hour. We could reasonably expect an adverse impact on this person's motivation.

Now consider someone receiving $10.1 million a year, which according to Edmans and colleagues was the median CEO compensation in the S&P 500 companies in 2014. Our wand performs the same trick, and this is reduced to $5.05 million. It requires more imagination to believe that someone paid this amount would somehow be insufficiently motivated to perform his or her job.

The fact is that, for long periods of time, CEOs and other top executives were paid very considerably less than they are now, and the economy nevertheless performed well.

Edmans and colleagues cite calculations by Frydman and Saks (2010) on the real levels of compensation of the three highest paid executives in the fifty largest US companies since the late 1930s. From the end of World War II to

around 1970, the median level of total compensation was just under $2 million a year in 2014 dollars. It is now well in excess of $10 million.

It may of course be important to an individual company that Mrs Smith rather than Mr Jones is appointed as CEO, and the markets can and do take an interest in such things. But the relevant context is not that of the individual company, but of the corporate sector as a whole, and within this the large companies which dominate the economy.

Given the increase in CEO compensation, we would expect to see an improved performance of the overall economy in recent decades. The evidence, however, points firmly in the opposite direction. The rate of growth of the economy as a whole has slowed. Between 1957 and 1987, real GDP in the US grew by 3.5 per cent a year, but by only 2.5 per cent between 1987 and 2017.

GDP of course measures the output of the economy, the amount of goods and services which are produced. The value of companies is more complex, and is not simply related to how fast their revenues are growing. It depends upon the expected future stream of dividends which investors might receive in return for holding equities.

Strong revenue growth might well increase the value of a company. But investors will look not just at the actual profits which are made on those revenues – because dividends can only be paid from profits – but will also form expectations about the future growth in profits.

Undoubtedly, the market value of major companies has grown spectacularly in recent decades. Stock markets across the West have boomed over the past three decades,

the crash of the late 2000s notwithstanding. In 1987, the Dow Jones stood at – using very round numbers here to make comparisons easy – about 2,000. In 2017 it averaged some 22,000: an elevenfold increase, which represents a dramatic acceleration of growth compared with the three decades immediately preceding, from 1957 to 1987. Over this earlier period, the Dow Jones rose from 500 to 2,000.

The contrast is even more marked when we examine the average annual rate of growth in real terms, after allowing for inflation. Between 1957 and 1987, the Dow Jones rose by 4.6 per cent a year, but inflation averaged 4.3 per cent. Effectively, it was more or less unchanged in real terms. For the period 1987–2017, the comparable numbers are 8.3 per cent and 2.1 per cent, a real increase of over 6 per cent a year.

Valuing financial assets is of course a complex business. The American economist Robert Shiller got the Nobel Prize for his analysis of financial markets. Using over a century's worth of data, he showed that the short-term fluctuations in share prices were much greater than those of the dividend streams which were paid out.[7]

But over longer periods, certainly over several decades, the excess swings in equity prices, both up and down, should to a substantial extent cancel each other out. We should expect to see over the longer run a closer, albeit still not perfect, relationship between changes in the stock market and changes in the prospect of firms' earnings.

7 Robert J. Shiller Prize Lecture, The Nobel Prize, 8 December 2013 (https://www.nobelprize.org/prizes/economics/2013/shiller/lecture/).

The rate at which output is expanding has slowed in recent decades, so we might imagine that the growth in the earnings of companies had also slowed. However, the amount of profit which companies make from any given level of output has risen.

The share of wages in national income was essentially the same in the late 1980s as it was in the late 1950s. But since then, there has been a fall of some four percentage points, and a similar rise in the profit share.[8] This may seem small, but in money terms it amounts to almost $1 trillion a year.

The rise in asset prices – and hence the value of companies – is underpinned by the fact that profits have grown faster than the economy as a whole over the past three decades.

Might this, then, be a rational explanation for the boom in top executive pay? The stock market value of companies has also boomed, and executive remuneration has risen commensurately.

This explanation is, however, rather undermined by a detailed analysis of CEO pay at the individual firm level carried out by Ethan Rouen of Harvard Business School. Rouen (2017) obtained confidential establishment-level annual data provided by the US Bureau of Labor Statistics for a large sample of firms in the S&P 1,500 from 2006 to 2013. He found 'no statistically significant relation between the ratio of CEO-to-mean employee compensation

8 In the national economic accounts, there are some minor categories of income other than wages and profits, so that changes in the two are not exactly equal and opposite.

and performance'. In other words, in publicly traded companies there was no connection between how much the CEO was paid relative to the average worker and how well the company did.

More generally, there has been a very powerful factor operating across the developed world over the past three decades or so which has held wages down relative to profits. The absorption of China and India into the global economy from around 1990 onwards added over a billion workers to world labour supply. This process created downward pressure on wage rates, particularly among the less skilled.

Less dramatically, many of the countries in Eastern Europe which were under Soviet domination entered into the capitalist world after the fall of the Berlin Wall in 1989. Christian Dustmann of University College London and colleagues provide detailed evidence on the impact on wages in the old West Germany of the opening up of economies such as Poland and the Czech Republic.

The implication of their empirical evidence is that CEOs *in general* have not created the massive rise in stock markets, and therefore the value of companies, by their skill and enterprise. They appear to have simply ridden on the backs of the powerful economic forces which generated the growth in the profitability of companies (Dustmann et al. 2014).

A network perspective

Earlier we argued that many prices – and executive compensation is one such price – are set not at a unique level

determined by supply and demand, but by a process of bargaining. There is a potentially wide range of indeterminacy in the price which emerges from any process of bargaining.

The existing literature within economics on executive pay largely neglects the role of social norms, and how these may evolve and influence the bargaining process.

I've combined these two sentences - OK?

One of the most rapidly growing areas of knowledge in the past two decades or so has been the field of networks, more specifically, knowledge about how ideas, beliefs, behaviours and the like either spread or are contained in networks of connected agents.

The phrase 'connected agents' means agents who are connected to each other in the sense that one, or both, of them has the potential to influence the behaviour of the other in a specific context. The group of others who may influence any given agent may, indeed almost certainly will, vary from setting to setting. I may, for example, pay attention to the opinions of certain people on restaurants, but when it comes to thinking about financial products, I look to an entirely different set of people.

Two key assumptions of standard economic theory are that the tastes and preferences of each 'agent' (individual or business entity) are both formed independently and do not vary over time. The decisions which an agent takes can certainly be affected by what others do, but only indirectly via prices. I may like bananas and go to the market intending to buy some, only to find that demand has been high today and the price has gone up so much that I choose not to buy any. But my liking for bananas has not altered. If the price is lower tomorrow, I will buy.

In a networked world of the type discussed here, these assumptions no longer hold. A seminal paper, many years ahead of its time, was by Nobel Laureate Thomas Schelling (1973). With the intriguing title 'Hockey helmets, concealed weapons and daylight-saving time', Schelling was inspired by a story in the sports stages of his local newspaper.

An ice hockey player had not been wearing a helmet and had suffered serious head injuries when struck by the puck. The rational choice was to wear a helmet. When a star player was asked why he continued not to wear one, he replied 'I don't because the other guys don't'. In other words, his preferences were not fixed. They were determined by the behaviour of others.

Twenty years later, the now famous paper by Bikhchandani and colleagues (1992) described how information cascades can grow through 'rational herding' in a sequential social learning process, with each agent balancing what he or she already knows against what others can be seen to be doing.

A landmark paper on how behaviours and opinions spread across networks was published by the mathematical sociologist Duncan Watts. In this model (Watts 2002), agents pay no attention at all to the attributes of the alternatives presented to them, in complete contrast to standard economic theory. Their decisions are based entirely on what the agents to which they are connected – those which can potentially influence them – do.

A key point from both the Bikhchandani et al. paper and the much more formal analysis by Watts is that the

information which spreads, the idea which becomes popular, are not necessarily those which have superior qualities to the rest.

Indeed, in the Watts paper by definition no single alternative is better than any of the others. Agents can only be in one of two 'states of the world'. To give a contemporary UK example, an agent could either want to remain in the EU or be in favour of leaving. By construction in Watts's model, at the start all agents hold the same view. Then a few, at random, change their mind. Agents base their decisions solely on the opinions of those to which they are connected.

Most of the time, the switch of opinion fizzles out. But occasionally, there is a 'global cascade' across the network, and almost all agents alter their initial opinion. In essence, the eventual size of the change depends upon some rather subtle mathematical properties of the structure of the network.

Of course, in practice things are much more complex. Scientific models such as those of Bikhchandani and Watts make deliberate simplifications in order to get a better understanding of reality.

The key point to draw from this is that the optimal choice among the alternatives in any given situation does not necessarily win out. This goes against the arguments put forward by, for example, Mankiw. He implicitly assumes that the market always leads to optimal outcomes. Modern technology simply enables rewards to a small number of talented individuals to be leveraged on a massive scale.

There could be justification for the size of the reward if the best really did emerge. In sport, for example, it is usually pretty clear. Either Bolt runs faster than you or he does not. It is much less clear cut among CEOs as to who is actually the best.

What executives *have* been really good at is in ensuring that the narrative that they in some way deserve their pay has become the dominant one in the relevant network. The network here is the pool of people from whom non-executive directors, management consultants with the large firms, remuneration 'experts' and the like are drawn.

Before the financial crisis, Piketty and Saez (2006: 204) advanced the argument that executive pay had grown so rapidly because of 'the increased ability of executives to set their own pay and extract rents at the expense of shareholders'. They did not suggest exactly how this had happened. But the view is consistent with the idea that a set of values which had previously held top pay in check has been replaced, in the relevant network, by a set in which traditional constraints no longer obtain.

Conclusion

Although the arguments of this chapter have perhaps been complex, the conclusion to be drawn can be stated rather simply.

The dramatic rise in the remuneration of athletes, film stars, musicians and entrepreneurs has a sound justification in economic theory. It has been created by the huge

advances in technology, principally communications technology in the broader sense.

The increase in executive pay, in contrast, is hard to justify, whether from a theoretical or from an empirical perspective. In my view, it is difficult to escape the conclusion that it has been primarily based on successful rent-seeking.