

Reflections on Macroeconomics and Share Systems

Martin L. Weitzman

It was with a strange mixture of delight and dismay that I set about trying to reflect on my 'contributions' to the *Economic Journal (EJ)* (the articles 'Increasing Returns and the Foundations of Macroeconomic Theory' in December 1982 and 'Some Macroeconomic Implications of Alternative Compensation Systems' in December 1983) – delight because being asked is obviously an honour; dismay because, well, what was I going to say that I hadn't already said before, in the recent past, and typically more than once (Weitzman, 1984, 1985).

For better or worse I decided to build my reflections around a narrative of how these ideas developed in my own mind and how with benefit of hindsight I might recast their essence today. That is at least superficially different from repeating the models themselves. And such introspective personal history leads naturally to introspective personal ramblings about what it all means. Since the latter is what I think I am supposed to be doing here, this way seems as good as any other.

Involuntary unemployment has been, is now, and perhaps always will be an extraordinarily difficult concept for economists (and almost anyone else) to handle. It seems that just about every sort of explanation has been tried but no single approach comes

up with an entirely consistent, convincing, consensus-building account. In prosperous times, it is perhaps not so entirely preposterous to think of unemployment as an economically normal, largely voluntary, free choice activity. The rational expectations school has made this sort of market-clearing equilibrium vision the centrepiece of a now well-known paradigm of unemployment and business cycles. For a wide variety of reasons, most of which have been articulated by others in the literature, I have not found this approach very palatable as a description of economic reality even during relatively good times. And during bad times the idea that very high unemployment rates are essentially voluntary seems way off the mark. Then the phrase 'coordination failure' seems to me to have a much more authentic ring. But what, exactly, is the idea supposed to mean?

If there is such a thing as a coordination failure, then the Great Depression was surely the biggest coordination failure of all time. This extreme event has long been in the back of my mind, asking to be explained. The puzzle of it all served as an image of what I thought I was trying to model.

I was too young to live through the Great Depression, but I feel as if I practically did because of the vivid profound effects it had on my parents and their contemporaries. As the orphaned son of a very poor tubercular mother, I later perceived some vague connection between our overall condition and the hard times everyone seemed to have passed through. My adoptive parents were among those who experienced first hand the mass unemployment of the 1930s, and it left them with a bitter cynicism about the ability of the capitalist system to provide a stable environment and to meet basic needs. To them the market system seemed irrationally and dangerously out of control. While their radical views moderated over time, as a young boy growing up in the 1940s and 1950s I was regaled with plenty of first-hand stories about 'failures of coordination'.

What, to my father, was the cause of the Great Depression? His basic story of what was actually happening comes down to me through the years as something like the following. The shoe factory was shut down because customers could not afford to buy shoes. The unemployed worker who needed shoes could not afford to buy them because he did not have a job. Since the system was unplanned and anarchistic, there was no good way to get workers,

factories and buyers together. If the worker would agree to take all of his pay in some of the shoes he produced, then just maybe something could be arranged. But that, of course, was ridiculous because what would the worker do with all those shoes?

The paper 'Increasing Returns and the Foundations of Unemployment Theory' was conceived in the spring of 1981 and a first draft was written that summer. The basic idea was then in the air. Around that time, Peter Diamond's (1982) seminal paper 'Aggregate Demand Management in Search Equilibrium' had a catalysing influence on my thinking.

Diamond's work showed that multiple equilibria could be endemic in a non-Walrasian environment. His milieu was a search-theoretic coconut economy. Individuals could either work for themselves and abstain from the market or shake down coconuts from trees. Unfortunately, anyone's shaken-down coconuts could not be consumed right away but had to be traded for somebody else's shaken-down coconuts, which could then be consumed. This was meant to reflect some sort of division of labour specialization.

The model displays a kind of externality that gives rise to multiple equilibria. If a lot of people are shaking down coconuts, it will be more worthwhile for me to shake down coconuts and try to trade with them because it will be relatively easy for me to find a trading partner. But if others are not shaking down coconuts, then neither should I, because it will be difficult to locate a trading partner and I may be stuck with my own coconuts, which I cannot consume.

Diamond's model struck me as an extremely interesting and perhaps important allegory. But it was at too high a level of abstraction for my tastes. What, exactly, was the source of the division of labour trading constraint? Where did preferences fit in? Where were prices and wages? What was the role of increasing returns (which the system as a whole displayed)? Where were ordinary looking firms and workers?

This intellectual inspiration led me to think seriously about how to integrate monopolistic competition into a general equilibrium type framework that would permit closer examination of Keynesian style coordination failures. It had always seemed to me, as it did to many others, that the natural habitat of a Keynesian macroeconomy is an imperfectly competitive microeconomy, just as classical macroeconomics and perfect competition are natural counterparts.

But the details were not really worked out. In principle it seemed as if it should be interesting to look carefully at the macroeconomic side of a Hotelling–Chamberlin monopolistically competitive circle economy. A model with endogenously determined product diversity and division of labour would appear to be a natural vehicle for studying failures of coordination. The heart of the matter seemed to me to be the idea that, while firms are specialists in production, workers are generalists in consumption; therefore an unemployed worker is unable to communicate to his potential employer his effective demand. Such failures of coordination are difficult to describe, much less to explain, in an artificially aggregated economy that produces essentially one good. From the beginning I thought that modelling the failure of coordination implicit in an inability to communicate effective demand required increasing returns and product diversity. It is true that, with enough ingenuity, unemployment can be generated in models by various forms of asymmetric information combined with adverse selection or moral hazard, but division of labour struck me as a rather more relevant foundation.

Naturally, with benefit of hindsight I would have presented the model somewhat differently. The transportation cost function on the circumference was assumed for simplicity to be linear; it makes more sense and is even neater for the appropriate loss function to be of exponential decay. Because I wanted to emphasize features of the real economy, I left out monetary considerations altogether; in retrospect it would have been better to bite this bullet by building in some crude representation of a monetary economy. A more serious drawback is that the model was vague about what exactly is supposed to be the definition of an unemployment equilibrium and just what had been proved. At that time I took it as sufficient to show how the model was hinting strongly at some problems in automatically adjusting to full employment, and left it at that. Since then there have been a number of serious attempts to deal carefully with expectations and to give a rigorous game-theoretic account of unemployment equilibrium.¹

Despite these and other misgivings, I believe there remains a basic germ of truth from this paper concerning the relative difficulty of full employment adjustments in an environment of increasing returns and imperfect competition. What comes out of the model, I hope, is a framework explaining how increasing returns,

product diversity, monopolistic competition, failures of coordination and hysteresis-like unemployment all fit together in a fairly consistent story. With benefit of hindsight, let me restate the essential argument about the macroeconomic role of increasing returns.

Suppose, just for the sake of argument, that we lived in a world where there was strict constant returns to scale—down to the level of a person or even a grain of sand, or even, if that were necessary, an atom. The standard reasons for macroeconomic failure would continue to hold. Money might be non-neutral, there could be a genuine role for government as a Pareto-improver of social welfare, and so forth. But it is truly difficult to imagine involuntary unemployment with strict constant returns to scale in all aspects of technology, including borrowing and lending of capital. Note that the claim is that some form of increasing returns is a necessary, but by no means sufficient, condition for genuine involuntary unemployment.

Under strict constant returns, the macroeconomic inefficiencies would show themselves in the form of 'wrong' labour-leisure choices (or, more generally, wrong substitution choices among various factors and commodities). An economy of blueberry pickers, mushroom gatherers, clam diggers and the like cannot exhibit involuntary unemployment no matter what else is present or absent. It can show fluctuations, inefficiencies, poverty, even starvation, but it cannot show involuntary unemployment. Any 'involuntarily unemployed' resource would merely form itself into a mini-firm, hire (with non-increasing returns to scale in borrowing) a few grains of cooperating input and sell its mini-output on competitive markets. Balanced expansion would take care of the rest. Involuntary unemployment is logically impossible in a strict constant returns to scale world of one-person firms.

Now it seems to me that the really damaging macroeconomic inefficiencies of advanced capitalist countries are caused by involuntary unemployment, and not by the wrong labour-leisure choice. I doubt that Keynes would ever have written his book if he had lived in a strict constant-returns-to-scale world because there probably never would have been a Great Depression. It seems to me that if we could magically turn our economy into a constant-returns system — if the automobile worker laid off from a 1,000-man plant could produce in his home workshop one-thousandth

of what that plant produces (by using in his home workshop one-thousandth of its capital) – we would have eliminated the lion's share of macroeconomic losses due to coordination failures. Other coordination inefficiencies admittedly might remain, but my casual empiricism tells me they would be orders of magnitude smaller in terms of welfare losses.

So I see increasing returns and imperfect competition as not just another minor detail, but as crucial aspects of the Keynesian story. That story simply cannot be told credibly or completely without something like increasing returns blocking unemployed labourers from working on their own or in small groups. It was to focus as sharply as possible on the underlying 'real' role of increasing returns and imperfect competition that I attempted (perhaps unsuccessfully) to trim away as much as possible of all else from the model. (Certainly there is a crucial place, in any complete story, for expectations, money, sticky wage contours and so forth – my basic point is that increasing returns constitutes a necessary, but by no means sufficient, condition for the existence of involuntary unemployment.) Furthermore, as I tried to show in the paper, the quantity-adjustment mechanisms which play such an important role in the operational part of Keynesian theory can be grounded more solidly in imperfect competition than in perfect competition, where they really do represent an artificial intrusion.

In most reasonable models of an economy with non-trivial increasing returns to scale, there is going to be a theorem showing that higher levels of equilibrium employment are associated with higher real wages. This aspect comes out quite clearly in the model under discussion and I believe it obtains under fairly general circumstances. The existence of economies of scale will generally mean that higher levels of long-run economic activity go together with higher real pay. This has at least two important implications, one for the long run and the other for the short run.

On the long-run side, consider the implications of an upward-sloping supply of labour schedule for the monopolistically competitive sector. The upward-sloping supply might come from the idea that labour not employed in the advanced increasing-returns sector can employ itself with diminishing returns in various small-scale self-contained activities. The combination of an upward-sloping long-run supply of labour schedule with a condition relating higher levels of equilibrium employment to higher real wages means that

it is easy to generate multiple equilibria. The situation is analogous to combining upward-sloping supply and demand curves. There can easily be multiple intersections. Between each pair of stable equilibria will be an unstable equilibrium. The stable equilibria can be ranked, so that some of them are inferior to others. All this is roughly consistent with a Keynesian world view and provides some support for the idea of government policy to move an economy from a bad to a good equilibrium.

On the short-run side, pro-cyclical real wages is a very unclassical feature which not only corresponds empirically to what we frequently observe in the real world but makes it theoretically difficult to accept the idea that the economic system can automatically, and relatively easily, adjust itself toward full employment. After all, the classical argument is that unemployment will be eliminated by downward pressure on wages. Arguing where burdens of proof lie is always tricky, but it seems to me that the burden of proof here might rest on whomever would assert that downward-pressed (money) wages spontaneously cause the increased real wages that accompany higher employment. For that to happen, prices must decline even faster than wages. It is of course possible, but some good stories have to be told. Here is yet another indication that economies of scale form a natural backdrop for Keynesian macroeconomics.

As I perhaps did not sufficiently stress in my original paper, but have been at pains to emphasize here, it is not increasing returns alone that causes involuntary unemployment. Other ingredients like money, expectations of other firm's responses, reasonable specifications of tastes, a sticky wage contour of equal pay for equal work, and so forth are also needed to give a credible account of involuntary unemployment. Yet, I would maintain it is important not to lose sight of the forest for the trees. Large-scale division of labour makes for an economic environment in which, contrasting with constant returns, it really is quite fundamentally difficult to tell reasonable stories about how a market economy naturally adjusts to create full employment.

While writing this paper on the foundations of unemployment theory in the summer of 1981, I was once again struck by how the seemingly innocuous assumption of a wage system actually carries a lot of implications. The institutional detail of a wage system seemed to matter. The basic argument about how an economy can

get stuck in a low level equilibrium trap reflecting a failure of coordination was quite dependent upon the wage system of paying labour. I vaguely made a connection with the depression story about the unemployed shoe factory worker being hireable but for the fact that he would not take payment in shoes. This story always seemed to trail off into oblivion at some point and the basic strand of argument was never fully developed. At the time I did not attach any great significance to the observation. Some time way back, perhaps in student days, it had seemed peculiar how, when you got right down to it, so much of standard marginal productivity theory seemed to depend on the institution of a wage contract. Labour was supposed to be hired to the point where the marginal revenue product of an extra worker was equal to the wage rate. But what happened if labour was paid entirely in shares of revenues or profits instead of wages? Then the maximand changed. There were no wages as such. Was more labour then hired, to the point where the marginal revenue product of labour was zero? This question was filed away as one more economic puzzle that had no practical implications and presumably could be resolved if one took the time to think about the problem the right way.

○ The basic ideas behind the share economy came in a sudden flash of Eureka-style inspiration. It happened at the end of January 1982. I can place the date fairly accurately because I remember it was right around the time of the one-hundredth anniversary of Franklin D. Roosevelt's birth. I recall thinking through the ideas while watching a documentary television film on the New Deal.

○ The beginning of 1982 was a very inauspicious time for the economy. Unemployment was soaring to heights that were unmatched since the Great Depression. A severe tightening had been instituted by monetary authorities who could see no other way to break the back of inflation. I well recall that several colleagues who later came grudgingly to praise Paul Volcker at that time criticized him bitterly for raising interest rates so high and causing so much unemployment. People were scared because no end to the recession was in sight. Economists were widely condemned for not coming up with better solutions to the then seemingly intractable problem of stagflation.

○ The automobile industry was hit very hard by this most serious recession since the Great Depression. It was quite clear that there were going to be even more layoffs and plant shutdowns. The

United Automobile Workers (UAW) union proposed to General Motors (GM) that they, the UAW, would moderate their wage demands provided GM kept down the prices of its autos; but if the prices of new models went up the UAW would be their usual aggressive selves in asking for higher wages. This proposal was floated as a trial balloon in late January of 1982. Probably the union was motivated by a combination of trying to appear generally helpful in the fight against stagflation and attempting to reach out specifically in a novel face-saving way to preserve jobs. Nothing much came of the idea, and the story itself was meagrely reported as a little-noted article buried in the business pages.

As soon as I saw this rather obscure news story, a flash went off in my mind. If workers agreed to contracts that in effect tied their pay to the prices of what they were producing, that could help with stagflation because it was just like product wages. While from the beginning I never thought this kind of idea represented a magic bullet, I did immediately become quite excited by the prospects. But first it was off to the drawing board to work out the details.

A bunch of questions immediately suggested themselves. It seemed intuitively clear that product wages, revenue sharing, profit sharing, cost sharing etc. were generically related to each other – but just what was the connection? And then I recollected that Japanese workers were paid by some kind of unusual system that somehow resembled profit sharing – could that have anything to do with why Japan seemed always to be able to maintain such high rates of employment? Most importantly, what were the theoretical properties of these kinds of unorthodox payment systems and how did they differ from wage systems?

When I tried to plug in product wages or profit sharing into the general equilibrium monopolistic competition model previously developed, some paradoxical results came out. Suppose a wage economy was in 'unemployment equilibrium'. Workers are unemployed because there is not adequate demand for what they produce and there is not adequate demand for what they produce because other unemployed workers lack purchasing power. Now substitute an identical-pay level of profit or revenue sharing for wages. The effect is equivalent to simultaneously lowering the base wage paid by a firm and transfer-taxing its pure profits so that each employed worker in the previous equilibrium configuration is paid the same

as before. But now the firm will want to expand production and employment, because the marginal revenue product of labour exceeds its marginal cost. If every firm in the economy converts over to the new system, a balanced expansion occurs out of the unemployment equilibrium trap with final pay actually at higher levels than before.

But now comes the paradox. Where does the expansion end? Obviously when the economy runs out of labour. But then how can there be an equilibrium? The firms want to hire more labour, but no more is available. Would there not then be a bidding up of profit-sharing coefficients? To what level? Intuitively there ought to be some kind of isomorphism with a full employment wage system, but a sharing system looks as if it would then be displaying an excess demand for labour. To what kind of macroeconomics does this correspond? What is going on here anyway? Things had reached that supremely exciting point where the theory I had created was taking on a life of its own, becoming strong enough to talk back to me while I, its creator, was not quite sure what the artificial creature was trying to say. Of course, later everything became obvious. And, after all, this is economics, so that no amount of theorizing is a match for the complications of the real world. But at the time I felt the unmistakable thrill of discovery. I really felt I was on to something.

It seemed fairly obvious that in a classical, deterministic, frictionless equilibrium, profit-sharing and wage economies should be isomorphic. Workers and capitalists should end up with the same remuneration under both arrangements. But how was this compatible with there being a seemingly eternal desire for expansion under profit-sharing? The essence, I reckoned, had to do with the difference between a Hicksian short run where pay parameters were essentially quasi-fixed, for whatever reason, and a Marshallian long run where they were subject to basic market forces. In the short run it could make sense to talk about an excess demand for labour even if the concept was essentially meaningless for the long run. Equilibrium, I reminded myself, does not mean demand equals supply. After all, the monopolist in equilibrium always wants to sell more at the price he has chosen than his customers want to buy, but that does not mean he wishes to lower his price.

In my view the concept of 'excess demand for labour' under a share economy is essentially a heuristic device that may be useful

as a way of thinking about what is happening in a share economy in a short run when pay parameters are quasi-fixed relative to everything else. But nothing substantive depends on using this phrase or even understanding what it means. The properly fastidious way to think about the matter is, I think, as follows. A short-run equilibrium is defined for a situation where pay parameters are quasi-fixed and every other variable in the system, including labour, can be freely changed. (This is obviously an extreme situation but it serves to make a basic point that hopefully holds for less extreme situations.) Then a long-run equilibrium, where pay parameters are also free to vary, is defined. The basic result is that small changes in the neighbourhood of a long-run equilibrium may produce short-run unemployment in a quasi-fixed wage system but not in a quasi-fixed share system.

Because the concept of 'excess demand for labour' is controversial, let me rephrase the essential argument without ever making use of that phrase. Suppose there are two kingdoms, Old Lakeland and New Lakeland, which are physically identical in every way. The economies of both identical-twin kingdoms consist exclusively of fishing from the numerous privately owned lakes and exporting all the fish at given world prices. For each lake, fishing production functions are stable and known.

In Old Lakeland, the monarch has decreed that the money wages to be paid throughout the year at each lake are to be posted on 1 January of that year and cannot be altered until 1 January of the next year. In New Lakeland, the monarch has decreed that payment of each lake shall consist of a share of the value of the fish caught per worker; the share fraction applying throughout the year is to be posted on 1 January of that year and cannot be altered until 1 January of the next year. In both economies, once the pay parameters (wages or share fractions) are posted, workers are free to migrate to that highest-paying lake which will employ them.

Suppose that the world price of fish has been steady for as long as anyone cares to remember. Then Old Lakeland and New Lakeland will settle into a (long-run) competitive equilibrium that is exactly identical in every respect except that pay is called 'wages' in Old Lakeland and 'shares' in New Lakeland.

Suppose next that, suddenly and without warning, in the middle of one year the world price of fish drops. By royal decree, pay parameters cannot be changed to reflect the new situation until

next 1 January. What happens in this (short-run) disequilibrium? Lake owners in Old Lakeland will choose to lay off workers, so that Old Lakeland exhibits unemployment. But at the same time New Lakeland remains at full employment.

This basic parable can be amended in a number of ways without destroying its essential message. A share economy will have a tendency to remain at full employment after contractionary shocks, because share employers want to retain workers, while a wage economy will probably exhibit unemployment during a recession because it is then profitable for wage firms to shed labour.

Let me turn to the issue of how a share economy might affect the NAIRU. In the highly idealized frictionless world of a perfectly competitive labour market with perfect information, long-run equilibrium is the same under wage and share systems. In an idealized long run, Old Lakeland and New Lakeland are isomorphic and both have zero rates of unemployment. But what about somewhat more realistic situations. Is the 'share natural rate' of unemployment lower than the 'wage natural rate'? The formal analysis of unemployment comparisons between Old Lakeland and New Lakeland that I have described has been based on short-run disequilibrium considerations, when pay parameters are quasi-fixed. But might widespread sharing also lower the natural rate under a more realistic concept of long-run equilibrium than was treated in the Lakeland example?

The answer is: Yes, it presumably would. Furthermore, the short-run and long-run unemployment problems are probably related.

In order to talk meaningfully about the effects of profit-sharing on the natural rate of unemployment, one has first to have some idea about what is causing a positive natural rate in the first place. There are several theories. Some are more persuasive than others, and they are not mutually exclusive.

A leading theory contends that long-term unemployment is largely inertial or hysteresis-like. Whatever initial disequilibrium caused the increased unemployment in the first place, once unemployment continues long enough it almost gets built into the system, perhaps because the long-term unemployed outsiders cannot or do not act effectively as a disciplining force in wage setting, perhaps because working skills or desires atrophy without work, perhaps because the plight of the unemployed eventually gets

forgotten by the electorate, perhaps for other reasons. In this view the rate of change of unemployment typically has a more powerful effect on wage settlements than the absolute level of unemployment.

If this kind of inertial effect lies behind the too-high natural rate, then presumably widespread profit-sharing would lower or eliminate it. The long-term unemployment would have difficulty developing in the first place out of an initial contractionary shock because profit-sharing firms are more reluctant to let go of workers. Taking as given this kind of natural rate unemployment, leaving aside how it got started in the past, the ingrained expansionary bias of a profit-sharing system should act as a built-in counterforce to help absorb the unemployed. The absorption process could of course be speeded by traditional expansionary macroeconomic policies which, under profit-sharing, presumably pose less danger of causing prices to accelerate because the employment-inflation trade-off has been improved. So any way you look at it, profit-sharing looks as if it ought to help diminish long-term inertial unemployment.

Another theory of why the natural rate is so high is that labour has too much bargaining power. Whether a switch from a wage system to profit-sharing would lower this kind of NAIRU depends on what it is that labour and management bargain over. If they bargain over pay parameters but management controls the employment decision, a switch to profit-sharing would lower the NAIRU. If labour and management bargain over both pay parameters and employment levels, the NAIRU would be the same under either system. In-between bargaining would yield in-between results, with the NAIRU then being somewhat lower under profit-sharing than under a wage system.

A third class of theories, based on the so-called 'efficiency wage hypothesis', holds that long-term unemployment is caused by companies themselves choosing to pay above market-clearing wages to discourage bad worker behaviour like shirking or quitting. Within this kind of model the equilibrium natural rate would in principle be the same under a wage or a profit-sharing system, although short-term disequilibrium dynamics could differ.

To the extent that too-high unemployment in some economies is aided by 'overly generous' unemployment and welfare benefits, which creates some voluntary unemployment, presumably the lab-

our payment mechanism *per se* makes little or no difference. So 'the revenge of the welfare state' kind of unemployment should not be affected by a switch to profit-sharing.

Finally, there is the long-standing identification of the 'natural rate' with semi-permanent frictional or structural unemployment, due to continuously occurring microeconomic changes. This kind of unemployment, it is usually said, cannot be reduced by pure macroeconomic policies except temporarily and at the cost of increasing inflation. As with inertial unemployment, however, the wage system is heavily implicated in frictional or structural concepts of the NAIRU. After all, both wage and profit-sharing systems respond to shifts in relative demands by sending a signal that eventually transfers workers out of a losing firm or sector and over to a winner. With a wage system the signal to workers that their firm is a loser in the game of capitalist roulette and that it is time to look for a new job with a winning firm is the boot – the worker is laid off and must suffer through an unemployment spell of some duration while searching for the new job. Under a profit-sharing system, the firm does not voluntarily let go of a worker because of weak demand. Instead it is the worker who chooses to leave because pay is too low relative to what is available elsewhere at relatively more successful firms.

Summing up, in none of the standard scenarios does a profit-sharing system cause a higher NAIRU than a wage system, and in most of the more reasonable descriptions a profit-sharing system generates a lower NAIRU than a wage system. In addition, the profit-sharing system has better disequilibrium properties when pay parameters are sticky in the neighbourhood of the NAIRU unemployment rate.

From all of these theoretical exercises considered together it seems difficult not to draw the conclusion that a profit-sharing economy is more likely to have lower unemployment than a wage economy. Yet these are theoretical exercises. What, realistically, is the possible role that profit-sharing might play in improving economic performance?

When it comes down to it, when all is said and done, I guess I think the moral of the story is that profit-sharing is likely to be a 'good thing'. The formal theory gives a veneer of respectability, a flavour of overall coherence, to a general idea that is easy enough to understand on a common-sense level. Profit-sharing can help

the unemployment–inflation trade-off because the built-in profit-sharing cushion makes it less cost-saving for the employer to lay off workers during bad times.

What causes unemployment or slack labour markets? There is one basic answer but, like a coin, the answer has two sides. Side one is that unemployment is caused when firms face insufficient demand for their products relative to their marginal costs of production. Side two is that unemployment is caused when firms have too-high marginal costs of production relative to the demand for their products. Sometimes it is useful to stress one side of the coin, sometimes the other. But it is essentially the same coin.

In either case, the key to non-inflationary full employment is an economic expansion that holds down the marginal cost to the firm of acquiring more labour. Macroeconomic policy alone – the purposeful manipulation of financial aggregates – can be very powerful in achieving full employment or price stability, but cannot be reliably depended upon to reconcile both simultaneously. Why? Because of the two-headed monster – stagflation. Illusions of being able to fine tune aside, we know how to get unemployment down and output up by the usual expansionary monetary and fiscal measures. We also know how to break inflation by policy-induced recessions. What we do not know – and this is perhaps the central economic dilemma of our time – is how simultaneously and reliably to reconcile reasonably full employment with reasonable price stability. Expansionary policies often dissipate themselves, to an excessive degree, in too-large wage and price increases rather than expanded employment and output.

When one thinks seriously about it, the wage system of paying labour does not seem like a very good idea. We try to award every employed worker a predetermined piece of the income pie before it is out of the oven, before the size of the pie is even known. Our ‘social contract’ promises workers a fixed wage independent of the health of their company, while the company chooses the employment level. This stabilizes the money income of whomever is hired, but only at the considerable cost of loading unemployment on low-seniority workers and inflation on everybody – a socially inferior risk-sharing arrangement that both diminishes and makes more variable the real income of workers as a whole. An inflexible money wage system throws the entire burden of economic adjustment on employment and the price level. Then macroeconomic policy is

called upon to do what is often very difficult – reconcile full employment with low inflation.

Any economy is full of uncertainty. There are no absolute guarantees, and if the uncertainty does not come out in one place it will show up in another. I am saying that it is much better, much healthier, if everyone shares just a little bit of that uncertainty right at the beginning rather than letting it all fall on an unfortunate minority of unemployed workers who are drafted to serve as unpaid soldiers in the war against inflation. It is much fairer if people will agree that only 80 per cent of their pay is going to be tied directly to the funny looking green pieces of paper – which are themselves an illusion, although a very useful illusion – and 20 per cent will be tied to company profits per employee. Then the economy can be more easily controlled to have full employment and stable prices. Society will be producing, and hence consuming, closer to its full potential. If people will face up to the uncertainty, and if everyone accepts some small part of it, then society as a whole can end up with higher income and less uncertainty overall.

A profit-sharing system, where some part of a worker's pay is tied to the firm's profitability per employee, puts in place the right incentives to resist unemployment and inflation. If workers allow some part of their pay to be more flexible by sharing profits with their company, that could improve macroeconomic performance by directly attacking the economy's central structural rigidity. The superiority of a profit-sharing system is that it has enough built-in flexibility to maintain full employment even when the economy is out of balance from some shock to the system. When part of a worker's pay is a share of profits, the company has an automatic inducement to take on more employees in good times and, what is probably more significant, to lay off fewer workers during bad times. A profit-sharing system is not anti-labour and does not rely for its beneficial effects on lowering workers' pay. The key thing is not to get total worker pay down – it could even go up within reason – but to lower the base wage component relative to the profit-sharing component. The marginal cost of labour is approximately the base wage, more or less independent of the profit-sharing component.

While it is possible to dream up unlikely counterexamples and to interpret the existing evidence perversely, the bulk of economic theory, empirical evidence and common sense argue that wide-

spread profit sharing will help to improve macroeconomic performance. The bottom line is that it is easy to envision situations where profit-sharing helps macroeconomic performance while it is difficult to imagine scenarios where profit-sharing damages an economy, which is as much as can be claimed for any economic idea.

The British Chancellor of the Exchequer stated the case for profit-sharing as follows in his 1986 annual budget speech before the House of Commons:

The problem we face in this country is not just the level of pay in relation to productivity, but also the rigidity of the pay system. If the only element of flexibility is in the numbers of people employed, then redundancies are inevitably more likely to occur. One way out of this might be to move to a system in which a significant proportion of an employee's remuneration depends directly on the company's profitability per person employed. This would not only give the workforce a more direct personal interest in their company's success, as existing employee share schemes do. It would also mean that, when business is slack, companies would be under less pressure to lay men off; and by the same token they would in general be keener to take them on.

It is no mystery why profit-sharing makes the employer view things fundamentally differently. In a profit-sharing system the young school graduate looking for work comes with an implicit message to the employer saying: 'Hire me. I am reasonable. Your only absolute commitment is to pay me the base wage. That is my marginal cost to you. The profit-sharing bonus is like a variable cost, depending to some extent on how well the company is doing. So you have a built-in cushion or shock absorber if something should go wrong. You won't be under such pressure to lay off me or other workers during downswings.' By contrast, the young school graduate looking for work in a wage system now comes to a potential employer with the implicit message: 'Think very carefully before you hire me. I am expensive and inflexible. You will have to pay me a fixed wage independent of whether your company is doing well or poorly.' Is it difficult to deduce in which situation companies might be expected more eagerly to recruit new hires and to retain them, and in which situation new hiring commitments

are likely to be avoided when possible? The essence of the case for profit-sharing is the basic idea that on the margin the profit-sharing firm is more willing than the wage firm to hire new workers during good times and, more importantly, to lay off fewer workers during bad times. From a social point of view, a wage system is poorly designed because it is inherently so rigid. There has to be a precise relation between the wage level and the level of aggregate demand to hit the full employment target exactly without causing inflation. By contrast, a profit-sharing system is inherently much more forgiving. Full employment can be maintained even if base wages and profit-sharing parameters are somewhat 'too high' relative to aggregate demand or, equivalently, aggregate demand is 'too low' relative to pay parameters.

Without getting into the more controversial issue of whether profit-sharing causes increased hiring and excess demand for labour, even a one-sided worst-case scenario where profit-sharing 'merely' dampens economic downturns by encouraging employers to lay off fewer workers during recessions still represents a potentially large economic benefit. In periods of recession and other kinds of squeeze, the 'insiders' risk becoming 'outsiders' and they may well be glad of a system which, without painful renegotiations, will enable an automatic adjustment in pay to be made – which would be self-reversing in recovery – to preserve jobs. Also, even in periods of normal growth there will always be firms under pressure to reduce employment, and anything which lessens that pressure will help overall employment. To ratchet an economy toward a tight labour market and improve the employment–inflation trade-off so that macroeconomic policies can be used more effectively requires only that, on the margin, during downswings a few less old workers are laid off and during upswings no fewer new workers are hired.

The basic story being told is subject to a host of potential criticisms and objections. The model is incomplete and *ad hoc* at certain crucial points. Nevertheless, I think the basic message rings true. Sharing arrangements generally make it easier to consummate and less desirable to dissolve an economic union of those who are hired with those who do the hiring.

When someone asks of an ordinary wage system 'why is company X hiring so many workers rather than fewer or more?', the first instinctive response of the economist is to say that labour is being

hired to the point where its marginal revenue product to the company is equal to its marginal cost. We economists know there are many caveats to be made and that a large number of models can be constructed that weaken or even negate this simple classical answer. Still, in the end, most of us come back to the marginal revenue product equals marginal cost explanation as the basic 'big picture' employment story we keep in the back of our minds.

The profit-sharing story that I am trying to tell attempts to push through to its logical conclusion the simple standard classical marginalist paradigm under conditions when there is an alternative 'sharing' payment mechanism. It then turns out that the resulting share system adheres closer to full employment than a corresponding wage system as various underlying parameters (including pay parameters) are perturbed around their long-run equilibrium values. If we try to tell the simple Econ-10 employment story with a money wage replaced by an equivalent amount of profit-sharing, the resulting macroeconomy has better employment-stabilizing properties because the marginal cost of labour is lower. The standard IS-LM story can be told 'as if' the underlying short-run cost of labour is the base wage.

The idea that profit-sharing is generally a 'good idea' has obviously caught on to some degree. I think there must be several reasons for this. People seem to be much more aware that the well-being of a company's workers depends, ultimately, on the well-being of the company itself, and that it is probably healthier to recognize this mutual dependence explicitly by some form of profit-sharing than to go on pretending that it does not exist. There is much speculation, and some evidence, that profit-sharing leads to improved productivity because workers feel that they have more of a stake in the outcome.² (At the very minimum, workers in a profit-sharing company seem to take an interest in company profits, whose existence and determinants they are often only dimly aware of otherwise.) Finally, there is the sense that macroeconomic performance may be improved under profit-sharing. Not, of course, exactly in the mechanical highly formal description of the models. Reality is more than a match for any model, at least in economics. My hope is that the overly crisp models of profit-sharing are at least pointing in the right direction of an internally consistent economy with better macroeconomic properties than the more conventional wage system.

NOTES

- 1 See, for example, Cooper and John (1988), Kiyotaki (1988), Roberts (1987), and the further references cited therein.
- 2 See Weitzman and Kruse (1989) for a broad survey of the evidence.

REFERENCES

- Cooper, R. and John, A. (1988): 'Coordinating Coordination Failures in Keynesian Models', *Quarterly Journal of Economics*, 103, 441-63.
- Diamond, P. (1982): 'Aggregate Demand Management in Search Equilibrium', *Journal of Political Economy*, 90, 881-94.
- Kiyotaki, N. (1988): 'Multiple Expectational Equilibria Under Monopolistic Competition', *Quarterly Journal of Economics*, 103, 881-94.
- Roberts, J. (1987): 'An Equilibrium Model with Involuntary Unemployment at Flexible, Competitive Prices and Wages', *American Economic Review*, 77, 856-74.
- Weitzman, M.L. (1982): 'Increasing Returns and the Foundations of Unemployment Theory', *Economic Journal*, 92, 787-804.
- Weitzman, M.L. (1983): 'Some Macroeconomic Implications of Alternative Compensation Systems', *Economic Journal*, 93, 763-83.
- Weitzman, M.L. (1984): *The Share Economy* (Cambridge, MA: Harvard University Press).
- Weitzman, M.L. (1985): 'The Simple Macroeconomics of Profit Sharing', *American Economic Review*, 75, 937-53.
- Weitzman, M.L. and Kruse, D. (1989): 'Profit Sharing and Productivity', in A. Blinder (ed.), *Worker Compensation and Productivity* (Washington, DC: Brookings Institution).

Index

- Acworth, William 40, 53
Adams, Henry Carter 93, 94, 95, 102
Addis, Sir Charles 150, 161
*ad hoc*ery, expectations and 253–5
agents: in expectation theory 233;
 infinitely lived 242, 244–5; Ramsey
 249, 254
Allen–Antonelli integrability conditions
 284
Allen, Roy 208
all-or-none offer functions 269–70, 271,
 272, 275; derivation of 291–2
American Economic Association (AEA)
 7, 9, 27, 34, 93, 172, 175
American economics profession 9,
 92–105
American Federal Reserve Board's
 multicountry model 222–3
Anderson, G.J. 302
Ando-Modigliani equation 324
Argyll, Duke of 32, 51–2, 63 n10
Armstrong, W.E. 149
Arrow, K.J. 243, 257, 282–3
Arrow–Debreu model 236, 242–3
Ashley, William 14, 54–5, 63 n29, 143;
 on trusts 96–7
Association of University Teachers of
 Economics (AUTE) 19, 186
automobile industry 342–3
autonomy *see* invariance
Axelrod, Robert 218, 225

Baird, Joyce 173, 175
Balducci, R. 221

Balfour, Arthur J. 32, 163
Ball, Sidney 83–4
Bank of England model 324, 325, 327
Banzhaf–Coleman index 207
bargaining 213–16; in cooperative
 games, Nash solution 204, 222, 227
bargaining power of labour 347
bargaining set 207
Barone, Enrico 266
Basar, Tamer 223
Bastable, C.F. 41, 112
Bayesian learning 254
Bean, C.R. 308
Bergson, Abram 264, 271, 273, 282, 286
Bernstein, Eduard 77
Bertrand, J.L.F. 210
Beveridge, William 16, 166, 170
Bickerdike, C.F. 56–7, 117, 149–50, 151
bimetallism 39
Binmore, K. 224
Birmingham Airport 219
Böhm-Bawerk, Eugen von 76, 112
Boldrin, M. 245
Bonar, James 6, 32, 41, 78–9, 112, 161,
 166, 168–9, 188
Bortkiewicz, Ladislaus von 123–4, 126,
 128–9
Bover, O. 305, 306
Bowley, Arthur 113, 118, 210
Bowley's conjectural variation 211
Brand, Lord 173
Bray, Margaret M. 240, 256
British Association for Advancement of
 Science, Section F 7–8, 12–13, 29, 31